Chapter 5

Multitap PWM pre-emphasis

5.1. Introduction

In Chapter 4, it was shown that the simple PWM pre-emphasis filter fits very well to cables and PCBs, and can provide up to 30dB of loss compensation. An interesting question is whether the pulse-width filtering technique can be extended to equalize channels with a more complex transfer function. As is well known, with a FIR filter, this is done by increasing the number of taps. Using multitap FIR pre-emphasis, more poles and zeros can be added to the filter transfer function. We can then for example make band pass filters, increase the filter order, and cancel reflections, e.g. on backplane PCBs [Stonick].

The question that we address in this chapter is: would the same be possible for the PWM filter? The goal is to make a transmitter that, unlike a FIR pre-emphasis filter, still switches between only two discrete output voltages, while still offering the degrees of freedom of a multitap FIR filter. This would enable a switching transmitter architecture, with possible advantages in the light of CMOS scaling.

The key question now becomes: can we find a useful function to calculate the pulse-widths? These should be a function of multiple bits, as in the FIR filter, where the amplitudes are a function of multiple bits. If we are able to do so, which extra degrees of freedom do we gain? Are these extra degrees of freedom orthogonal? In this chapter, we explore the functionality that these multitap PWM filters can offer and we analyze their power spectral density (PSD) functions.

The remainder of this chapter is divided into four sections as follows. Section 5.2 starts by introducing the principle of the multitap PWM filter. We first convert the output of a multitap symbol-spaced FIR (SSF) filter into a PWM output by using the ‘same-area approach’. We show with transient simulations that, for an example channel, using such a multitap PWM transmit filter results in nearly identical eye diagrams at the receive side as for a multitap SSF transmit filter. Next, Section 5.3 introduces a method to calculate the autocorrelation function and the PSD of such a multitap PWM filter. Next, in section 5.4, an optimization of the PWM function is made that lowers the number of transitions per second, decreasing the width of the PSD. Finally, conclusions are drawn in section 5.5.

Throughout the sections, the example of a three-tap PWM filter is used. However, the calculation methods can be extended to more than three taps.

5.2. Principle of the multitap PWM filter

In this section, we develop the concept of a multitap PWM pre-emphasis filter. As a start, in subsection 5.2.1, we try to convert the output of a multitap SSF filter into a PWM output by making sure that the area under the transmitted waveforms is equal. This is first done using a PWM pre-emphasis transmitter with three output voltage levels (-1V, 0V and 1V). The second step (in subsection 5.2.2) is a PWM filter with two output voltage levels (-1V and 1V), which would be more straightforward to implement. Transient simulations and eye diagrams are shown to demonstrate the feasibility of the multitap PWM filter. This section follows an intuitive approach; mathematical models are developed later in this chapter.

5.2.1. Same-area approach to translate FIR filter to PWM filter

A first step can be taken as follows. The basic idea for the multitap PWM filter is that the pulse-width is a function of multiple bits, in the same way that the amplitude of a SSF filter is also a function of multiple bits. The question then arises: what is the correct function for this pulse width as a function of multiple bits? We replace the pulse amplitude modulated signal coming out of a SSF filter (which has variable amplitude and a fixed switching interval), with a pulse width modulated (PWM) signal that has a fixed amplitude and a variable switching interval. Our key assumption in making the translation is that if the area under both is the same, in terms of (absolute) voltage integrated over time, the channel response will also be the same. This is a general assumption behind many well-known PWM schemes [Nielsen].

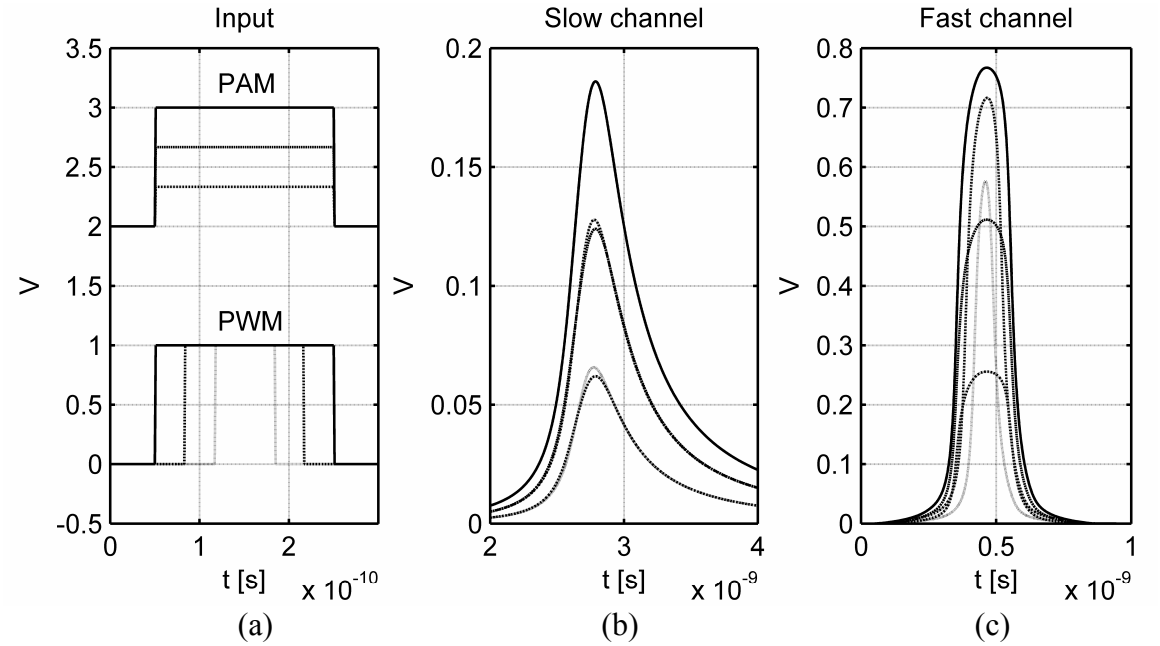

Fig 1. Replacing PAM pulses with PWM pulses. a) PAM and PWM input pulses. The PAM pulses have amplitudes of 0.33V, 0.66V and 1V. The PWM pulses have pulse-widths of 33%, 66% and 100%. T_s=200ps. b) Response of low-bandwidth example channel. c) Response of high-bandwidth example channel.

This idea is shown in Fig. 1. The responses of a low-pass channel to a single PWM pulse and to a single SSF pulse are shown. This is done for two channels: one with a low bandwidth relative to the bit rate, and the other with a high bandwidth relative to the bit rate. We can see that for the first situation, where we need an equalizer, the channel response is indeed almost the same. For the situation with a fast channel, the response to the PWM signal is different from that to the PAM signal. However, equalization is not needed for a channel that is fast enough when compared to the bit rate.

