

Chapter 4

PWM pre-emphasis

4.1. Introduction

To compensate for channel losses, transmitter pre-emphasis or receiver equalization can be applied [Farjad-Rad], [Lee], [Kudoh], [Gai], [Dally-1]. Receiver equalization typically involves several analog blocks which impose speed, accuracy and noise requirements. On the other hand, transmitter pre-emphasis allows the use of a simple receiver that only needs to sample binary values [Dally-1]. Pre-emphasis methods found in the literature are commonly based on symbol-spaced Finite Impulse Response (SSF) filtering [Farjad-Rad], [Lee], [Kudoh], [Gai], [Dally-1].

In this chapter, we describe the application of PWM pre-emphasis to the equalization of copper cables and PCBs. It provides an alternative to FIR pre-emphasis with advantages in the light of developments in CMOS scaling. In our group, PWM pre-emphasis was successfully applied to the equalization of on-chip wires [Schinkel-1], [Schinkel-2]. In this chapter, it is shown that PWM pre-emphasis provides a higher maximum loss compensation (30dB) than the commonly used 2-tap SSF filter (20dB), because its transfer function happens to fit very well to the copper channel. Only one parameter needs to be adjusted to the channel (the duty-cycle), unlike with 5-tap SSF filters, or multi-stage receiver equalizers, which can also achieve 30dB loss compensation but with higher complexity. This chapter is based on our existing publications [Schrader-4], [Schrader-5], [Schrader-6] and [Schrader-7].

Section 4.2 of this chapter explains the principle behind PWM pre-emphasis, and section 4.3 analytically compares FIR and PWM pre-emphasis in the frequency domain. Section 4.4 shows time domain simulations. In section 4.5, the design and test results of three chips are discussed. Measurements are given for cables and for PCBs. Finally, conclusions are drawn in section 4.6.

4.2. Pulse-width modulation pre-emphasis

In Fig. 1(a), the output voltage waveform for the PWM-PE filter is shown. The output is normalized to ± 1 V. We assume that the modulation scheme is 2-level Pulse Amplitude Modulation (PAM). The PWM pulse shape resembles a Manchester-coded signal, but whereas the Manchester duty-cycle is fixed at 50%, the PWM signal instead has a tunable duty-cycle. A duty-cycle of 100% corresponds to transmission of a normal polar NRZ data signal without pre-emphasis, and 50% to transmission of a Manchester-coded data signal (maximum pre-emphasis setting). The optimum duty-cycle is somewhere in between, depending on the channel characteristics. In comparison, Fig. 1(b) shows the output of a 2-tap symbol-spaced FIR (SSF) transmitter, again normalized to ± 1 V. Both types of PE need only one adjustable parameter to fit them to the channel, making a coefficient-finding algorithm converge in a quick and straightforward manner. (For the 2-tap SSF this parameter is the ratio between the two tap weights.)




Fig. 1. Output signals for PE transmitters, $T_s = 200\text{ps}$. (T_s is defined as the symbol length.)
(a). PWM pre-emphasis. (b). 2-tap SSF pre-emphasis.

For a quick insight into PWM-PE filtering, the simulated time domain response of a 25m low-cost, low-end, standard RG-58CU cable to PWM pulses with several duty-cycles and $T_s = 200\text{ps}$ is shown in Fig. 2. This cable is used later in the experiments. The time domain cable model is explained in Chapter 2 and includes both skin-effect and dielectric loss. The sample moment t_s is shown with a triangle, and the ISI contributions are shown with circles. Note that for the duty-cycle setting of 53%, the cable output pulse becomes much narrower than the response to a plain polar NRZ pulse (100%). This reduces the ISI contributions significantly. It is seen that an optimum setting can be found at which the ISI is minimized. Second, note that the optimum duty-cycle for this particular channel is near - but not equal to - 50%. As a comparison, in Fig. 3 the response of the channel to 2-tap SSF pulses is shown. (Parameter r is described below.) It can be seen that PWM-PE is capable of narrowing the channel pulse response, similar to FIR pre-emphasis.

In practice, PWM duty-cycle d can be adapted to the channel automatically using return channel communication and a control algorithm. The need for a return channel is a disadvantage compared to receiver equalization, but it is common among all pre-emphasis approaches. A sign-sign block least mean squares (LMS) algorithm can be used as shown in [Stonick]. Such a control algorithm could also compensate for temperature and channel variations. Convergence of the LMS algorithm for the single-coefficient PWM-PE filter is more straightforward than it would be for a filter with multiple coefficients.

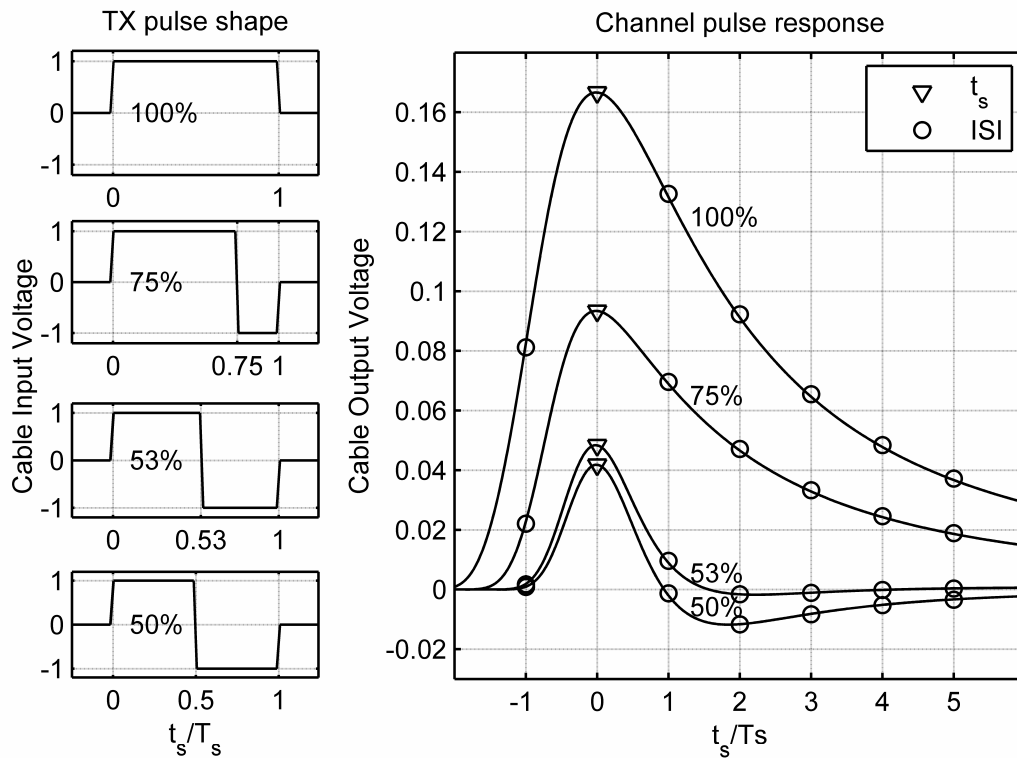


Fig. 2. TX pulse shapes ($T_s=200\text{ps}$) of PWM-PE filter with varying duty-cycles and simulated responses of 25 m RG-58CU cable.

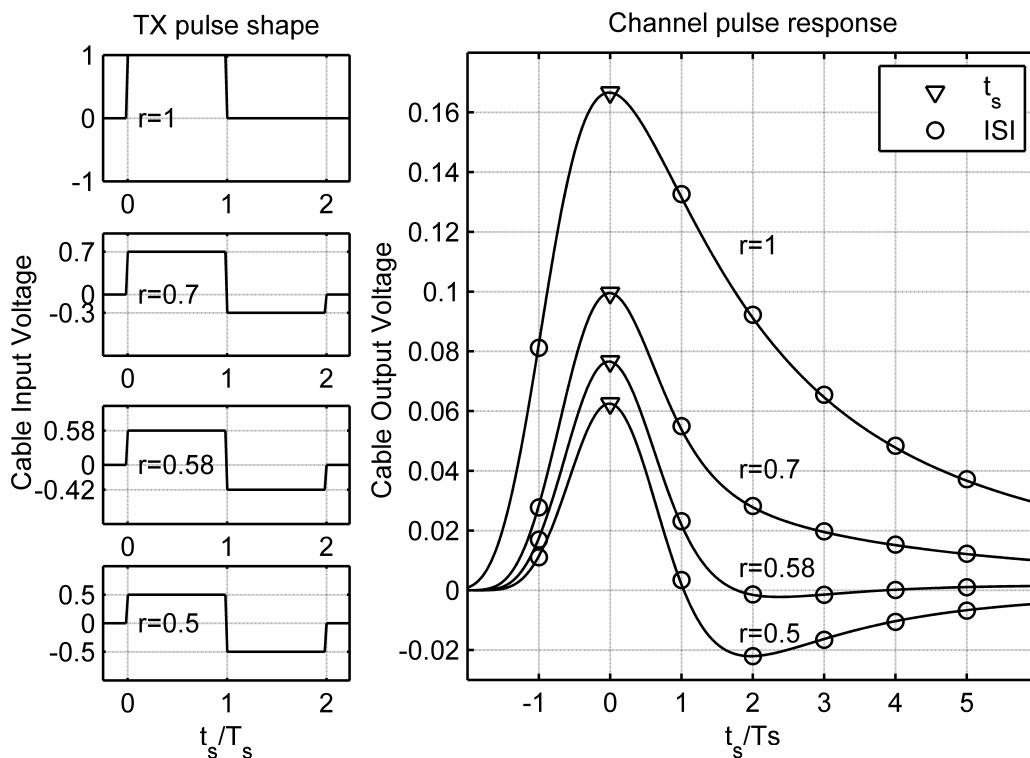


Fig. 3. TX pulse shapes ($T_s=200\text{ps}$) of SSF-PE filter with varying r parameter and simulated responses of 25 m RG-58CU cable.

By replacing FIR pre-emphasis with PWM pre-emphasis (PE), amplitude resolution requirements are replaced with timing resolution requirements, which is beneficial in the light of technology scaling. Future CMOS technologies will have lower voltage headroom and offer higher switching frequencies. PWM-PE exploits the timing resolution available in modern CMOS processes.

