DSP Dimensioning of a 1000Base-T Gigabit Ethernet PHY ASIC

Pedro Reviriego, Carl Murray
Pedro.Reviriego@massana.com

Massana Technologies,
Arturo Soria 336, Madrid, Spain.

Abstract

This paper describes the strategy followed for dimensioning the coefficient width of the adaptive filters (FFE, echo and NEXT Cancellers) in a Gigabit Ethernet PHY Silicon implementation. These filters account for an important part of the circuit area and power consumption. It is therefore critical to ensure that they are properly dimensioned.

1. Introduction

Gigabit Ethernet over copper or, more precisely, 1000BASE-T is a part of the IEEE 802.3 standard for the Ethernet protocol [1]-[2]. This standard was put in place to provide a common method for the exchange of data within local area networks (LANs) and over the past 18 years has slowly been extended to provide increasingly faster data rates [3]. These rates have been made possible by the improvements in VLSI and DSP techniques. 1000Base-T is the most recent in a series of improvements in the Ethernet standard and as such presents some of the more complex system and VLSI challenges ever required by a baseband communication system.

In this paper, we describe the approach followed to determine the appropriate coefficient width of the adaptive filters used in the Gigabit PHY. This task is critical to ensure an efficient implementation in silicon as these filters account for a significant part of the circuit area and power. We first give a brief overview of the DSP architecture of the Massana Gigabit Ethernet PHY [4] and then in subsequent sections we describe how the most significant bits (MSBs) and the least significant bits (LSBs) for the different filters were determined. The intention of the paper is to give an overview of the approach, therefore some details are omitted for clarity. Also some dynamic effects and filter interactions are not discussed as they are beyond the scope of this paper.

2. Overview of Massana Gigabit Ethernet PHY DSP Architecture

Figure 1 shows the block diagram of the Massana Gigabit Ethernet PHY (from now on M-PHY1000) architecture. Data is simultaneously transmitted and received on each of the four UTP-5 cable pairs (signal dimensions or channels A, B, C and D) at 125 Msymbols/sec using 5 level PAM. The dashed box in Figure 1 contains the functionality required to implement the receiver for dimension A. Similar receivers are implemented for dimensions B, C and D.

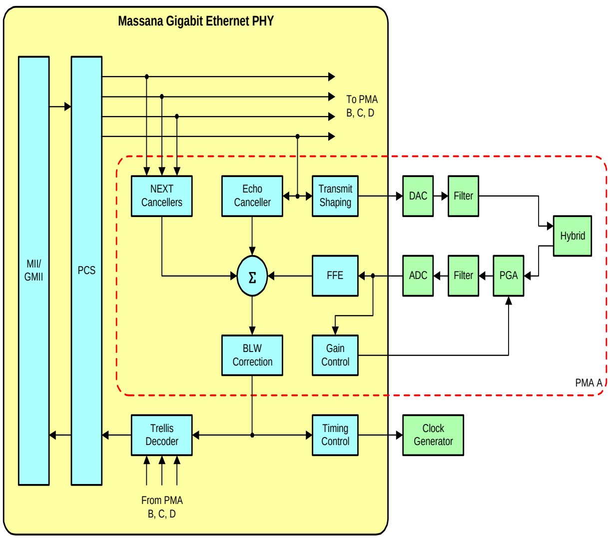

Figure 1 Block Diagram of the M-PHY1000 Architecture**.**

Because the four twisted pairs are bundled together and are not shielded from each other, each of the transmitted signals is coupled onto the other three cables and is seen at the receiver as Near End Cross-Talk (NEXT). Far End Cross-Talk (FEXT) also occurs for the same reason but due to the