It appears that the area under the pulse, whether it is a PWM or PAM pulse, does indeed determine the response of a 'slow' enough low-pass channel. The channel is sufficiently slow that it 'sees' each of the pulses as an impulse.

We now work with this same-area idea to make a multitap PWM filter that outputs the same area as a multitap SSF filter. To do this we first look at the SSF filter to determine which amplitude values it produces at its output as a function of the input bit stream. Next we replace these in the above fashion with PWM pulses and see whether the channel response is the same. The mathematics are dealt with in a later section; for now we look at the transient simulations to obtain an intuitive understanding.

A theoretical channel was used to perform the time domain simulations. The loss of the channel is monotonously increasing and is approximately 30dB at the Nyquist frequency of 2.5GHz. This theoretical channel is designed to have more pre-cursor ISI than e.g. 25m RG-58CU so that a three-tap filter is necessary (one pre-cursor tap is added to the two-tap FIR). The channel loss is shown in Fig. 2(a), and the response to a 200ps pulse is shown in Fig. 2b.

We now look at the transient simulations. First, in Fig. 3(a), the output of a 3-tap symbol-spaced FIR (SSF) filter is shown. In the same figure, the channel response to it is shown. Note that there are 2^3=8 possible amplitude levels at the filter output. The bit rate is 5Gb/s, which gives a symbol length (and bit length, assuming 2PAM) of T_s=200ps. The tap settings used for this specific example are w={-0.15, 0.55, -0.29}. (These tap settings are optimized for the example channel.)

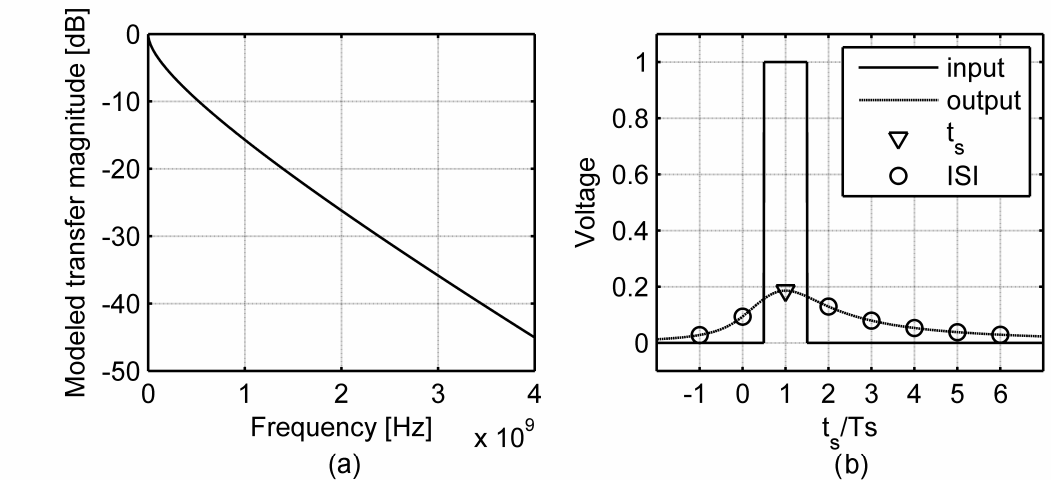

Fig. 2. Theoretical channel used for time domain simulations. (a) Channel loss. (b) pulse response (T_s=200ps).

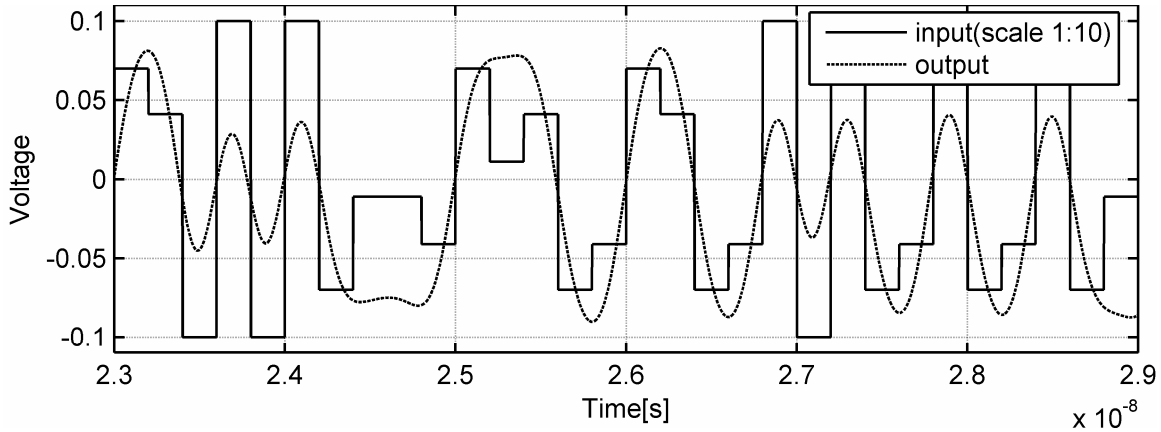

(a)

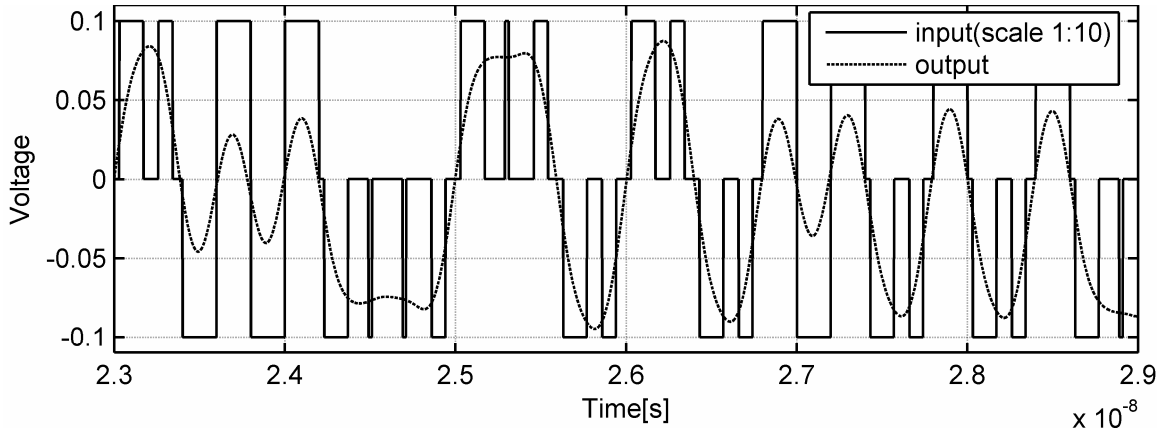

(b)

Fig. 3. Simulated response of example low-pass channel to 3-tap SSF and 3-tap PWM filter, and transmitter outputs (=channel inputs). (a) 3-tap SSF. (b) 3-tap, 3-voltage-level PWM.