4.3. Frequency domain comparison

To be able to calculate the frequency transfer functions of the filters and compare them, the well-known stochastic method for calculation of the Power Spectral Density (PSD) is used as a basis [Couch]. The stochastic approach gives the PSD for a random data sequence. This method is described briefly below. We first calculate the PSD and transfer function of PWM pre-emphasis (subsection 4.3.1), then calculate the PSD and transfer function of the 2-tap SSF filter (subsection 4.3.2). These transfer functions are compared in 4.3.3. Next, in 4.3.4 we compare the flatness of the equalized channel transfer functions. Finally, in subsection 4.3.5 we comment on the high-frequency (HF) behavior of the PWM pre-emphasis filter.

4.3.1. PSD and transfer function of PWM pre-emphasis

First, the transmitted data pattern $data_{tr}(t)$ is defined as:

$$data_{tr}(t) = \sum_{n=-\infty}^{\infty} a_n p(t - nT_s), \quad (1)$$

where $a_n \in \{-1,1\}$ denotes the random data, T_s is the symbol duration, and $p(t)$ is the pulse shape. Both levels -1 and 1 for a_n are equally likely to occur. For example, $a = \{\dots, 1, -1, -1, 1, -1, 1, 1, \dots\}$. The PWM pulse $p(t) = p_{pwm}(t)$ is defined as follows (illustration in Fig. 4(a)):

$$p_{pwm}(t) = \begin{cases} 0, & t < -T_s / 2, \\ 1, & -T_s / 2 \leq t < (d - 1/2)T_s, \\ -1, & (d - 1/2)T_s \leq t < T_s / 2, \\ 0, & T_s / 2 \leq t, \end{cases} \quad (2)$$

where d denotes the duty-cycle ($0.5 < d < 1$ fits best to copper cables) and T_s denotes the symbol duration.

The PSD $PSD(\omega)$ of the stochastic signal can now be calculated using the formula [Couch]:

$$PSD(\omega) = \frac{|P(\omega)|^2}{T_s} \sum_{k=-\infty}^{k=\infty} R(k) e^{jk\omega T_s}, \quad (3)$$

where $P(\omega)$ is the Fourier transform of $p(t)$, or in other words the pulse spectrum, and $R(k)$ is the autocorrelation of the random data sequence a_n .

The autocorrelation $R(k)$ is the same as for polar NRZ signaling, and is calculated as follows [Couch]:

$$R(k) = \sum_{i=1}^I (a_n a_{n+k})_i P_i = \begin{cases} A^2, & k = 0, \\ 0, & k \neq 0, \end{cases} \quad (4)$$

where a_n and a_{n+k} are the multiplication factors for the data pulses at the n th and $n+k$ th symbol positions, respectively, and P_i is the probability of having the i th $a_n a_{n+k}$ product. We use polar signaling (a_n can be either $-A$ or A). $A=1$ in our case. I is the number of possibilities for the product values, which is equal to 4 for polar NRZ (and the PWM pre-emphasis filter).

We can now calculate the PSD. The TX output swing is normalized to $\pm 1V$ (as illustrated in Fig. 2). The spectrum $P_{pwm}(f)$ of the PWM pulse is calculated by taking the Fourier transform of $p_{pwm}(t)$:

$$P_{pwm}(f) = \int_{-\infty}^{\infty} p_{pwm}(t) e^{-j\omega t} dt = \int_{-T_s/2}^{(d-0.5)T_s} e^{-j\omega t} dt + \int_{(d-0.5)T_s}^{T_s/2} -e^{-j\omega t} dt. \quad (5)$$

Simplifying leads to:

$$P_{pwm}(f) = \frac{2}{j\omega} \left(\cos(\omega T_s / 2) - e^{-j\omega(d-0.5)T_s} \right). \quad (6)$$

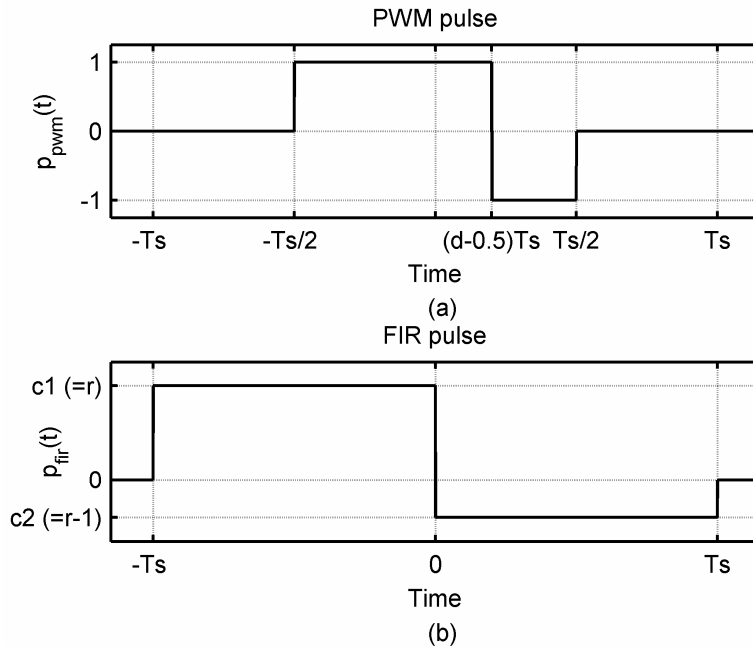


Fig. 4. Definition of pulse shapes. (a) PWM-PE. (b) 2-tap SSF-PE.




Fig. 5. Power spectral densities for the PWM signal.

Now we calculate the power spectral density PSD_{pwm} for the PWM filter:

$$\begin{aligned}
 PSD_{pwm}(\omega) &= \frac{|P_{pwm}(\omega)|^2}{T_s} \sum_{k=-\infty}^{k=\infty} R(k)e^{j\omega k T_s} \\
 &= 2 \frac{\cos(\omega T_s) - 2 \cos(\omega d T_s) - 2 \cos(\omega(d-1)T_s) + 3}{\omega^2 T_s}.
 \end{aligned} \tag{7}$$

This function is shown in Fig. 5. The situation with $d=1$ corresponds to normal polar NRZ line coding, and the situation with $d=0.5$ corresponds to Manchester line coding. Note how the spectrum is shaped and moved to higher frequencies by the (time-varying) PWM filter.

The frequency domain transfer function $H_{pwm}(f)$ of the PWM-PE filter can now be calculated as follows (because $R(k)$ is the same for PWM as for polar NRZ):

$$H_{pwm}(f) = \frac{P_{pwm}(f)}{P_{NRZ}(f)}, \tag{8}$$

where $P_{NRZ}(f)$ is the spectrum of a normal polar NRZ pulse of width T_s and height 1, which is well known to be:

$$P_{NRZ}(f) = \frac{2}{\omega} \sin(\omega T_s / 2). \tag{9}$$

The expression for $H_{pwm}(f)$ now becomes:

$$H_{pwm}(f) = \frac{\cos(\omega T_s / 2) - e^{-j\omega(d-0.5)T_s}}{j \sin(\omega T_s / 2)}. \tag{10}$$

Taking the modulus yields:

$$\begin{aligned}
& |H_{pwm}(f)| \\
&= \sqrt{2 \frac{\cos(\omega(d-1)T_s) + \cos(\omega d T_s) - 2}{\cos(\omega T_s) - 1}}.
\end{aligned} \tag{11}$$

This transfer function is illustrated in Fig. 7(a) for several values of d . (Note that 0.5 on the x-axis is the Nyquist frequency f_N .) The values for d were chosen to show the range of possible equalizer gains. In Fig. 8(a), the same function is given with the y-axis in dBs. It can be seen from this figure that a duty-cycle closer to 50% results in a steeper transfer function. Changing the duty-cycle from 100% to 50% attenuates the low-frequency components of the pulse spectrum as compared to the spectrum of a polar NRZ pulse. The result is pre-emphasis filtering. The next subsection shows how this transfer function compares to that of the 2-tap SSF.

4.3.2. PSD and transfer function of 2-tap SSF

The function $|H_{pwm}(f)|$ is now compared to the modulus of the transfer function $|H_{fir}(f)|$ of the 2-tap SSF. In Fig. 1(b), TX waveforms for a 2-tap SSF-PE filter are shown. The 2-tap FIR equalized pulse $p_{fir}(t)$, shown in Fig. 4(b), is defined as follows:

$$p_{fir}(t) = \begin{cases} 0, & t < -T_s, \\ c_1, & -T_s \leq t < 0, \\ c_2, & 0 \leq t < T_s, \\ 0, & T_s \leq t, \end{cases} \tag{12}$$

where c_1 and c_2 denote the values of the first and the second FIR taps respectively, and T_s again denotes the symbol duration. A good fit to copper cables is obtained by choosing $c_1 > 0$ and $c_2 < 0$, while $|c_1| > |c_2|$. The sum of the absolute values of all tap weights has to be limited to a certain value (at or below the supply voltage) to avoid compression at the TX output. When the output swing is again normalized to $\pm 1V$, like in the previous subsection, and the whole of the available swing is used, c_1 and c_2 have to comply with:

$$|c_1| + |c_2| = 1. \tag{13}$$

These coefficients can be rewritten as $c_1 = r$ and $c_2 = r - 1$, with one coefficient r , chosen from the interval $[0.5, 1]$, that completely controls the shape of the filter transfer function. Function $p_{fir}(t)$ now becomes:

$$p_{fir}(t) = \begin{cases} 0, & t < -T_s, \\ r, & -T_s \leq t < 0, \\ r - 1, & 0 \leq t < T_s, \\ 0, & T_s \leq t. \end{cases} \tag{14}$$

As mentioned earlier, in practice, this coefficient can be determined automatically using a control loop.

The spectrum $P_{fir}(f)$ of the PWM pulse is calculated by taking the Fourier transform of $p_{fir}(t)$:

$$P_{fir}(f) = \int_{-\infty}^{\infty} p_{fir}(t)e^{-j\omega t} dt = \int_{-T_s}^0 re^{-j\omega t} dt + \int_0^{T_s} (r-1)e^{-j\omega t} dt. \quad (15)$$

Simplifying leads to:

$$P_{fir}(f) = \frac{1}{j\omega} (re^{j\omega T_s} + (1-r)e^{-j\omega T_s} - 1). \quad (16)$$

Now we calculate the power spectral density PSD_{FIR} for the 2-tap SSF filter:

$$\begin{aligned} PSD_{fir}(\omega) &= \frac{|P_{fir}(\omega)|^2}{T_s} \sum_{k=-\infty}^{k=\infty} R(k)e^{j\omega k T_s} \\ &= 2 \frac{(r^2 - r)(1 - \cos(2\omega T_s)) - \cos(\omega T_s) + 1}{\omega^2 T_s}. \end{aligned} \quad (17)$$

This function is shown in Fig. 6. Note that the situation with $r=1$ corresponds to normal polar NRZ line coding.