remote transmitter. In order to transmit and receive simultaneously on each pair a hybrid circuit is used. If the transmitter is not perfectly matched to the line, a signal component will be reflected back as an echo. Reflections can also occur at other connectors or cable imperfections. Using the locally transmitted symbols, the M-PHY1000 cancels the echo and NEXT by subtracting an estimate of these signals from the equalizer output at the symbol rate. In the M-PHY1000 FEXT is not cancelled as there is considerable complexity involved in doing so and it is very much a secondary effect (wrt NEXT) at long cable lengths where performance is most required.
The M-PHY1000 uses an oversampled architecture, i.e. it samples at two times the symbol rate. A fractionally spaced Feed Forward Equalizer (FFE) adapts to remove Inter-Symbol Interference (ISI) and to shape the spectrum of the received signal to maximize the SNR at the trellis decoder input. The FFE equalizes the channel to a fixed target response. Note that the output of the FFE is at the symbol rate and, therefore both the FFE and the rest of the receiver operate at the symbol rate.
From figure 1 we can see that the adaptive filters for echo and NEXT cancellation and the FFE will represent an important part of the circuit. This becomes even more apparent if we know that in order to meet the BER specifications for different channels we need a significant number of taps in each of these filters that are replicated for each of the four cable dimensions. This clearly requires a significant amount of area and power (the filters operate at 125Mhz) which is influenced by many factors like the number of taps used, the coefficient range of the taps, the filter implementation , etc.. As part of our efforts to minimize the area/power needed to implement these filters, we have looked in detail at the coefficient width for each of them.

3. Overview of the approach

It is assumed for the purposes of the following discussion that the coefficients are represented with the sign magnitude format s[MSB:LSB]. The MSB determines the maximum value that a coefficient can take while the LSB determines the precision of the coefficient.
As by design each canceller tap is independent, if we set the MSB so that an echo or NEXT tap can not take its value, we will not cancel part of the echo or NEXT signal at that tap, but other coefficients in the filter will not be affected. The amount of interference not cancelled will be proportional to the value of the MSB. Therefore, we can be more aggressive on far echo taps that have a small MSB than on near echo taps that can potentially take big values. For the FFE saturation is more critical as each tap is not independent of the others and saturation in one tap may affect the overall functionality of the FFE. This suggest a more conservative approach in setting its MSB.
The LSB effect can be modeled as an addition of noise that is proportional to the value of the LSB. So a noise target can be assigned for the LSB contribution based on the system level noise budget.
In the following sections we describe in more detail the approach used for setting the MSBs/LSBs in the different filters. Many of those results rely on system level simulations of the Gigabit PHY and the UTP-5 cable. Massana has developed an extensive model of the cabling environment and the PHY using Matlab/C++. This model has been extended to facilitate our analysis (including for example max/min coefficient tracking, some special cable settings,etc.) and used to verify our assumptions.

4. Determining Most Significant Bits (MSBs)

This section focus on how to determine the MSB which determine the biggest values that the coefficients can take.

Echo and Next Cancellers

When looking at the MSBs in these filters we have to take into account that :

- Cancellation is performed after the FFE. This means that the canceller coefficients are a function of the FFE response. For a given interference at the A/D it is assumed that the greater equalisation required of the FFE (long channel lengths) the greater the interference at the output of the FFE.
- Canceller coefficients can to some extent saturate without causing performance to degrade catastrophically. This is due to the fact that each canceller tap is independent of all other taps, as by design data input to all cancellers are uncorrelated.

Next Cancellers

With the NEXT cancellers it is difficult to determine a set of cable characteristics that are guaranteed to maximise the peak of the NEXT. However it is possible to propose cable characteristics that maximise the NEXT power as measured at the output of the FFE with the following assumptions :

- The peak due to NEXT is maximum due to a single connector. Multiple connectors destructively interfere (this is an observation from extensive simulations designed to maximize NEXT).
- The closer the connectors are to the PHY under consideration the greater the interference.

- The longer the cable length, the more the net interference is enhanced by the FFE.