It is now explained how the PWM filter output in Fig. 3(b) was made. To translate the amplitude modulated signal output of the SSF filter into a pulse-width modulated signal, and to be able to calculate the autocorrelation later, we first chop the signal up into 'singlets'. A precise mathematical definition of a singlet is given later. For now it is sufficient to know that the duration of a singlet is equal to the symbol duration T_s (=200ps in our example), and that the edges between singlets are at those points in time where the SSF signal is allowed to change in amplitude. For example, looking at Fig. 3(a), we can see that there are five singlets between 25ns and 26ns. The same goes for Fig. 3(b), except that those singlets have a different shape. The amplitude of the SSF output is fixed during one singlet (Fig. 3(a)) because it is a symbol-spaced filter. The pulse-width of the PWM filter is fixed during one singlet (Fig. 3(b)) because it is also a symbol-spaced filter. Note that there are 2^3=8 possible singlets at the filter output of the 3-tap PWM filter.

To calculate the singlets for the PWM filter, we convert each SSF singlet into a PWM singlet with an equal area, as was conceptually shown in Fig. 1. We make sure that the area under the PWM 'spikes' is the same as under the 'flat tops' of the SSF filter. The polarity of the SSF singlet is preserved by letting the PWM pulse go to either 1V or -1V, depending on the polarity of the SSF singlet. This PWM filter transmits three output levels (-1V, 0V and 1V), and therefore we call it 3PWM.

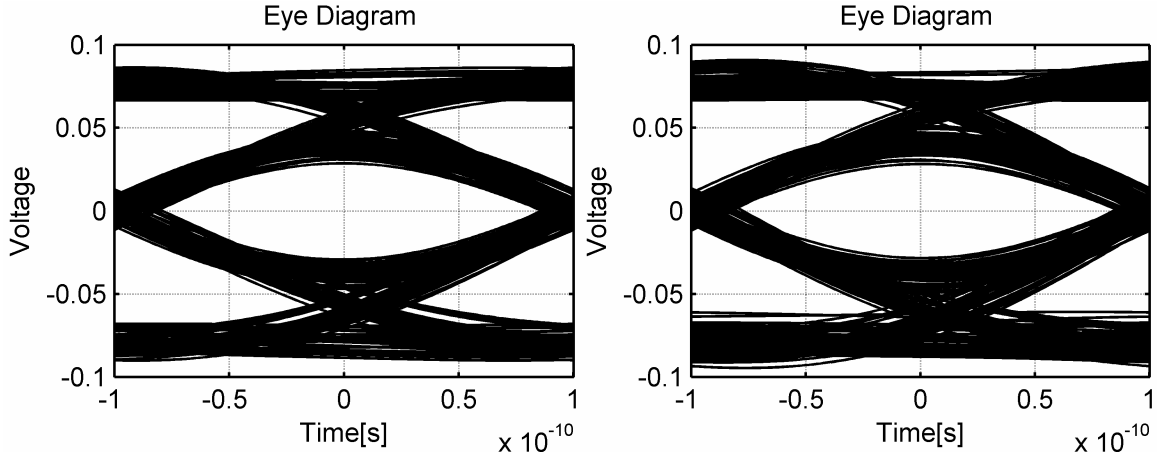

(a) (b)

Fig. 4. Eye diagrams of response of example channel to 3-tap SSF input and to 3-tap PWM input, T_s=200ps. (a) 3-tap SSF. (b) 3-tap, 3-voltage-level PWM.

As hoped, the channel response is almost the same as that to the 3-tap SSF signal, even though the transmitter output has quite a different shape. The eye diagrams in Fig. 4 confirm this. The eye diagram for the 3PWM filter looks practically the same as that for 3-tap SSF.

In conclusion, we have seen that we can obtain a transient cable response nearly identical to that of a SSF filter by replacing the PAM singlets with PWM singlets, using the same-area approach. The channel is successfully equalized. A prerequisite is that the channel has a low-pass characteristic.

5.2.2. 2-level, 3-tap PWM

It might not be so straightforward to implement the 3-voltage-level PWM scheme presented above in a transmitter. For example, in a voltage mode implementation it would require 3 power supply rails. Therefore, it would be more convenient to have two instead of three allowed levels at the output of our PWM transmitter. This might make it easier to implement the circuit with a switching transmitter. Below we show how the number of output levels can be reduced to two without degrading the performance of the multitap PWM equalizer.

During one singlet, we now let the signal first go negative, then positive, and then negative again. We now need another translation from pulse amplitude to pulse-width that gives us two instead of three output levels. To find this new translation we can again use the same-area principle: the total area (positive and negative parts) should be equal to the area of the SSF singlet. We can perform these calculations and provide a new 2-voltage-level scheme. We use the term ‘2PWM’ because it outputs two voltage levels: -1V and +1V. In the next section we present these calculations. Here we first show the results of the transient simulations.

The simulated transient response of the example channel to the 2PWM scheme is given in Fig. 5(b). For comparison, the channel response to the 3PWM scheme is repeated in Fig. 5(a). Both the 2PWM and 3PWM schemes are 3-tap filters; they only differ in the number of output levels. The eye diagrams are compared in Fig 6. As can be seen, the channel responses are almost identical.