The transfer function $H_{fir}(f)$ of the 2-tap SSF can be found by dividing P_{fir} over P_{NRZ} , resulting in:

$$H_{fir}(f) = \frac{1}{2} \frac{re^{j\omega T_s} - (r-1)e^{-j\omega T_s} - 1}{j \sin(\omega T_s / 2)}. \quad (18)$$




Fig. 6. Power spectral densities for the FIR signal.

Taking the modulus yields:

$$|H_{fir}(f)| = \sqrt{(r^2 - r) \frac{\cos(2\omega T_s) - 1}{\cos(\omega T_s) - 1} + 1}. \quad (19)$$

This function is illustrated in Fig. 7(b) for several values of r . The values for r were chosen to show the range of possible equalizer gains. Fig. 8(b) illustrates the same function with the y-axis in dBs. The closer r is to 0.5, the more low-frequency attenuation the filter exhibits.

4.3.3. Comparison of transfer functions

We now try to understand what makes the PWM filter fit better to the channel transfer function than the 2-tap SSF filter. We know that this must be the case because of the higher measured loss compensation, as is shown later.

Comparing the PWM and 2-tap SSF transfer functions (Fig. 7(a) / 8(a) and Fig. 7(b) / 8(b) respectively), it can be seen that in the low frequency (LF) range the PWM-PE filter behaves like the 2-tap SSF filter (first order behavior). However in the high frequency (HF) range, at frequencies just below f_N , the PWM filter behaves differently. Due to the time-varying behavior of the PWM filter, near $2f_N$ the effective gain of the PWM filter increases to infinity. As a result, at f_N (at 0.5 on the x-axis) the calculated magnitude of the PWM filter transfer function has a value of one while its derivative is non-zero, unlike with the 2-tap SSF filter. The PWM filter transfer function can be seen as a higher order filter with only one parameter.

While both filters have a different slope at f_N , their transfer at f_N is equal: both filters leave the amplitude of the fastest data transitions ('101010') unchanged. This is precisely what a pre-emphasis filter for these channels should do. In fact, de-emphasis is a better word. Ideally, the equalizer should attenuate all frequencies below f_N and leave the fastest transitions untouched. This is because the channel has an attenuation that increases monotonically with frequency, as is shown in chapter two. The cable loss at f_N therefore roughly determines the eye height at the cable output. This eye height is approximately the same for both filters up to ~ 20 dB cable loss at f_N . After that, the eye of the 2-tap SSF closes and PWM-PE still has an open eye, up to ~ 30 dB, as is shown later.

Finally, one could ask whether PWM filtering is essentially the same as 2-tap half-symbol-spaced FIR filtering. The modulus of its transfer function $|H_{fsf}(f)|$ is the same as that for 2-tap SSF, except that ω is replaced by $\omega/2$:

$$|H_{fsf}(f)| = \sqrt{(r^2 - r) \frac{\cos(\omega T_s) - 1}{\cos(\omega T_s / 2) - 1} + 1}. \quad (20)$$

This is shown in Fig. 7(c), and with the y-axis in dB in Fig. 8(c). The 2-tap half-symbol-spaced filter has a frequency response two times wider than its 2-tap SSF equivalent. However, as a result of that, the transfer function is stretched out and is not as steep as that of the PWM-PE filter in the frequency range from zero to f_N , in turn resulting in a lower loss compensation. Whereas the PWM and 2-tap SSF filters attenuate only the signals below f_N , a 2-tap half-symbol-spaced FIR filter also attenuates the signals at f_N . Half-symbol-spaced FIR filters are sometimes used at the receiver side [Farjad-Rad].




Fig. 7. Calculated magnitude of filter transfer. Note that f_N is (per definition) at 0.5 on the x-axis. (a). PWM-PE filter. (b). 2-tap SSF-PE filter. (c) 2-tap HSF-PE filter.




Fig. 8. Calculated magnitude of filter transfer, y-axis in dB. Note that f_N is (per definition) at 0.5 on the x-axis. (a). PWM-PE filter. (b). 2-tap SSF-PE filter. (c) 2-tap HSF-PE filter.

4.3.4. Equalized channel transfer function

Next, we calculate the equalized channel transfer function for the 2-tap FIR filters and the PWM filter. A theoretical first-order channel can be perfectly equalized with PWM pre-emphasis as calculated in [Schinkel-2]. A cable does not have a first order transfer function but still can be equalized with PWM pre-emphasis, and better than with 2-tap SSF, as is shown later. The equalized transfer function is calculated by taking the measured cable transfer of 25m RG-58CU cable and multiplying it with the calculated theoretical transfer of the pre-emphasis filters. A bit length of $T_s=200\text{ps}$ is chosen. The results are shown in Fig. 9. The values for r and d are chosen around the values that minimize the peak distortion. (See section 4.4.) We can now determine the flatness in the frequency interval $[0, f_N]$. Again, note that f_N is at 0.5 on the x-axis. The channel response for PWM pre-emphasis is flat to within 5dB (Fig. 9(a)), while the FIR response is only flat to within 10dB, shown in Fig. 9(b). PWM-PE clearly outperforms FIR-PE. For shorter cable lengths the responses become flatter. The half-symbol spaced 2-tap FIR is shown in Fig. 9(c); it is flat only to within 8dB.




Fig. 9. Equalized transfer of 25m RG-58CU. $H_c(f)$ indicates the measured transfer function of 25m RG-58CU. (a). PWM. (b). 2-tap SSF. (c) 2-tap HSF.

4.3.5. Filter DC behavior: move power to higher frequencies

It is important to know how the output of the PWM filter differs from that of the SSF filter when transmitting many 1s in succession, and when transmitting a quickly alternating pattern of 1s and 0s. Using a simple Fourier series calculation, we calculate the spectrum for the LF pattern $\{...,1,1,1,1,...\}$. Next, we calculate the spectrum for the HF pattern $\{...,1,-1,1,-1,...\}$ and compare the two. We then do the same for the 2-tap SSF filter.

The Fourier coefficients for a periodic waveform are calculated as follows:

$$c_n = \frac{1}{T_0} \int_0^{T_0} y(t) e^{-jn2\pi t/T_0} dt, \quad (21)$$

and the waveform then can be reconstructed using:

$$y(t) = \sum_{n=-\infty}^{\infty} c_n e^{jn2\pi t/T_0}. \quad (22)$$

For the PWM filter, we define two periodic waves, one for the LF pattern and one for the HF pattern. First, the LF $\{1,1,1,1,...\}$ pattern $y_{pwm,LF}(t)$:

$$y_{pwm,LF}(t) = \begin{cases} 1, & 0 \leq t \bmod T_s < dT_s, \\ -1, & dT_s \leq t \bmod T_s < T_s. \end{cases} \quad (23)$$

Next, the HF $\{1,-1,1,-1,...\}$ pattern $y_{pwm,HF}(t)$:

$$y_{pwm,HF}(t) = \begin{cases} 1, & 0 \leq t \bmod 2T_s < dT_s, \\ -1, & dT_s \leq t \bmod 2T_s < (d+1)T_s, \\ 1, & (d+1)T_s \leq t \bmod 2T_s < 2T_s. \end{cases} \quad (24)$$

Examples are given in Fig. 10(c) and 10(f), for $d=0.7$. For the 2-tap SSF filter, we again define two periodic waves, one for the LF pattern and one for the HF pattern. First, the LF $\{1,1,1,1,...\}$ pattern $y_{fir,LF}(t)$:

$$y_{fir,LF}(t) = r - (1-r) = 2r - 1. \quad (25)$$

Next, the HF $\{1,-1,1,-1,...\}$ pattern $y_{fir,HF}(t)$:

$$y_{fir,HF}(t) = \begin{cases} 1, & 0 \leq t \bmod 2T_s < T_s, \\ -1, & T_s \leq t \bmod 2T_s < 2T_s. \end{cases} \quad (26)$$

Examples are given in Fig. 10(b) and 10(e), for $r=0.7$. We choose d and r to be equal because then the DC levels of the outputs are equal for both filters.

In 10(a) and 10(d), the non-equalized polar NRZ patterns are shown.




Fig. 10. LF and HF patterns. (a), (b) and (c) LF patterns for NRZ, FIR and PWM respectively. (d), (e) and (f) HF patterns for NRZ, FIR and PWM respectively.

Note that the repetition period needs to be chosen as $T_0=2T_s$. Note also that $n=1$ corresponds to the Nyquist frequency f_N . Now the coefficients can be calculated using Eq. 21. For the PWM LF wave, the results are:

$$c_{n,PWM,LF} = \begin{cases} 2d - 1, & n = 0, \\ \frac{j}{n\pi} (e^{-j2\pi nd} - 1), & n \neq 0. \end{cases} \quad (27)$$

Note that all coefficients are dependent on d . For $n=0$, the magnitude is dependent on d . For $n \neq 0$, both magnitude and phase are dependent on d . The magnitude of the coefficients is plotted in Fig. 11(a).