Taking these assumptions to be valid we have simulated a number of configurations that maximize those factors and obtained a reasonable upper bound on the NEXT interference power. The power per dimension in the signal cancelled by the NEXT cancellers may be given by the following equation.

$$\sigma_{NEXT}^2 = \sum_{i=1}^{N} 2c_{1,i}^2 + \sum_{i=1}^{N} 2c_{2,i}^2 + \sum_{i=1}^{N} 2c_{3,i}^2$$

where $c_{x,i}$ is the i[th] NEXT tap coefficient value from canceller x and 2 is the symbol power during IDLE. Based on the NEXT models, it can be seen that the interference has a large peak to RMS ratio. This makes it reasonable to assume that all the NEXT interference is concentrated in a single tap for each contributing dimension. Therefore, the magnitude of this tap can be determined as (assuming the power is equally distributed across the three contributing dimensions).

$$c_{\max} = \sqrt{\frac{\sigma_{NEXT}^2}{3*2}}$$

This upper bound was used to determine the MSB for the NEXT cancellers.

Echo cancellers

For echo cancellation the peak may be more accurately determined than for NEXT as the dominant components contributing to the echo can be more easily determined.

There are two significant components (when considering MSB selection) to the echo

- Echo at the connectors
- Echo due to the hybrid

The hybrid will only have an effect on the dimensioning of the first echo canceller taps. For the other taps we need only consider the effect of reflections due to cable mismatched boundaries at the connectors.

In the case of the connectors, each will contribute an echo component. In most of the situations, as a first approximation, the total echo can be estimated as the sum of the echoes at each of the connectors. These echoes can constructively/destructively interfere depending on the relative position of the connectors and on the sign of the impedance mismatches at each of them.

The peak value may be bounded considering the following arguments :

- The remote hybrid may be viewed as an impedance mismatch boundary.
- The reflection coefficient for an impedance mismatch is given by $\Gamma = \frac{z_1 - z_2}{z_1 + z_2}$.
- The reflections from successive connector reflection boundaries must be considered in two ways.
    - impedance mismatches are monotonic
    - reflection coefficient alternates in sign
- We can assume a limit on each individual impedance mismatch boundary 100/120Ω.

Consider first monotonic impedance mismatches (i.e. Γ of the same sign). The maximum peak is bounded by the sum of the echos occurring at all boundaries, which in turn is related to the sum of the reflection coefficients at each boundary i.e.

$$\begin{aligned}
\Gamma_1+\Gamma_2+\Gamma_3+\Gamma_4+\Gamma_4 &= \frac{z_1-z_2}{z_1+z_2}+\frac{z_2-z_3}{z_2+z_3}+\frac{z_3-z_4}{z_3+z_4}+\frac{z_4-z_5}{z_4+z_5}+\frac{z_5-z_6}{z_5+z_6} \\
&\le \frac{z_1-z_2}{2z_{\min}}+\frac{z_2-z_3}{2z_{\min}}+\frac{z_3-z_4}{2z_{\min}}+\frac{z_4-z_5}{2z_{\min}}+\frac{z_5-z_6}{2z_{\min}} \\
&\le \frac{z_1-z_6}{2z_{\min}} \\
&\le \frac{z_{\max}-z_{\min}}{2z_{\min}}
\end{aligned}$$

This is equal to the maximum reflection coefficient at a single boundary. This latter equation tells us that if the echo at a boundary is due to a single maximum impedance mismatch then its peak value is greater than the sum of the peak values due to the impedance mismatches at the 5 possible connector boundaries – i.e. the peak echo due to multiple boundaries is bounded by the peak echo due to the maximum impedance mismatch at the first boundary.

In the following figures the maximum echo for a single connector reflection boundary at different positions along a 100m cable is plotted. The first figure shows the impulse response of the echo at the A/D. The second shows the optimum echo canceller coefficients. The differences between these figures is due to the fact that the canceller acts on the signal after the FFE and also that the peak of the echo may lie between two cancellers taps or be close to one tap. This latest effect explains the drops seen for some cancellers taps positions.