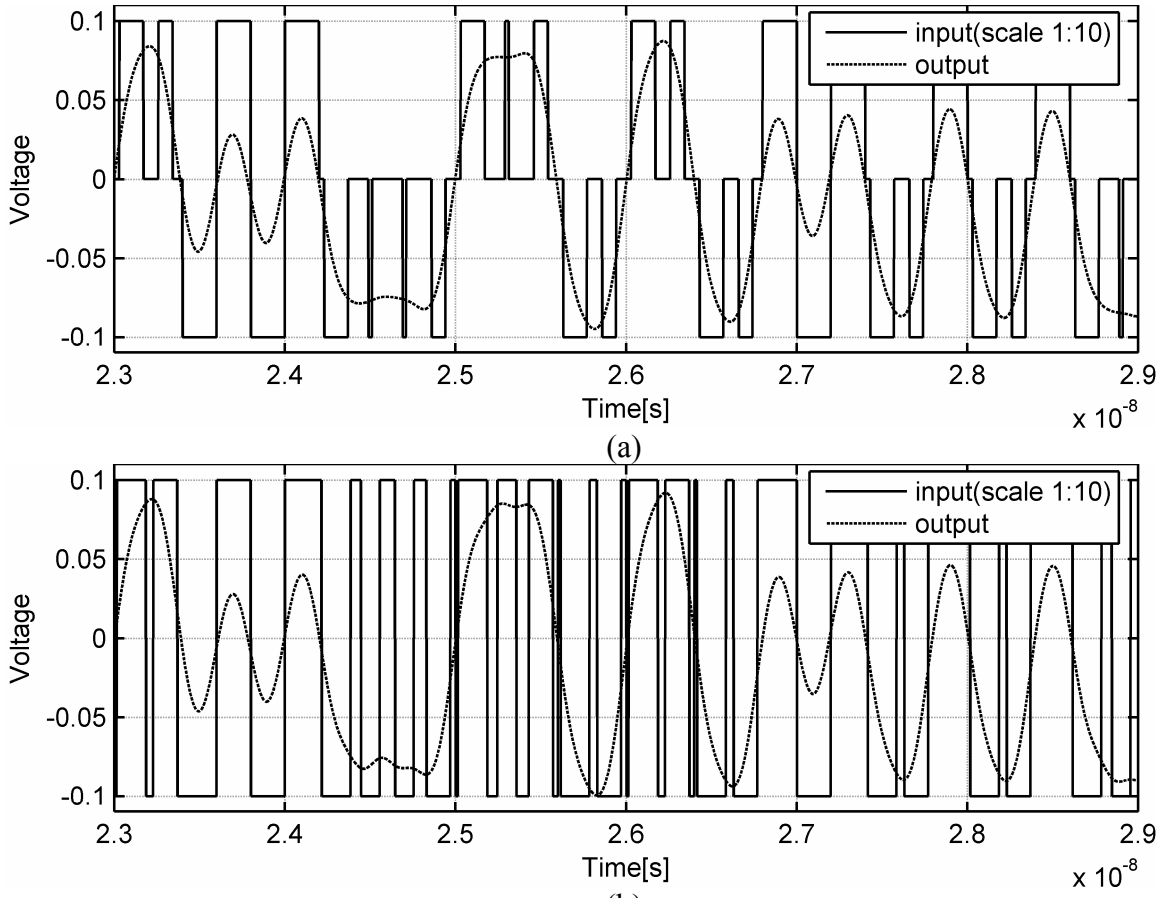

Fig. 5. Simulated transient response of example channel to 3PWM and to 2PWM, and transmitter outputs (=channel inputs), T_s=200ps. Both PWM filters have three taps. (a) 3-level PWM (3PWM). (b) 2-level PWM (2PWM).

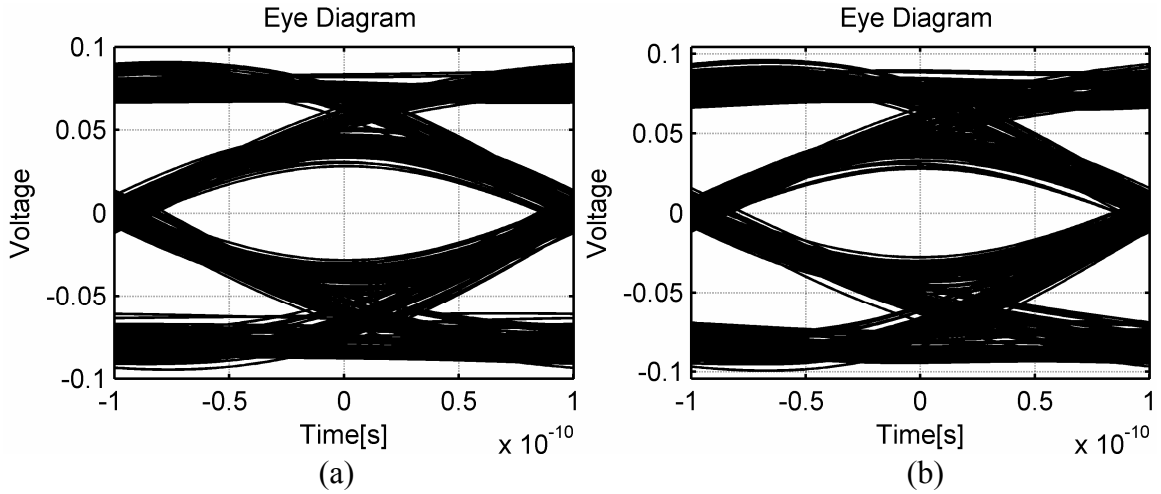

Fig. 6. Eye diagrams of response of example channel to 3PWM input and to 2PWM input, T_s=200ps. (a) 3-tap, 3-level PWM (3PWM) (b) 3-tap, 2-level PWM (2PWM)

We can conclude that the 3-tap SSF, the 3-tap 3-level PWM (3PWM), and 3-tap 2-level PWM (2PWM) schemes all result in almost the same output eye diagram, even though the TX shapes are vastly different. As shown in the time domain simulations, the channel responds in the same way to both the PWM and the SSF pre-emphasis filters due to the channel's low-pass behavior.

Furthermore, we conclude that it was straightforward to find the duty-cycles because they were derived from the singlet heights of the SSF filter using a 'same-area' approach.

5.3. Power spectral density (PSD) functions of the multitap PWM filters

Now that we have established an intuitive time-domain understanding of the multitap PWM filters, we analyze the filter from a frequency-domain point of view. The power spectral density (PSD) functions of the PWM filters are calculated and compared to the PSD of the 3-tap SSF filter.

We can obtain the power spectral density (PSD) in the following way:

1) Calculate the autocorrelation $R_{xx}(\tau)$ (subsection 5.3.1).
2) Take the Fourier transform $F\{\}$ of the autocorrelation to calculate the PSD (subsection 5.3.2):

$$PSD = F\{R_{xx}(\tau)\}. \tag{1}$$

5.3.1. Autocorrelation method

In this subsection, we calculate the autocorrelations of the 3-tap PWM filters. We use the method first to calculate the autocorrelation of the well-known 3-tap SSF filter. This might seem cumbersome, because for the FIR filter we could more easily use (for example) a z-transform, but in this way the result (and the correctness of the method) can be checked against the known PSD function of the 3-tap SSF filter. For calculating the autocorrelation of the PWM filters, we cannot use methods commonly used for FIR filters, because of the time-varying nature of the PWM filter. Therefore we need to devise a calculation method that accepts pulse-width modulated signals. This method uses the 'singlets', as described above. It can be extended to more filter taps. As an additional check, we compare the calculated autocorrelations with statistical simulations.

This subsection is divided as follows. First, in 5.3.1.1, we explain how to calculate the singlets needed for our method, and as an example we calculate the singlet shapes for the 3-tap SSF filter. Next, in 5.3.1.2, it is explained how to calculate the autocorrelation from these singlets. In 5.3.1.3, we calculate the singlet shapes for the 3PWM and 2PWM schemes. Finally, in 5.3.1.4, we use the method to calculate the autocorrelations of 3PWM, 2PWM, 2-tap SSF and polar NRZ.