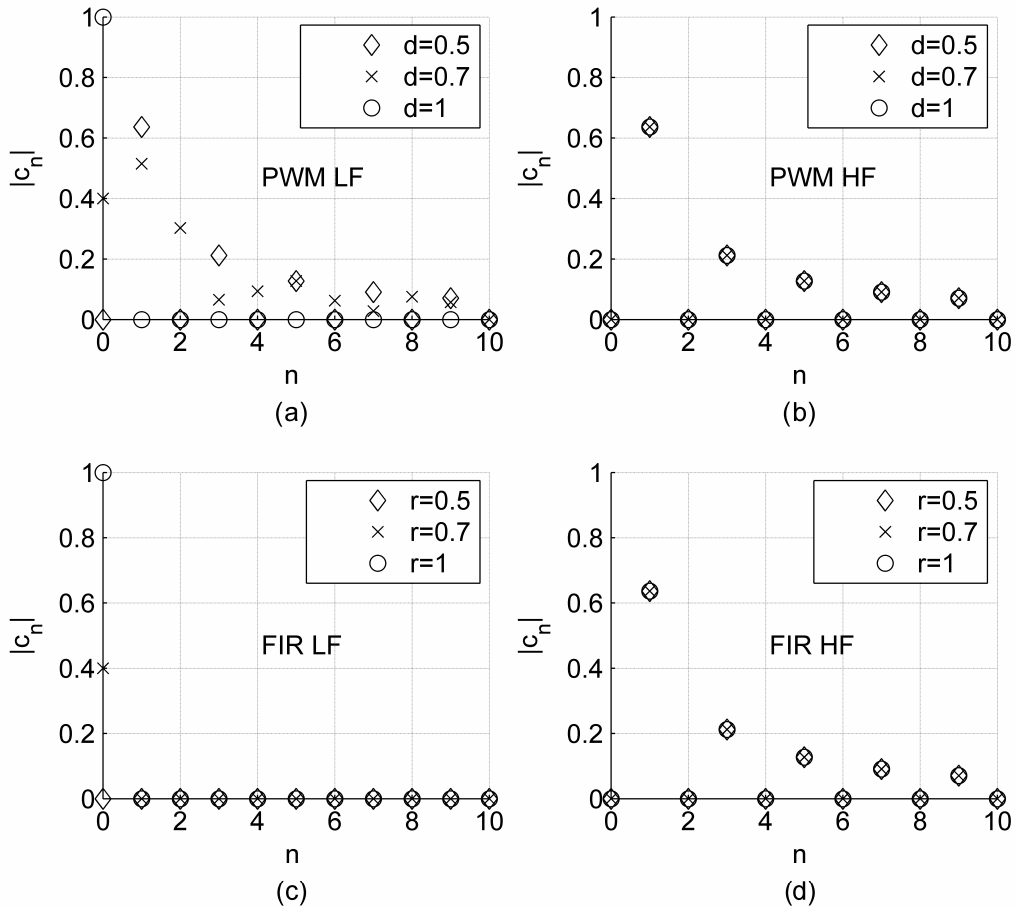


Fig. 11. Magnitude $|c_n|$ of Fourier coefficients. (a) PWM LF. (b) PWM HF. (c) FIR LF. (d) FIR HF.

For the PWM HF wave, the results are:

$$c_{n,PWM,HF} = \begin{cases} 0, & n = 0, \\ \frac{2j}{n\pi} e^{-j\pi nd}, & n = \text{odd}, \\ 0, & n = \text{even}. \end{cases} \quad (28)$$

These coefficients are simply equal to the coefficients of a (phase-shifted) square wave. Note that the phase (time shift) is dependent on d . The magnitude of the coefficients is illustrated in Fig. 11(b).

For the FIR filter, we also calculate the Fourier series coefficients. The results for the LF FIR wave are:

$$c_{n,fir,LF} = \begin{cases} 2r - 1, & n = 0, \\ 0, & n \neq 0. \end{cases} \quad (29)$$

Only for $n=0$ are the coefficients nonzero. This is illustrated in Fig. 11(c).

For the FIR HF wave, the results are:

$$c_{n,fir,HF} = \begin{cases} 0, & n = 0, \\ -2j, & n = \text{odd}, \\ \frac{n\pi}{0}, & n = \text{even}. \end{cases} \quad (30)$$

This is the same as for the PWM HF wave, except for a phase shift. Note that there is no dependency on r whatsoever. The magnitude of the coefficients is shown in Fig. 11(d).

In conclusion, the difference between PWM and 2-tap SSF is as follows. For the quickly alternating $\{1,-1,1,-1,1,-1\}$ waves, there is only a difference in phase (time shift). Neither filter attenuates the HF pattern. This is the required behavior for a de-emphasis filter that needs to compensate a monotonically decreasing cable transfer function, as mentioned in subsection 4.3.3. For the LF $\{1,1,1,1,1,1\}$ pattern, the behavior is different. A 2-tap SSF outputs less power for lower switching frequencies. Where the 2-tap SSF is decreasing the output power for low frequency output patterns, the PWM filter instead outputs the same power as for its HF pattern, because the absolute squared output voltage remains the same (see Fig. 10). Instead of decreasing the power, the PWM filter transfers that power to higher frequencies, where it is dissipated by the cable.

Could the additional HF noise, generated by the PWM filter, become a cause of crosstalk? Near-end crosstalk (NEXT), from the transmitter to a nearby, sensitive receiver, could become especially problematic. A possible solution is as follows. Most of this HF power is just dissipated by the cable, so we might as well filter this out before it leaves the transmitter, for example by using an (adjustable) low-pass filter at the TX output. As long as this low-pass filter has a pole that is well above the dominant pole in the rest of the channel, crosstalk can be effectively decreased without harming the equalization. The pole of this low-pass filter can be located at the Nyquist frequency f_N .

4.4. Numerical time-domain simulations of ISI

In the previous section, it was shown from a frequency domain perspective that PWM-PE provides more relative HF boost than 2-tap SSF. In this section, the two pre-emphasis methods are compared from a time-domain perspective by calculating the remaining ISI at the output of the channel as a function of the bit rate. We calculate the peak distortion to obtain an indication of the eye opening. The goal of the simulations is to find the maximum achievable bit rate for the pre-emphasis filters at which the remaining ISI is still acceptable.

The following approach is used. First, the channel responses to the pre-emphasis pulses $p_{pwm}(t)$ and $p_{fir}(t)$ are calculated for a number of bit rates, at the optimum pre-emphasis setting (subsection 4.4.1). Next, at these points, the remaining ISI in the received signal is calculated as a function of pre-emphasis setting, symbol length, and channel time constant using a peak distortion analysis. The remaining ISI is then plotted and compared between the two different equalizers (subsection 4.4.2). Finally, the sensitivity to jitter of the PWM pre-emphasis scheme is analyzed in subsection 4.4.3.

4.4.1. Model of skin-effect-only channel

To calculate the time domain channel response, the skin-effect impulse response $h_I(t)$ is used, as described in chapter 2, subsection 2.4.1:

$$h_I(t) = \frac{\sqrt{\tau_I}}{2t\sqrt{\pi t}} \cdot e^{-\frac{\tau_I}{4t}}, \quad (31)$$

where τ_I is a channel time constant. In Fig. 12, $h_I(t)$ is shown. The x-axis shows time divided by τ_I . The y-axis shows $h_I(t) \cdot \tau_I$. The axes are chosen in this way to clearly show the maxima and time span of the functions. It can be seen that $h_I(t)$ is asymmetrical over time with a very long tail.




Fig. 12. Skin-effect impulse response $h_I(t)$.

We use this skin-effect-only channel as a good approximation of the response of many types of cable, as these are mostly skin-effect dominated. Also, to an extent, the dielectric impulse response has a similarly asymmetrical shape¹. See Chapter 2 for a more detailed estimation of the accuracy of this simplification.

The Fourier transform of Eq. 31 gives us the frequency domain transfer function $H(j\omega)$ for the skin-effect:

$$H(j\omega) = e^{-\sqrt{j\omega\tau_I}}. \quad (32)$$

The magnitude of this transfer function at the Nyquist frequency f_N for a skin-effect-only channel can be calculated by substituting $\omega=2\pi f_N=\pi/T_s$ in Eq. 32 and taking the magnitude:

$$|H(f_N)| = \left| e^{-\sqrt{j\pi\frac{\tau_I}{T_s}}} \right| = e^{-\frac{1}{2}\sqrt{2\pi\frac{\tau_I}{T_s}}}. \quad (33)$$

Note that $|H(f_N)|$ is solely dependent on the ratio T_s/τ_I . The dependence is illustrated in Fig. 13.

¹ When writing our paper [Schrader-6] we were not yet aware of recent advances in dielectric loss modeling and we used an inaccurate model [vdPlaats] for the time-domain simulation of dielectric loss. Thus a symmetrical impulse response was used, while – on the contrary – we should have used an asymmetrical one, much like the skin effect impulse response. (See also Chapter 2). As a consequence, we were led to the conclusion that the PWM equalization method would not work so well on channels dominated by dielectric loss. As we now know, this was far too pessimistic. In this chapter, we will show with measurements that the PWM equalization method can indeed provide 25dB loss compensation on an FR4 PCB trace.




Fig. 13. Magnitude of transfer function at f_N of skin-only channel, as a function of T_s/τ_1 . (a) Linear. (b) In dBs.

4.4.2. Peak distortion analysis

A theoretical channel with only skin loss and no dielectric loss ($\tau_1 \neq 0$, $\tau_2 = 0$) is analyzed. The response $y_{pwm}(t)$ of the channel to a single PWM pulse is calculated as follows:

$$y_{pwm}(t) = p_{pwm}(t) * h_c(t), \quad (34)$$

where the asterisk denotes the convolution operation. To quantify the level of remaining ISI, the peak distortion is calculated, which is defined as [Lucky]:

$$Ds_{peak,ISI} = \frac{1}{|y_{pwm}(t_s)|} \sum_{\substack{n=-\infty \\ n \neq 0}}^{n=\infty} |y_{pwm}(t_s + nT_s)|, \quad (35)$$

where t_s is the sample moment. For example, a value of $Ds_{peak,ISI} = 0.2$ means that the eye diagram for the worst-case data pattern is 20% closed. In a simple receiver with a bang-bang phase detector, the received signal's median zero crossing is used for time reference, and then data is sampled at a fixed time, generally $T_s/2$, away from the zero crossing. After that, the ISI contributions can be found at distances of $n \cdot T_s$ from t_s . The peak distortion $Ds_{peak,ISI}$ is a function of the symbol length T_s , the channel time constant τ_1 , the sample moment t_s and the duty-cycle d (or parameter r for the SSF filter).

A reduction in the number of variables can be made by acknowledging that $Ds_{peak,ISI}$ only depends on the ratio T_s/τ_1 . A large value for T_s/τ_1 means that the bit rate is very low compared to the channel speed, whereas a small value means the opposite. The expectation is that for very small ratios the ISI becomes unacceptably large, whereas for very large ratios it converges to zero.

The mathematics involved in the manipulation of Eqs. 34 and 35 is rather complicated and does not yield concise symbolic results. Therefore, we resort to numerical computation. In the simulation, the zero-forcing criterion is applied to $Ds_{peak,ISI}$ to find the optimum duty-cycle setting d_{opt} (and the optimum FIR filter parameter r).