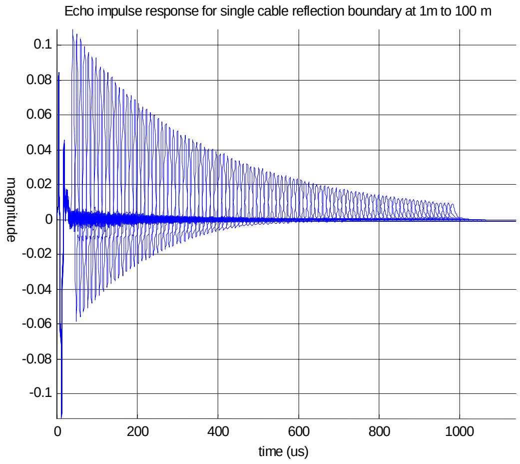

Figure 2 Echo impulse response for single cable impedance mismatch (100–120Ω) at 1 to100 m

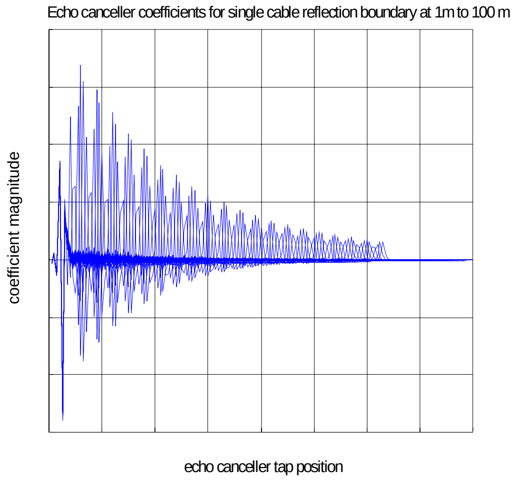

Figure 3 Echo canceller taps for single cable impedance mismatch (100 –120Ω) at 1 to 100 m

If we consider the above results only it can be seen that the canceller can be divided into banks of N coefficients and the range of each successive bank can be reduced by a factor of 2 on the previous bank.

We must now consider the effect when successive reflection boundaries alternate in sign. The following figure shows the impulse response for echo due to a single cable mismatch boundary (at ~ 150 us) plus the echo due to the local hybrid. The key feature of interest is the fact that echo due to the boundary has a single positive large peak and a single negative smaller peak. If echoes from successive boundaries alternate in sign the large peak from the second boundary may constructively interfere with the smaller peak from the first.

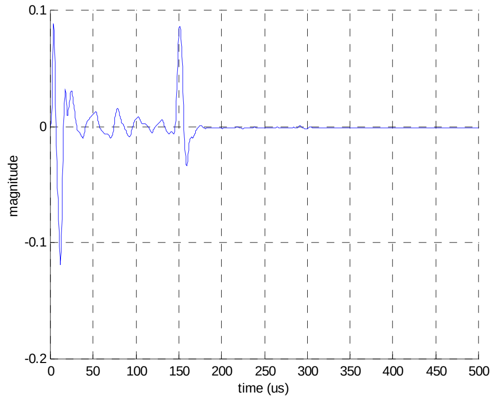

Figure 4 Echo impulse response for single impedance mismatch (100Ω–120Ω)

The next figure shows plots of the interference when a second reflective boundary with a reflection coefficient of the same magnitude but opposite polarity is added to the cabling structure considered in the previous figure. It is positioned at 0.1m and swept in 0.1m steps up to 1m from the first boundary. We can see that although they do interfere constructively the peaks of either 'lobe' are less than the peak that would happen due to the single boundary (also plotted). This is due to the fact that the constructive interference can only happen with the second (and smaller) lobe of the first echo. So that it can not be much bigger than the first lobe of the first echo.