5.3.1.1. From bits to singlets

As explained above, we cut the transmitted transient data signal for which the autocorrelation is to be calculated into singlets $y(t)$, each of length T_s (=symbol duration and bit duration, because 2PAM is used). The singlets $y(t)$ are defined as follows. They have a width of T_s, are zero outside of the interval $(0,T_s)$ and inside that interval they are a function of time t:

$$y(t) = \begin{cases} 0, & t < 0, \\ f(t), & 0 < t < T_s, \\ 0, & t > T_s. \end{cases} \tag{2}$$

Singlet name	Singlet number	Bit sequence $b_{n-1}...b_{n+1}$
A	$y_1(t)$	-1 1 -1
B	$y_2(t)$	-1 1 1
C	$y_3(t)$	1 1 -1
D	$y_4(t)$	1 1 1
-A	$y_5(t)$	1 -1 1
-B	$y_6(t)$	1 -1 -1
-C	$y_7(t)$	-1 -1 1
-D	$y_8(t)$	-1 -1 -1

Table 1. Singlets and represented bit sequences.

In a 3-tap filter, the singlet shape is a function of three bits. At the output of a 3-tap SSF filter, there are 8 possible singlets, with 8 different voltage amplitudes (2^3=8 combinations), as a function of the three-bit sequence $b_{n-1}...b_{n+1}$, where n is the bit index from the interval $(-\infty,\infty)$. The same applies to the 3-tap PWM filter: there are 8 allowed pulse-widths at the filter output (or 4 pulse-widths and 2 polarities). In the case of the SSF filter, each singlet has a specific height (and fixed width of 1 bit time) and in the case of PWM each singlet has a specific width and polarity (and a fixed amplitude of 1V). The advantage of this method is that the actual singlet shape $f(t)$ does not matter. It can be either a SSF or PWM singlet – or anything else.

We start with four of the eight singlets A, B, C and D and we can then make the other four by flipping them over the t-axis to obtain –A, -B, -C and –D. This symmetry eases the calculation of the autocorrelation. (There are other, non-symmetrical configurations possible but these lead to more complex autocorrelation calculations while not offering a clear benefit.)

In Table 1, the assignment of singlet names to bit sequences is shown. For example, if the bit sequence b_n to be transmitted is equal to {...,-1,1,-1,1,1,...}, then the singlet output of the 3-tap filter would be {...,A,-A,B,...}, where A-D stand for the singlets in Table 1.

To give an example, and for later use, we now calculate the exact shape of the singlets for a 3-tap SSF filter. The mathematical description of the singlets for the 3-tap SSF filter is:

$$y_{SSF}(t) = \begin{cases} 0, & t < 0, \\ \alpha, & 0 < t < T_s, \\ 0, & t > T_s, \end{cases} \tag{3}$$

where the singlet amplitude α is given by:

$$\alpha = w_1 \cdot b_{n+1} + w_2 \cdot b_n + w_3 \cdot b_{n-1}, \tag{4}$$

where w represents the tap weight values.

The calculated singlet shapes for a 3-tap SSF filter are illustrated in Fig. 7 below, for tap settings w={-0.15, 0.55, -0.29}. The amplitude is normalized to +/-1V to allow a same-voltage-headroom comparison to the PWM filter.

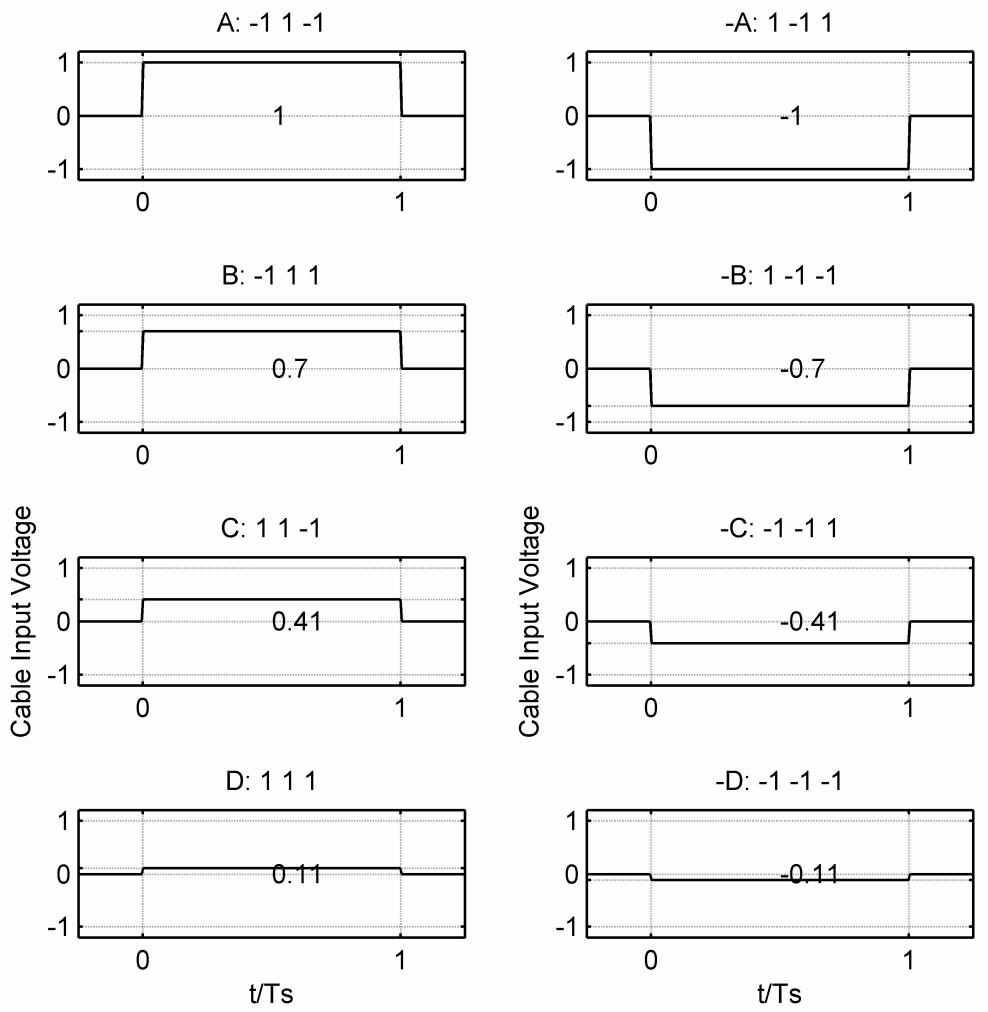

Fig. 7. All possible singlets for a three tap SSF filter, as a function of the 3-bit sequence. In the middle of each singlet, α is printed.