In Fig. 14(a), the value of $DS_{peak,ISI}$ for d_{opt} is plotted versus T_s/τ_I , and in Fig. 14(b) d_{opt} is plotted versus T_s/τ_I . An identical procedure is followed for the 2-tap SSF PE filter. Again, the minimized $DS_{peak,ISI}$ for r_{opt} for the SSF is plotted in Fig. 14(a), and the optimum SSF coefficient r_{opt} is plotted in Fig 14(b)².

Choosing $DS_{peak,ISI}=0.2$ for a reasonable eye opening, it can be seen from the figures that the FIR filter reaches this point at $T_s/\tau_I=0.19$ whereas the PWM filter reaches it at $T_s/\tau_I=0.09$. This means that the PWM-PE achieves twice the bit rate of the 2-tap SSF for the same peak distortion.

4.4.3. Sensitivity to jitter

Another point of interest is the sensitivity of PWM pre-emphasis to timing errors. An error in the duty-cycle will effectively change the pre-emphasis setting and thus have an effect on the eye opening at the receiver side. However, a certain error can be allowed in the duty-cycle setting.

Fig. 14(c) shows the width W_d of the duty-cycle range wherein $DS_{peak,ISI}<0.2$. For the 2-tap SSF-PE filter, the same figure shows the width W_r of the range of filter coefficient r wherein $DS_{peak,ISI}<0.2$.

The y-axis in Fig. 14(c) is explained in the following example. It can be seen that $W_d=0.057$ at $T_s/\tau_I=0.3$, indicated in the figure by a '+'. From Fig. 14(b), the optimum duty-cycle $d_{opt}=0.565$ indicates the middle of this range. This means that a duty-cycle from the interval (0.537,0.594) will yield a peak distortion $DS_{peak,ISI}<0.2$. Next, for 2-tap SSF, from Fig. 14(c), $W_r=0.054$ at $T_s/\tau_I=0.3$, indicated in the figure by a circle. From Fig. 14(b), the optimum r parameter $r_{opt}=0.610$ indicates the middle of this range. This means that a value for r from the interval (0.583,0.637) will yield a peak distortion $DS_{peak,ISI}<0.2$.

It can be seen that both W_d and W_r become smaller with increasing bit rate. The identical behavior shown in Fig. 14(c) demonstrates the interchangeability of timing and amplitude precision.

² Several different regions can be identified in the plots. In each region, another ISI term dominates the response (pre-cursor, 1st post-cursor, 2nd post-cursor, etc.).

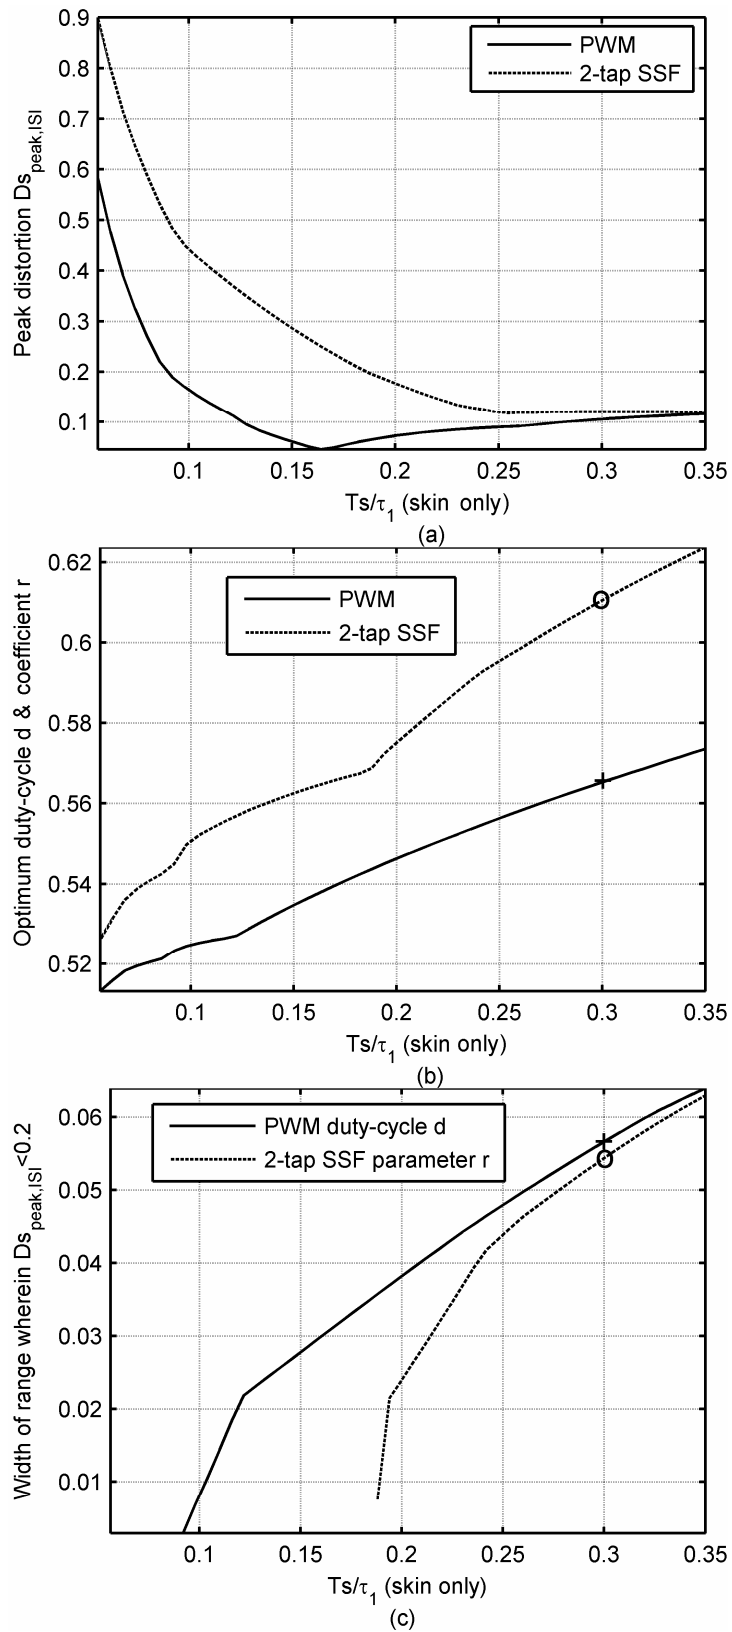


Fig. 14. (a) Simulated minimized peak distortion for skin channel. (b) Simulated optimum PWM duty-cycle and FIR coefficient for skin channel. (c) Simulated interval width for $D_{s_{peak,ISI}} < 0.2$ of d (PWM) and r (2-tap SSF).

4.5. PWM 1-edge chip design and measurements

In this section, three proof-of-concept chip designs and their measured performance are presented. In subsection 4.5.1, we describe the first $0.13\mu\text{m}$ chip and give measurement results. Next, we describe the 90nm chip and give measurement results (subsection 4.5.2). Finally, in subsection 4.5.3, the second $0.13\mu\text{m}$ chip is described which was used to equalize a PCB trace.

All chips were designed using current-mode logic (CML) to provide maximum supply noise rejection and minimum supply noise injection and to keep timing noise as low as possible. A difference between up and down slew rates would have a negative effect on the fit of the pre-emphasis to the channel. The use of differential CML guarantees symmetrical up and down slew rates. An advantage over FIR pre-emphasis is that non-linear (symmetrical) slewing effects do not affect the equalizer's fit to the channel because only two signal levels are used. Moderate bandwidth limitations in the circuitry do not pose a problem because they will just become part of the total channel transfer function that needs to be compensated by the pre-emphasis.

4.5.1. Chip 1: Design and Measurements

In this subsection, the first $0.13\mu\text{m}$ chip is described. First, a circuit overview is given in 4.5.1.1. Next, in 4.5.1.2, the measurement results are given.

4.5.1.1. Circuit overview

The principle of operation of chip 1 is shown in Fig. 15(a). The data is XOR'ed with a pulse-width modulated clock in order to provide pre-emphasized data. The PWM clock is generated using an OR gate and a delay circuit. The timing of the signals is illustrated in Fig. 15(b).

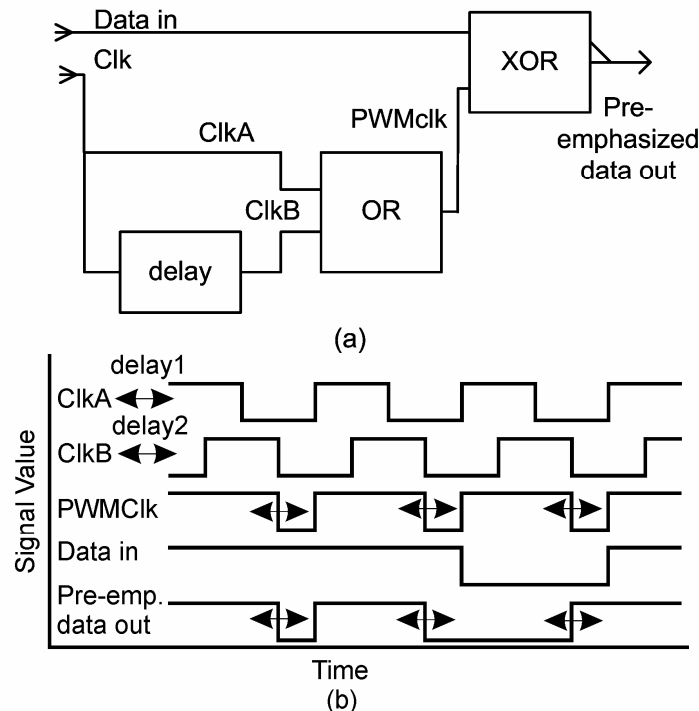


Fig. 15. (a). Circuit operation principle for chip 1. (b). Signals used in generating PWM signal.