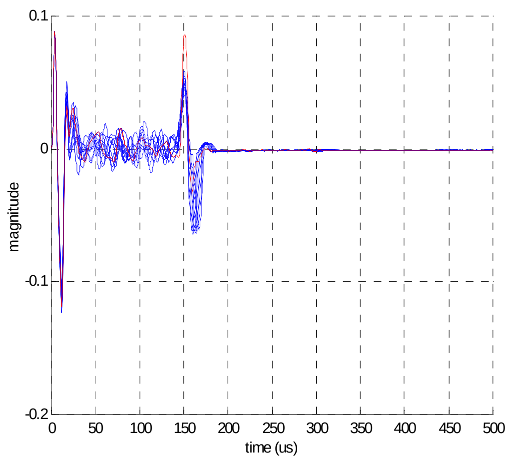

Figure 5 Illustration of interference between echos from successive reflection boundaries where reflection coefficient polarity is different in each.

The effect of the local hybrid has been considered by setting the worst case impedance match between the local hybrid and the immediate cable segment. Additionally a reflection boundary due to cable mismatched impedances is placed at positions sweeping from 0.1m in steps of 0.1m to 14m. The following plot shows the resultant echo canceller coefficients. This information has been used to set the MSB for the first echo coefficients.

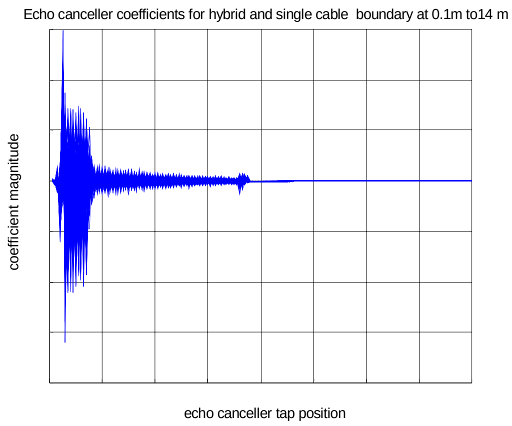

Figure 6 Echo canceller coefficients due to echo at local hybrid

Feed Forward Equalizer (FFE)

The main purpose of the FFE as it name suggests is to equalize the channel. In our case the UTP-5 cable. The first bound that we can use to set its MSB can be derived from (see section 2.7 in [5]):

$$|h(n)| = \left| \frac{1}{2\pi} \int_0^{2\pi} H(e^{jw}) \cdot e^{jwn} \cdot dw \right|$$

$$\leq \frac{1}{2\pi} \int_0^{2\pi} \left| H(e^{jw}) \right| . \left| e^{jwn} \right| \cdot dw \leq \left| H(e^{jw}) \right|_{\max}$$

This basically states the absolute value of any given coefficient in a filter will always be smaller than the maximum absolute value of its frequency response. So if we can determine what is the maximum frequency boost that the FFE needs to implement, we can safely determine a bound on the MSB.
The maximum frequency boost has been determined by running exhaustive simulations on a Matlab/C++ model of the Gigabit channel. As expected the bigger values appear with long cables, so we set the bound so that we can deal with the longest length of cable that the PHY is intended to operate at. As a reference the FFE response for a 180m channel is shown in figure 7.
The above discussion deals only with the steady-state values of the coefficients. During the different adaptation phases the coefficients fluctuate around those values. Again in simulation we have seen that the maximum values seen during adaptation are not very different that those of the steady state.
In the cancellers, saturation is not as critical as in the FFE, because as it was mentioned, each tap adapts independently. That is why the dynamic behavior of the cancellers has not been discussed.

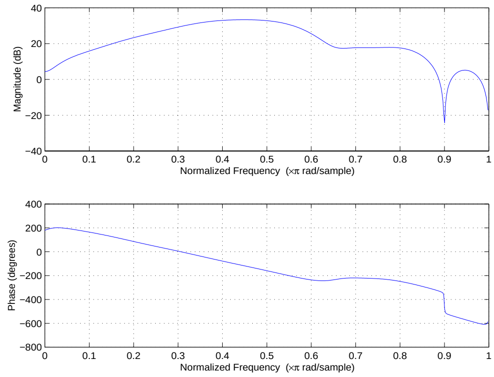

Figure 7 Frequency response of final coefficients (180m)