5.3.1.2. Autocorrelation from singlets

As is well known, the formula for calculating the autocorrelation of a signal *x(t)* as a function of time shift τ is [Couch]:

$$R_{xx}(\tau) = \lim_{T\to\infty} \frac{1}{2T} \int_{-T}^{T} x(t)x(t+\tau)dt \,. \tag{5}$$

We assume that the signal is ergodic; its time-average is equal to its ensemble average. The autocorrelation for our 3-tap symbol-spaced filter has a total width of $6T_s$; it is zero outside of the interval for τ of $(-3T_s,3T_s)$. This is because the singlets have a width of T_s and the filter is a 3-tap symbol-spaced filter.

Using the singlet shapes A-D together with the probabilities of their occurrence, we can calculate the autocorrelation for the filtered data stream. Because the singlets are a function of three bits (and each singlet is only one bit wide) not all combinations of singlets are allowed to follow each other. For example, the singlet corresponding to bit sequence {-1,1,-1} ('A') can only be followed by a singlet corresponding to {1,-1,*}, where * denotes either +1 or -1

('-A' or '-B'); e.g. 'D' (corresponding to {1,1,1}) is not possible. To understand this, shift the first bit sequence one bit to the left.

To obtain the autocorrelation, we need to check all the possible combinations and their probability of occurrence. Then, to calculate the autocorrelation using the singlets, we split the autocorrelation function into five (partially overlapping) parts:
-the first part ($R_2(\tau+2T_s)$) around a time shift of $-2T_s$
-the second part ($R_1(\tau+T_s)$) around a time shift of $-T_s$
-the third part ($R_0(\tau)$) around a zero time shift
-the fourth part ($R_1(-\tau-T_s)$) around a time shift of T_s
-and the fifth part ($R_2(-\tau-2T_s)$) around a time shift of $2T_s$.

Each of the five parts has a width of $2T_s$ because the width of each singlet is T_s. The five parts are illustrated in Fig. 8 for the 3-tap SSF filter.

The autocorrelation $R_{xx}(\tau)$ is symmetrical around τ=0, and is composed of the three components $R_0(\tau)$, $R_1(\tau)$ and $R_2(\tau)$ as follows:

$$R_{xx}(\tau) = R_2(\tau + 2T_s) + R_1(\tau + T_s) + R_0(\tau) + R_1(-\tau - T_s) + R_2(-\tau - 2T_s). \tag{6}$$

The components $R_0(\tau)$, $R_1(\tau)$ and $R_2(\tau)$ are defined as follows. First, $R_0(\tau)$ is defined as:

$$R_0(\tau) = \frac{1}{8}\sum_{i=1}^{8}\left(\frac{1}{T_s}\int_0^{T_s} y_i(t)y_i(t-\tau)dt\right), \tag{7}$$

where $y_x(t)$ are the singlets as defined in Table 1.

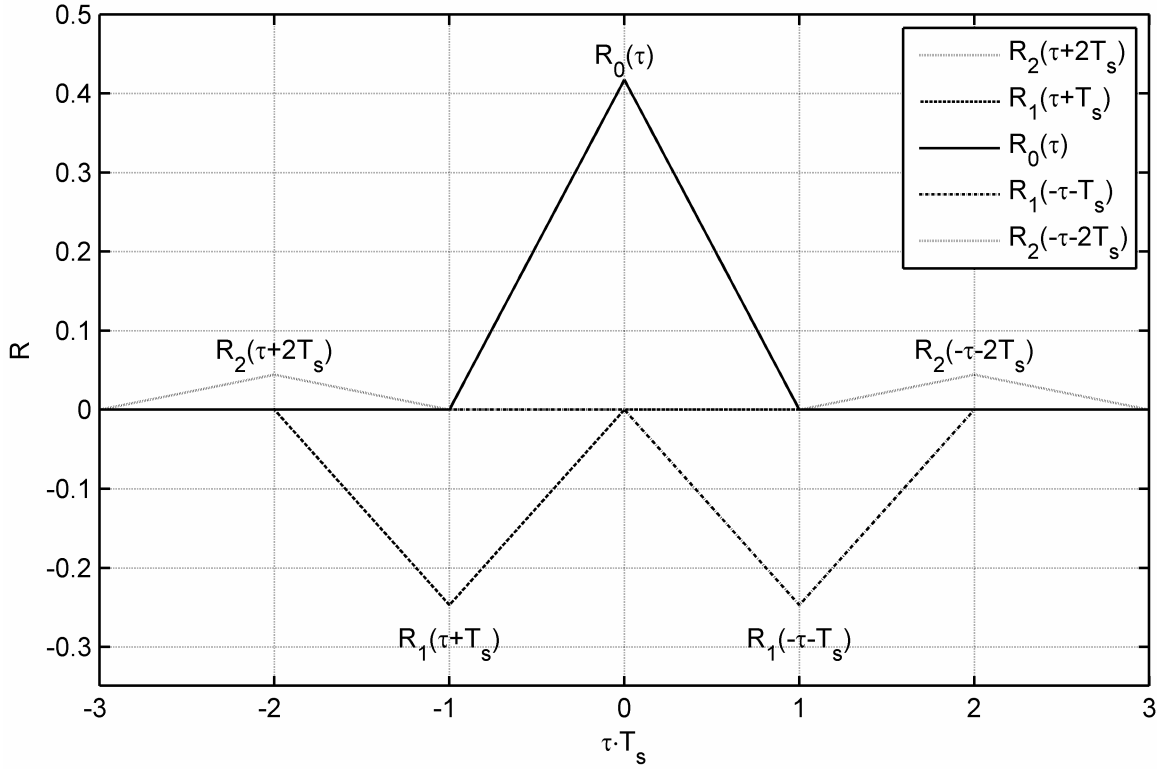

Fig. 8. Autocorrelation components $R_0(\tau)$, $R_1(\tau)$ and $R_2(\tau)$ for the 3-tap SSF filter.

Next, $R_1(\tau)$ is defined as:

$$R_1(\tau) = \frac{1}{16}\sum_{i=1}^{8}\left(\sum_{j=1}^{2}\left(\frac{1}{T_s}\int_0^{T_s} y_i(t)y_{v_1(i,j)}(t-\tau)dt\right)\right), \tag{8}$$

where $y_{v1(i,j)}$ is the jth allowed singlet to follow singlet y_i. Appendix C lists the values for $y_{v1(i,j)}$. The values for j are chosen from the interval [1,2] because, as explained above, for each of the eight (i=1..8) singlets, only two singlets are allowed to follow it. This is because the bit sequence is shifted to the left by one bit. There are 8·2=16 allowable combinations in total.