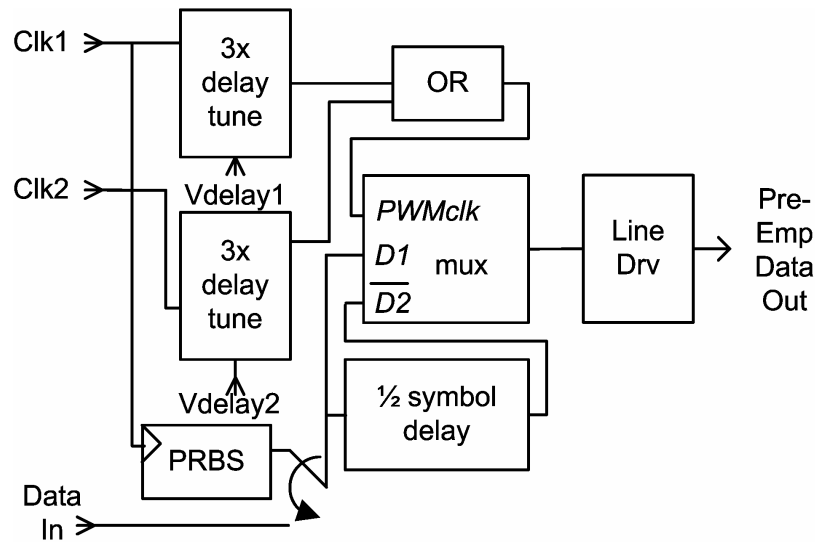


Fig. 16. Chip block diagram: PRBS generator, pre-emphasis circuit and line driver.

Fig. 16 shows the detailed block diagram. Because a small differential delay is more straightforward to generate than a short absolute delay, the relative delay for clock B is created by delaying $clk1$ by delay1 and delaying $clk2$ by delay2. The differential delay is thus equal to $(delay1 - delay2)$ and can be controlled by voltage $V_{delay1} - V_{delay2}$. Both are differential voltages.

The XOR, shown in Fig. 17(a), is implemented using a multiplexer that selects either non-inverted data $D1$ or inverted data $D2$. For an optimum timing margin, $D2$ is delayed by half a symbol time using a negative edge clocked flip-flop [Schinkel-2]. The duty-cycle of the PWM pulse shape can be tuned between 50%-100%, provided the relative phase-shift between clocks is adjustable between 0° and 180° .

The time-shifted clock is generated using a variable delay circuit, illustrated in Fig. 17(b) [Razavi]. The delay between input and output of this circuit is mainly determined by the RC time at the output. By adding a positive feedback circuit in parallel to the output, which effectively behaves as a negative resistance, the effective R can be changed and hence the RC-delay is changed. The value of the negative resistance is controlled by the differential delay control voltage V_{delay} ($=V_{delP} - V_{delN}$), which divides the total bias current between the input differential pair and the negative resistance pair. For $V_{delP} \gg V_{delN}$, the delay is minimized. As the total bias current through the output resistors is fixed, the output swing remains constant. The required tuning range of the delay circuit depends on the desired symbol length and on the necessary duty-cycle range for pre-emphasis. The continuous tuning range can be enlarged by cascading multiple delay stages. Three stages are cascaded in our prototype to provide sufficient tuning range. For very large delay ranges this becomes impractical and it is more effective to combine the continuously tunable delay with discrete fixed delay steps. The prototype is designed to give the flexibility needed to evaluate the new PWM-concept in various ways. Therefore external clocks can be provided, for example to accommodate very low bit rates for long, low-quality cables. However, during normal operation, both inputs $clk1$ and $clk2$ (Fig. 16) can just be connected to the same clock.

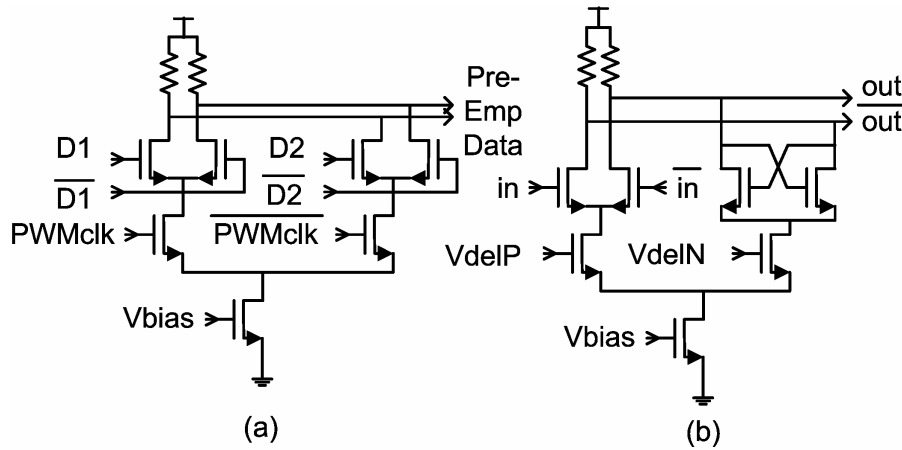


Fig. 17. (a) CML multiplexer. (b) CML delay tuning circuit.

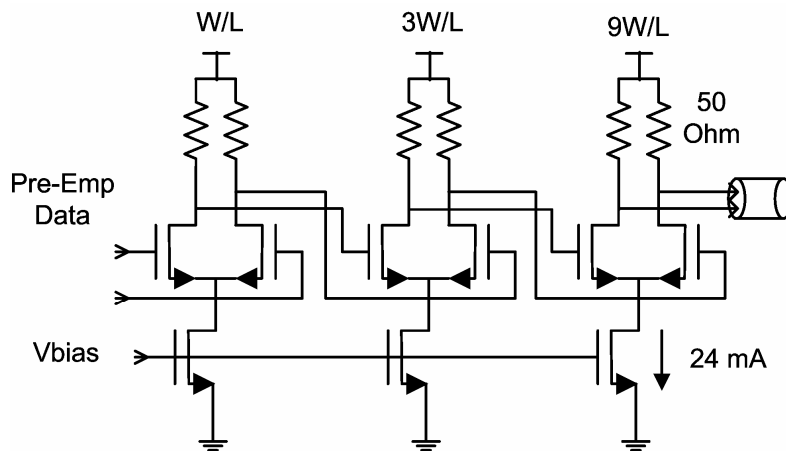



Fig. 18. Three-stage differential line driver of chip 1.

The line driver on chip 1 consists of three stages (Fig. 18). The final stage has a 50Ω on-chip termination resistance and a tail current of 24mA . This results in a nominal single-ended output swing of $600\text{mV}_{\text{p-p}}$ in a 50Ω cable. This corresponds to a differential voltage of $1.2\text{V}_{\text{p-p}}$. The line driver power is 42mW at 1.2V .


4.5.1.2. Measurements

A microphotograph of the first chip is shown in Fig. 19(a). This chip measures $1\times 1\text{mm}^2$, and has an active area of approximately $150\times 150\text{ }\mu\text{m}^2$. All chip I/Os have on-chip 50Ω termination and are ESD protected.

To evaluate the performance of the prototype chips, eye diagrams were created and BER tests run. Measurements were made using four cable assemblies and a PCB – all described and modeled in Chapter 2. For chip 1, the DC voltages/currents, including the power supply, were connected to a PCB by wirebonds, and the high speed inputs and outputs were probed. A 50Ω differential probe with 4 pins was used: ground-signal-signal-ground. If the high-speed data outputs were to be connected via a PCB, there might be an impedance mismatch between the PCB and the cable, causing reflections. (However, such reflections would be largely absorbed by the transmitter termination resistors. For the measurements we wanted to maximize signal integrity so we used probes.) The printed circuit board with the chip mounted is shown in Fig. 19(b).



(a)



(b)

Fig 19. Chip micrograph, printed circuit board and probe configuration. (a). Chip micrograph, 1x1 mm². (b) Printed circuit board with chip wire-bonded to it, and measurement probes for high speed signals.



Fig 20. Measurement setup.

The measurement setup is shown in Fig. 20. The list of equipment used is as follows:

- Agilent DCA86100A Digital Communications Analyzer (oscilloscope, eye diagrams),
- Anritsu pattern generator / BER tester,
- HP semiconductor parameter analyzer (used for generating bias currents & control voltages for setting the duty-cycle),
- several power supply units to give the required voltages,
- chip probe station.

Effect of Adjustments in PWM Duty-Cycle

In Figs. 21(a), (c) and (e), the effect of adjusting the PWM duty-cycle on the transmitter output can be seen. The left and right edges in the eye diagrams correspond to the symbol edges. In Figs. 21(b), (d) and (f), the responses of a 10-m RG-58CU cable to the pre-emphasized data stream with different pre-emphasis duty-cycles are shown. It can be seen that there is an optimum duty-cycle [Figs. 21(c) and (d)]. Under-emphasis is shown in Figs. 21(a) and (b) and over-emphasis in Figs. 21(e) and (f). Note that the time scale is the same in all the figures. The PWM pre-emphasis leaves the fastest data pattern ('101010') unchanged while it attenuates those data patterns with fewer transitions (e.g. a long string of 1s) .

Eye Diagrams at Max. Loss Compensation

The requirement for all measurements is to achieve a BER of $<10^{-12}$. The cable lengths are chosen such that this BER figure can be achieved at a bit rate of 5 Gb/s, the highest speed achievable using the TX circuitry. The channel loss at the Nyquist frequency can then be taken as a figure-of-merit for the equalizer. We call this the 'loss compensation' of the equalizer. In Figs. 22(a)-(d), measured eye diagrams of the cable outputs at 5 Gb/s are shown.

All four cable types, described in section 4.3, are shown. Figs. 22(a)-(c) are for the coaxial cables and Fig. 22(d) is for the differential cable. For the coaxial cables, one of the two transmitter outputs was used while the other was terminated with a 50Ω dummy load. For the differential cable, both transmitter outputs were used. The oscilloscope input is single-ended. Therefore, the eye-diagram shown in Fig. 22(d) was measured at the output of the differential limiting amplifier. For cables (a)-(c), the loss at 2.5GHz is ~30dB. For cable (d), the loss at 2.5GHz is 19dB. (See also Chapter 2.) The total channel loss is a few dB more than the cable loss alone, because of additional losses from probes, short wire, bias tee and connectors. Using an external pattern generator and tester, the BER was tested with all cable assemblies at 5 Gb/s and is $<10^{-12}$.