5. Determining Least Significant Bits (LSBs)

Echo and Next Cancellers

The quantization process for the LSB (assuming rounding) of a coefficient can be modeled as additive noise of the form:

$$\sigma_{lsb}^2 = \frac{\left(2^{lsb}\sigma_s\right)^2}{12}$$

where σ_s^2 is the power of the data in the associated filter delay line [5].
If the data in the delay line is uncorrelated and the LSB is small compared to the overall coefficient range, the noise due to N taps is given by :

$$N\sigma_{lsb}^2$$

If in the implementation of the cancellers in the design the coefficients are truncated rather than rounded when implementing the sum of products. The effect of truncation is to change the noise contribution of each coefficient to

$$\sigma_{lsb}^2 = \frac{\left(2^{lsb}\sigma_s\right)^2}{12} + 2^{2\cdot(lsb-1)}\sigma_s^2 = \frac{2^{2\cdot lsb+1}\cdot\sigma_s^2}{12}$$

Due to the bias in the coefficient quantisation. The net effect wrt the noise added is to increase the LSB by 1 bit.
This assumes that the coefficients are static. However due to the adaptation process the coefficients may be continually changing. In the best case the LSBs of the coefficients that are fed to the SOP are toggling. This would mean that in some cases the coefficient would appear to be rounded up wrt the optimum coefficient, other

times truncated. If this is the case then it is worth reconsidering the quantisation noise.
The mean quantisation error is given by :

$$a.x+(1-a).(x-q)$$

where q = 2^{lsb}, 'a' is the probability that the coefficient is truncated wrt the optimum coefficient, (1-a) is the probability that the coefficient is rounded up wrt the optimum coefficient and x is the difference between the coefficient and its truncated value.
For a given coefficient if adaptation has reached a steady state then the mean error due to the quantisation of the coefficient must be zero, i.e.

$$a.x+(1-a).(x-q)=0$$

which implies

$$a(x)=\frac{q-x}{q}$$

The quantisation noise becomes :

$$\sigma_{lsb}^{2} = \int_{0}^{q}\left[a \cdot x^{2} + (1-a)(x-q)^{2}\right] \cdot \frac{dx}{q}$$

$$= \int_{0}^{q}\left(q^{2} \cdot x - q \cdot x^{2}\right) \cdot \frac{dx}{q} = \frac{q^{2}}{6}$$

This would imply that if the coefficient is continually adapting then the effect of truncating the coefficient for purposes of implementing the filter is only a half a bit. Using this formula and knowing the number of filter coefficients and the power of the signal being filtered (which is the normal gigabit Ethernet transmit signal) we can set the LSBs if we know the amount of noise that we can allow the quantisation to introduce. The amount of noise that quantisation can introduce has been determined by estimating the total noise through extensive system simulations and then setting the limit for quantisation to be significantly smaller than that.

FFE

A similar analysis may be carried out for the FFE coefficients as for the canceller coefficients. The key differences to note is that the power of the signal in the FFE delay line is controlled by the RMS target of the PGA (The PGA continually monitors the incoming signal and adjust the gain to ensure that the input signal has always a given RMS target). Secondly the data is correlated. Therefore the theoretical analysis is less applicable. So the theoretical analysis was used as a bound and complemented with extensive Matlab/C++ system simulations to set the LSB for the FFE.

6. Conclusions

In this paper we have briefly described the approach used in the dimensioning of the coefficient widths of the adaptive filters in a 1000 BaseT silicon implementation. This is a very small part of the work facing anyone involved in gigabit Ethernet over copper designs and illustrate the sheer complexity of this challenge.

References

[1] IEEE 802.3 Standard, 1998 Edition, *1998.*
[2] IEEE Draft P802.3ab/D6.0, *1999.*
[3] Gigabit Ethernet: Migrating to High-Bandwidth LANs, Kadambi, Crayford and Kalkunte,Prentice Hall, *1998.*
[4] MA1110 Datasheet Massana Inc. 2002.
[5] Dicrete Time Signal Processing A.V. Oppenheim and R. W. Schafer Prentice Hall, *1999.*