Finally, $R_2(\tau)$ is defined as:

$$R_2(\tau) = \frac{1}{32}\sum_{i=1}^{8}\left(\sum_{j=1}^{4}\left(\frac{1}{T_s}\int_0^{T_s} y_i(t)y_{v_2(i,j)}(t-\tau)dt\right)\right), \tag{9}$$

where $y_{v2(i,j)}$ is the jth allowed singlet to follow singlet y_i, $2T_s$ after the start of this first singlet. Appendix C lists the values for $y_{v2(i,j)}$. The values for j are chosen from the interval [1,4] because the bit sequence is shifted to the left by two bits, so that for each of the eight (i=1..8) singlets there are now four singlets that are allowed to follow it, $2T_s$ after its start. There are 8·4=32 allowable combinations in total.

The above equations for $R_0(\tau)$, $R_1(\tau)$ and $R_2(\tau)$ can be generalized into one equation for $R_k(\tau)$:

$$R_k(\tau) = \frac{1}{8\cdot 2^k}\sum_{i=1}^{8}\left(\sum_{j=1}^{2^k}\left(\frac{1}{T_s}\int_0^{T_s} y_i(t)y_{v_k(i,j)}(t-\tau)dt\right)\right). \tag{10}$$

In Appendix C, the complete calculation is given. The results are reproduced here. The components $R_0(\tau)$, $R_1(\tau)$ and $R_2(\tau)$ are given by, respectively:

$$R_0(\tau) = \frac{1}{4}\left(R_{AA} + R_{BB} + R_{CC} + R_{DD}\right), \tag{11}$$

$$R_1(\tau) = \frac{1}{8}\left(-R_{AA} - R_{AB} + R_{BC} + R_{BD} - R_{CA} - R_{CB} + R_{DC} + R_{DD}\right), \tag{12}$$

and

$$\begin{aligned} R_2(\tau) = \frac{1}{16}(&R_{AA} + R_{AB} - R_{AC} - R_{AD} - R_{BA} - R_{BB} + R_{BC} + R_{BD} \\ &+ R_{CA} + R_{CB} - R_{CC} - R_{CD} - R_{DA} - R_{DB} + R_{DC} + R_{DD}). \end{aligned} \tag{13}$$

In the above equations, a shorthand notation is used, e.g. R_{AB} means:

$$R_{AB} = \int_0^{T_s} y_1(t)y_2(t+\tau)dt\,, \tag{14}$$

where $y_1(t)$ is the singlet 'A' and $y_2(t)$ is the singlet 'B' (see Table 1).

5.3.1.3. Calculation of singlets for the PWM filters

Now, a mathematical description is given of the singlets A,B,C and D for the 3PWM and 2PWM schemes.

Singlets for 3PWM

The mathematical description of the 3PWM singlets is

$$y_{3PWM}(t) = \begin{cases} 0, & t < (1-|\alpha|)T_s/2, \\ sign(\alpha), & (1-|\alpha|)T_s/2 < t < (1+|\alpha|)T_s/2, \\ 0, & t > (1+|\alpha|)T_s/2. \end{cases} \tag{15}$$

We can use the same variable α as for the 3-tap SSF filter, but now it denotes the pulse width and pulse polarity instead of the pulse height. It falls into the interval of [-1,1], because of the normalization that was done. As previously stated, the parameter α contains information about the pulse-width and its polarity. The absolute value of α, which is a number from the interval [0,1], corresponds to the pulse-width (0-100%) and the sign of α corresponds to the pulse polarity.

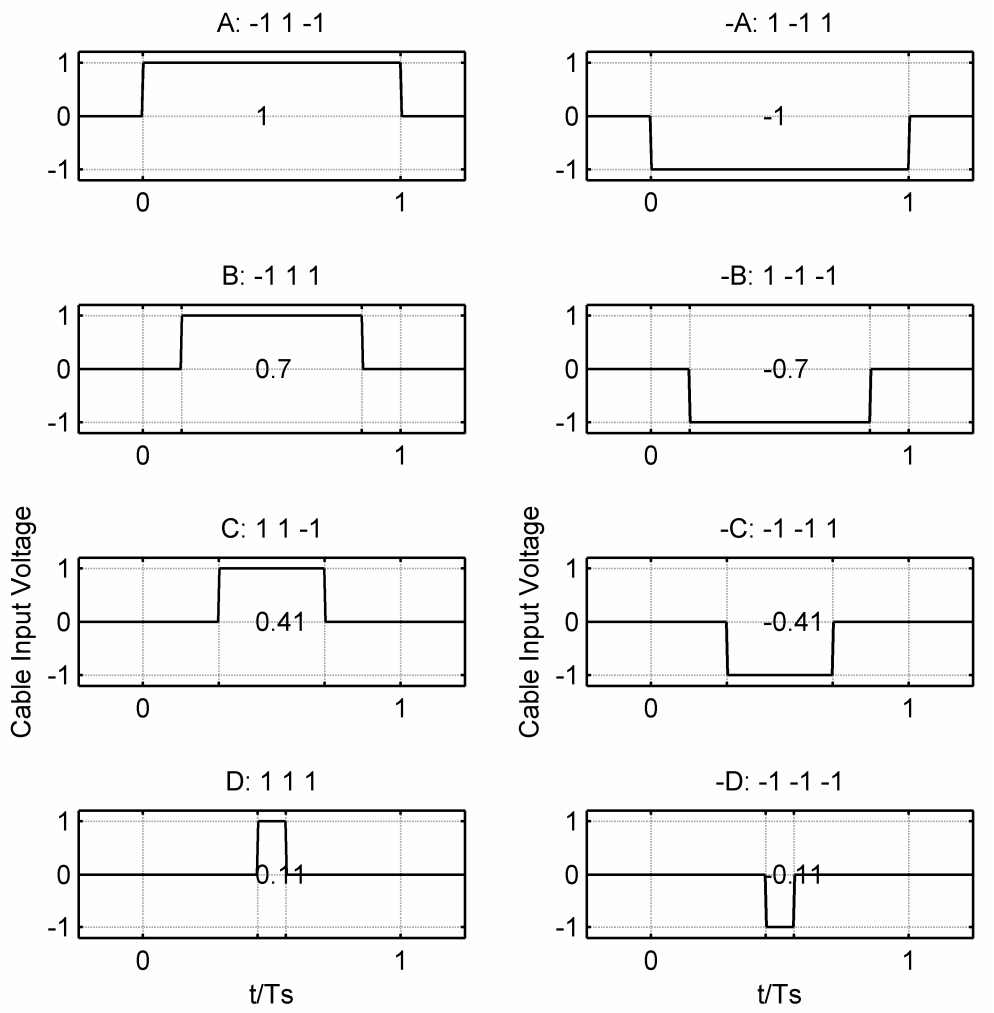

Fig. 9. All possible output singlet shapes for 3-tap 3-voltage-level PWM pre-emphasis (3PWM). α is printed in the middle of each singlet.

In Fig. 9, the singlets for 3PWM are shown. It can easily be proven that the areas under the 3PWM and SSF singlets (see Fig. 7) are equal.