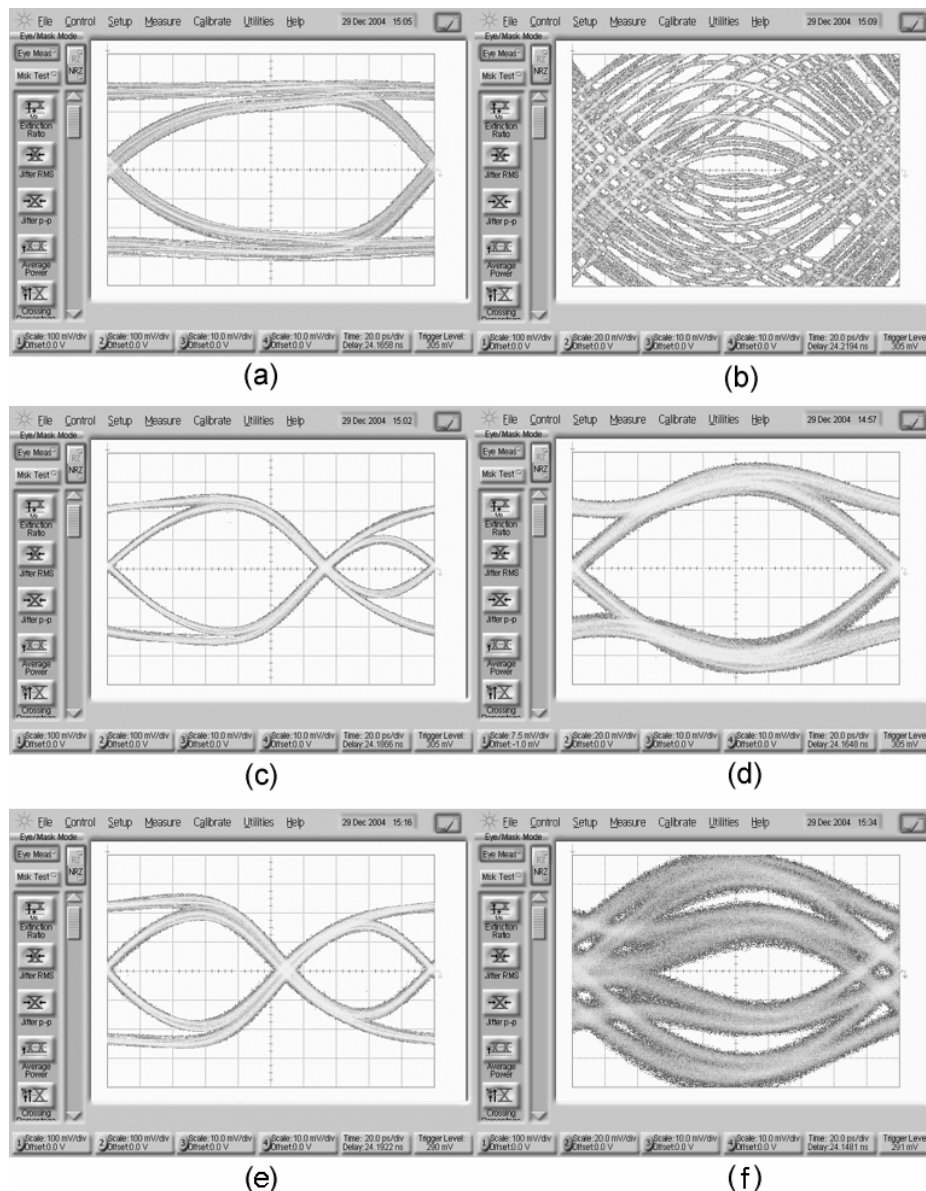


Fig. 21. Measured eye-diagrams for transmitter output and cable output (10 m RG-58CU) at 5 Gb/s. Horizontal axis = 20ps/div, vertical axis is 100mV/div. for (a), (c) and (e), and 20mV/div for (b), (d) and (f).

(a) TX, no pre-emphasis (100%). (b) RX, no pre-emphasis (100%). (c) TX, optimum pre-emphasis (66%). (d) RX, optimum pre-emphasis (66%). (e) TX, strong pre-emphasis (55%). (f) RX, strong pre-emphasis (55%).




Fig. 22. Measured output eyes of chip 1 at 5 Gb/s and $BER < 1 \cdot 10^{-12}$. Horizontal axis = 20ps/div, vertical axis is 4mV/div. for (a)-(c) and 20mV/div. for (d).
 (a) 25 m RG-58CU. (b) 130 m Aircom+. (c) 80 m Aircell7. (d) 15 m 10GBASE-CX4 24AWG after differential limiting amplifier.

At a channel loss of $>30\text{dB}$, the small cable output amplitude imposes a high demand on receiver sensitivity. For the BER measurements, a limiting amplifier was used with a 1mV input offset. Using the fully differential transmitter capabilities with cable (d) offers the advantage of a 6dB higher (differential) swing while also rejecting common mode noise.

In Table 1, a comparison with other published work is given. For the 2PAM systems, none of the pre-emphasis filters that use 2 taps offer more than 18dB loss compensation [Lee], [Kudoh]. A 5-tap FIR filter reached 30dB in a 2PAM system, but only at 3.125 Gb/s [Gai]. More taps can offer higher loss compensation but at the expense of increasing complexity, possibly causing accuracy and speed problems. Furthermore, algorithm convergence for automatically finding the optimum equalizer coefficients is more troublesome than for an equalizer with only one ‘knob’. The PWM-PE filter presented here offers a record loss compensation (33dB), at a bit rate of 5 Gb/s.

(In [Farjad-Rad] a 4PAM transmitter is described, which cannot be compared directly because it theoretically requires a 9.5dB higher SNR for the same BER, due to the smaller eye opening.)

In Table 2, the electrical characteristics of the transmitter are given. Power dissipation figures are hard to compare because most publications only give total figures. In the current proof-of-concept design, the clock buffering takes much of the power budget, but this drain can be reduced when internal clocks are available on the IC. Because of the simplicity of the pre-emphasis method, both the chip area and power consumption can be very low.

Ref.	R	Loss	Process	Type
[Farjad-Rad]	8 Gb/s	~10dB	0.3 μ m	2-tap FIR
[Lee]	4 Gb/s	~10dB	0.25 μ m	2-tap FIR
[Kudoh] TX only	5 Gb/s	18dB	0.13 μ m	2-tap FIR
[Gai] TX only	3.125 Gb/s	30dB	0.11 μ m	5-tap FIR
this work	5 Gb/s	33dB	0.13 μ m	PWM

Table 1. Pre-emphasis comparison with other work. In [Kudoh] and [Gai] the loss compensation from the TX pre-emphasis was measured by turning off the receiver equalizer.

Bit rate (2-PAM)	5Gb/s
U-I	200ps
TX amp. (V_{p-p}) nom.	1.2V (dif), 600mV (s-ended)
Channel loss @ 2.5GHz	33dB
V_{sup}	1.2V
Power (pre-emphasis)	12mW
Power (line driver)	42mW
Power (clock buffering)	39mW
Power (on-chip PRBS)	17mW

Table 2. Electrical characteristics of transmitter.

4.5.2. Chip 2: Design and Measurements

The high cable loss results in a low amplitude at the cable output, which can lead to problems at the receiver side. The SNR might be too low, or the received signal amplitude might be below the receiver offset. A usual value for the receiver offset is ~ 50 mV. Therefore, it would be convenient to have a transmitter with a higher output swing. With this in mind, a new line driver was designed, which is described in 4.5.2.1. The duty-cycle tuning circuit that was used is described in 4.5.2.2. Finally, the measurement results are given in 4.5.2.3.

4.5.2.1. Line Driver

The line driver circuit of chip 2 extends the line driver of chip 1 by cascading a cascode stage. High-voltage transistors are not needed. This enables a maximum output voltage of 2.5V. The output current is adjustable from 36 to 74mA, resulting in a 1.8V_{p-p} single-ended voltage swing in 50 Ω or 25 Ω , the latter in the case of an external TX termination resistor. Compared to the first chip, the TX power is almost 10dB higher. This would lead to high power dissipation in an on-chip TX termination resistor, dissipating $(1.8)^2/50=65$ mW. The excess heat could pose a problem when many high-speed links are operated in parallel. For that reason, the TX termination resistor was left out and the driver is open-drain. It is shown in Fig. 23. The cable needs to be terminated to 2.5V, in order to supply the driver bias current.




Fig. 23. Line driver of chip 2 with cable and termination.

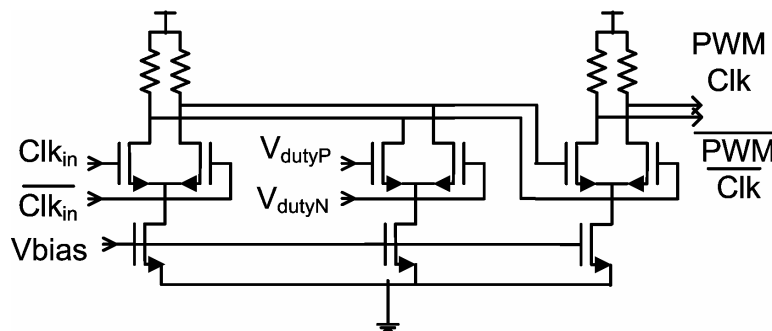


Fig. 24. Duty-cycle tuning circuit of chip 2.

4.5.2.2. Duty-cycle tuning circuit

A duty-cycle tuning circuit different from that used in chip 1 is chosen, because the latter has some drawbacks. It requires the adjustment of two voltages, and its duty-cycle is not nominally set at 50%, which would be convenient for a cable with significant loss. The circuit on chip 2 does have a nominal setting of 50% [Westergaard]. It outputs a duty-cycle of 50% at a differential voltage $V_{dutyP} - V_{dutyN} = 0V$, as shown in Fig. 24. The differential voltage $V_{dutyP} - V_{dutyN}$ controls the the offset of the middle differential pair. The dV/dt voltage slope then converts this into a duty-cycle change.

4.5.2.3. Measurements

A micrograph of chip 2 is shown in Fig. 25. The chip measures $1 \times 0.9 \text{ mm}^2$, and has an active area of approximately $100 \times 100 \text{ }\mu\text{m}^2$. The control voltages, bias currents, power supply, and the high speed in-/outputs were all connected via probes. Because our oscilloscope is terminated to ground, a bias tee was used at the cable output. In this way, we could still supply the driver bias current while terminating the cable to ground. The DC input of the bias tee was connected to 1.65V, which is the threshold voltage of the driver output for 1.8V swing ($2.5 - 1.8 / 2 = 1.65V$).

The output of the TX at 4 Gb/s is shown in Fig. 26, clearly showing its high swing. In Figs. 27(a)-(d), measured eye diagrams of the cable outputs are shown. The maximum speed achieved with cables (a)-(c) is 3 Gb/s for a BER $< 1 \cdot 10^{-12}$. The loss of cables (a)-(c) at 1.25GHz is $\sim 20\text{dB}$. For the differential cable (d), a speed of 5 Gb/s was achieved. The loss of

this cable at 2.5GHz is 19dB. A possible explanation for the slightly lower performance than chip 1 (30dB) is the data-dependent noise at the TX output. In Fig. 26, separate lines can be identified. (The 2^7-1 PRBS pattern repeats after 127 bits. The lines disappear when the PRBS pattern is switched from 2^7-1 to $2^{31}-1$.) The data-dependent noise is probably caused by the lack of an on-chip termination resistance. For canceling these reflections, the use of a 50Ω on-chip TX termination is commonly mentioned in the literature. To save power, we removed this termination resistor but, as stated earlier, this might have caused the decreased performance of chip 2. Ideally we would like to save power and cancel reflections at the same time, but for this we would need to find (for example) a suitable virtual impedance driver concept.

The measured eye openings in Figs. 27(a)-(c) are clearly larger than for chip 1, due to the larger TX output swing, even when the 8dB difference in channel loss at the (lower) Nyquist frequency is taken into account.




Fig. 25. Chip 2 micrograph, 1x0.9 mm².




Fig. 26. Single-ended transmitter eye of chip 2, for 2^7-1 PRBS pattern at 4 Gb/s. Horizontal axis = 25ps/div., Vertical axis=100mV/div. (10dB att.)




Fig. 27. Measured output eyes of chip 2 at $\text{BER} < 1 \cdot 10^{-12}$. Horizontal axis is 33.3ps/div for (a)-(c), and 20ps/div. for (d). Vertical axis is 30mV/div. for (a)-(c), and 20mV/div. for (d).
 (a) 25 m RG-58CU at 3 Gb/s. (b) 130 m Aircom+ at 3 Gb/s. (c) 80 m Aircell7 at 3 Gb/s.
 (d) 15 m 10GBASE-CX4 24AWG at 5 Gb/s after differential limiting amplifier.

4.5.3. Chip 3: measurements on a PCB trace

In data-communication over legacy FR4 backplanes (printed circuit boards), the most important speed limit is not the skin effect, as in cables, but dielectric polarization and relaxation. A 2-5 tap symbol-spaced FIR filter is commonly used for TX pre-emphasis (PE) on backplanes [Stonick], [Balan], [Payne], [Krishna], [Zerbe]. We have shown above that pulse-width modulation (PWM) TX PE can effectively replace FIR TX PE for copper cables. We show here that it also works for printed circuit boards. In subsection 4.5.3.1 the circuit of a $0.13\mu\text{m}$ chip, that was designed for this purpose, is described. In 4.5.3.2, the measurement results are given.

4.5.3.1. Circuit description

A test transmitter chip was designed in $0.13\mu\text{m}$ CMOS. At the time of the design of the chip, we thought that the impulse response from dielectric loss was more symmetrical than it is in reality (see also the footnote on page 79), so an extra degree of freedom was added. Where previously the pulse was always on the left side, we can now place it at an arbitrary position, e.g. in the middle. This is illustrated in Fig. 28(b), along with the block diagram and electrical signals in Fig. 28(a). The circuit was designed in differential CML. The duty-cycle control is different from the other chips and is based on a divider and XOR circuit [Nam]. The delays between signal 'I' and 1 and signal 'Q' and 2, delay1 and delay2 respectively, are generated with four cascaded tunable delay sections, shown in Fig. 17(b). The differential line driver, illustrated in Fig. 18, outputs a single-ended swing of 600mVp-p in 50Ω (differential 1.2Vp-p in 100Ω).




Fig. 28. Chip block diagram and signals. (a) Block diagram. (b) Signals.




Fig. 29. Chip micrograph. Line driver width is 110µm.

4.5.3.2. Measurements

In Fig. 29 the chip microphotograph is shown. The measurement goal was again to evaluate the maximum loss compensation capability of the PWM-PE technique. A very long PCB trace of 270cm (106”) of FR4 was used to be able to achieve a high channel loss at the Nyquist frequency. (See also Chapter 2 for a description and photograph.) The trace is single-ended. It is a microstrip with a characteristic impedance of 50Ω. There are no vias, and at the end there are SMA connectors. The total channel (PCB + 1.75m coaxial cable + bias tee) has 25dB loss at the Nyquist frequency of 2.5GHz. Measurements were made using chip probes. One of the two TX outputs was connected to the trace, and the free TX output was directly connected to the 50Ω oscilloscope.

In Fig. 30, the transmitter output obtained using the extra degree of freedom (the pulse position) is shown. Two bits are shown in the picture: the length of the x-axis is two bit times. Looking at the left half of the figure (the first bit), during one bit time, if a ‘1’ bit is transmitted, the voltage first goes negative, then positive, and then negative again. Comparing the two bits, we can see some asymmetry. The XOR implementation probably causes this. Looking at Fig. 28(b), we can see that, when the duty-cycle of the (half-frequency) signals 1 and 2 is not exactly 50%, signal 3 is different for even bits than for odd bits.

In Fig. 31, the eye diagrams from these measurements are shown, both without PE and for the optimum PE setting. A bit rate of 5Gb/s was achieved at a BER of $1 \cdot 10^{-12}$ over 270cm of FR4 (2^7-1 PRBS pattern). The 25dB loss compensation at the Nyquist frequency represents a good performance, compared to the ~20dB achieved in [Balan], [Payne], [Krishna]. The total power dissipation of the test chip is 56mW, of which the line driver consumes 42mW, and the PE circuit itself dissipates 14mW.

It should be taken into account that channel conditions are different than in the mentioned literature; in the mentioned literature the channels had impedance discontinuities, for example from vias. However, our figure was achieved without any RX equalization (described in the literature as used to cancel reflections), using only TX pre-emphasis. Reflections from impedance discontinuities can be compensated for by adding other, more conventional techniques, especially a decision feedback equalizer (DFE) at the RX side, as is done in [Balan], [Payne], and [Krishna].

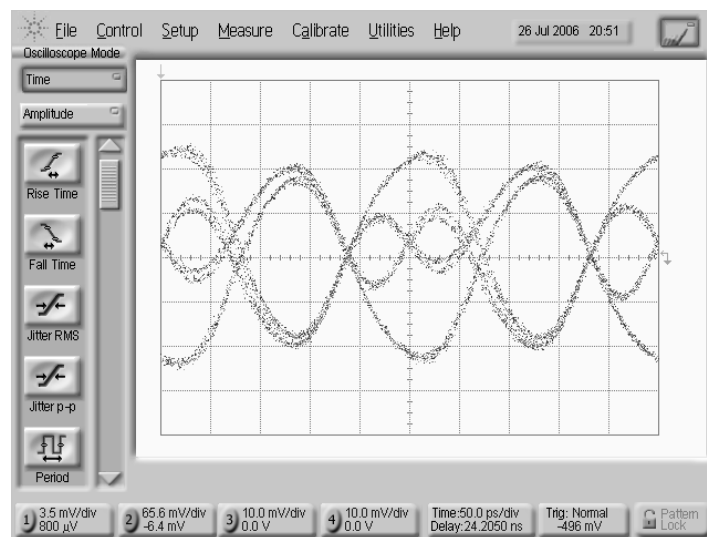


Fig. 30. Output of chip 3 (two bits), showing the extra degree of freedom: the pulse position.




Fig. 31: 5Gb/s over 270cm (106'') FR4, channel loss at Nyquist frequency is 25dB. Top left: input to pcb with no pre-emphasis, right: output from pcb without pre-emphasis. Bottom left: input to pcb with optimum pre-emphasis, right: output from pcb with optimum pre-emphasis. Note that the x-axis length is one bit time. Y-axis, left: 50mV/div, right: 7.5mV/div.

The additional degree of freedom that enables us to place the pulse at an arbitrary position in the bit turned out not to be necessary to compensate for the dielectric loss. The notion that a precursor filter tap is not necessary for PCBs is confirmed in [Ren], where analysis shows that the precursor tap often used for pre-emphasis for PCBs is in fact only necessary because the clock-and-data recovery (CDR) chooses a non-optimum sampling point. A better sampling point can be found, without the use of a precursor tap.

4.6. Conclusions

High-speed data links over copper channels can be effectively equalized using pulse-width modulation (PWM) pre-emphasis. This provides an alternative to the usual 2-tap symbol-spaced FIR. The use of PWM pre-emphasis allows a channel loss at the Nyquist frequency of ~ 30 dB, compared to ~ 20 dB for a 2-tap symbol-spaced FIR filter. The PWM method does not tune the pulse amplitude as for FIR pre-emphasis, but instead exploits timing resolution. This fits neatly with CMOS technology trends toward higher switching speeds and lower voltage headroom.

The PWM duty-cycle determines the shape of the transfer function. Therefore, the filter has only one adjustment parameter. This makes convergence of an algorithm for automatic adaptation straightforward. Spectral analysis illustrates that, compared to a 2-tap FIR filter, the PWM filter transfer function fits better to the copper channel. The equalized channel transfer function is the flattest for PWM pre-emphasis, compared to both half-symbol-spaced and symbol-spaced 2-tap FIR filters. At the Nyquist frequency, the calculated magnitude of

the PWM filter transfer function has a value of 1 while its derivative is non-zero, unlike with the 2-tap SSF filter. The PWM filter transfer function can be seen as a higher order filter with only one adjustment parameter. Time domain simulations with a skin-effect-only channel show that approximately twice the bit rate can be achieved compared to a 2-tap SSF for the same peak distortion.

Three transmitter chips were designed, manufactured and tested as a proof-of-concept. The higher the loss compensation demanded, the lower the allowable jitter in the duty-cycle. The pre-emphasis technique can be implemented on only a small chip area and with low power consumption. Chip 2 has a higher TX swing than chip 1, to compensate for the high channel loss. The TX termination resistance was left out of chip 2 to keep power dissipation low, however it suffers from some data-dependent noise at the TX output. The chips were tested with three types of coaxial cable and one differential cable. The cable lengths were 25 m, 80 m, 130 m and 15 m, respectively. The loss at 2.5GHz is ~30dB for the three coaxial cables and 19dB for the differential cable. A BER $<10^{-12}$ at 5 Gb/s (2-PAM) was achieved with chip 1 for all cable assemblies. Transmission of a 2-PAM 5Gb/s data signal over 25m of low-cost, low-end, standard RG-58U coaxial cable was demonstrated with a BER $<10^{-12}$. This corresponds to a record loss compensation of more than 30dB at the Nyquist frequency of 2.5GHz. On a PCB, a loss compensation of 25dB was achieved using chip 3.