Singlets for 2PWM

As mentioned earlier, we would like to reduce the number of output levels of our PWM transmitter to two levels (1V and -1V) instead of three. The total area is now a sum of the positive and negative parts, as described above. To keep the area equal to the SSF area, the pulse-width needs to be different from that of the 3PWM scheme. Again, we use a single variable to denote both the pulse-width and its polarity. The absolute value of this parameter ψ is the pulse-width and its sign is the pulse polarity. This scheme is termed 2PWM because it has two output voltage levels.

The definition of the 2PWM singlets is as follows:

$$y_{2PWM}(t)=\begin{cases}0, & t<0,\\ -sign(\psi), & 0<t<(1-|\psi|)T_s/2,\\ sign(\psi), & (1-|\psi|)T_s/2<t<(1+|\psi|)T_s/2,\\ -sign(\psi), & (1+|\psi|)T_s/2<t<T_s,\\ 0, & t>T_s,\end{cases} \tag{16}$$

where sign(ψ)=sign(α), per definition, and ψ is chosen from the intervals [-1,-0.5] and [0.5,1]. Values in the interval (-0.5,0.5) are not allowable. Using these intervals, we can achieve an equal summed area as that in the 3PWM scheme, as is shown below.

The total area under the new pulse is equal to:

$$A_{n1}=\int_{-\infty}^{\infty_s} y_{n1}(t)dt=sign(\psi)\left(-\int_{0}^{(1-|\psi|)T_s/2}1\ dt+\int_{(1-|\psi|)T_s/2}^{(1+|\psi|)T_s/2}1\ dt-\int_{(1+|\psi|)T_s/2}^{T_s}1\ dt\right)=(2\psi-sign(\psi))T_s\,. \tag{17}$$

We calculate how to translate from α to ψ, using the same-area approach:

$$A_{n1}=A_{3PWM}\,, \tag{18}$$

$$(2\psi-sign(\psi))T_s=\alpha T_s\,, \tag{19}$$

$$|\psi|=(|\alpha|+1)/2\,, \tag{20}$$

(and sign(ψ)=sign(α), per definition, as said above). Note how the interval for α of [-1,1] indeed leads to an interval for ψ of [-1,-0.5] and [0.5,1]. As soon as α<0, the 2PWM pulse is flipped around the t-axis, because of its sign definition, maintaining the symmetry between A,B,C,D and -A,-B,-C,-D singlets. The resulting singlet shapes are shown in Fig. 10.

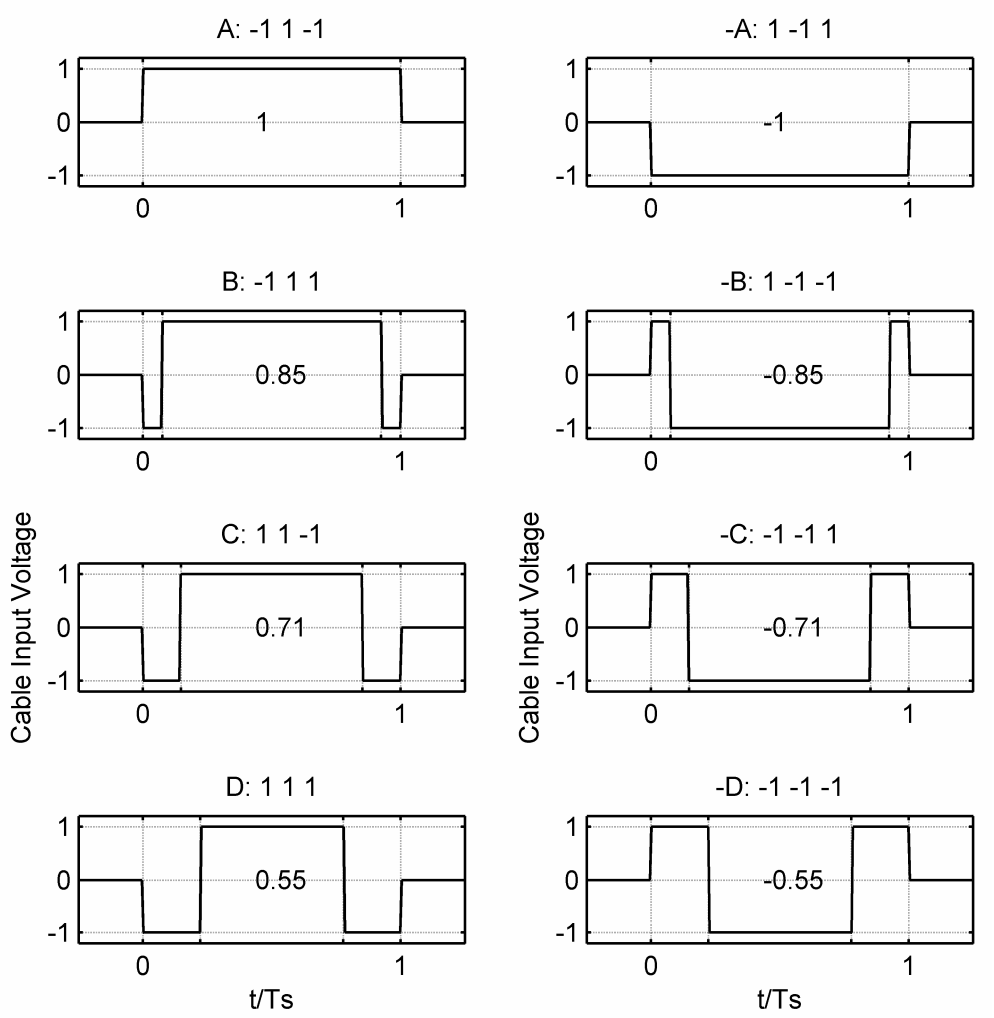

Fig. 10. All possible output singlet shapes for 3-tap, 2-voltage-level PWM pre-emphasis (2PWM). ψ is printed in the middle of each singlet.

The translation is ambiguous at only one point, which is at α=0 (where α has no sign). At this point, there are two possible singlet shapes for the 2PWM scheme, each with a 50% duty-cycle. The first is negative, positive, and then negative (ψ=0.5), and the second is exactly the opposite: first positive, then negative, then positive (ψ=-0.5). We define here that in that case (which is rare), the sign of ψ be equal to the sign of the most significant filter tap (the tap with the maximum absolute value).

5.3.1.4. Autocorrelation results

Knowing the description of all the singlet shapes, we can fill them into the previously given autocorrelation formula and calculate the autocorrelations. The method is used to calculate the autocorrelations for the 2PWM and 3PWM schemes. For comparison, also the autocorrelation for the 3-tap SSF filter output and for polar NRZ (2PAM, unequalized) is calculated.

To check the analytical calculations, we also compare the calculated autocorrelations to the results from a statistical simulation with 10,000 random symbols.

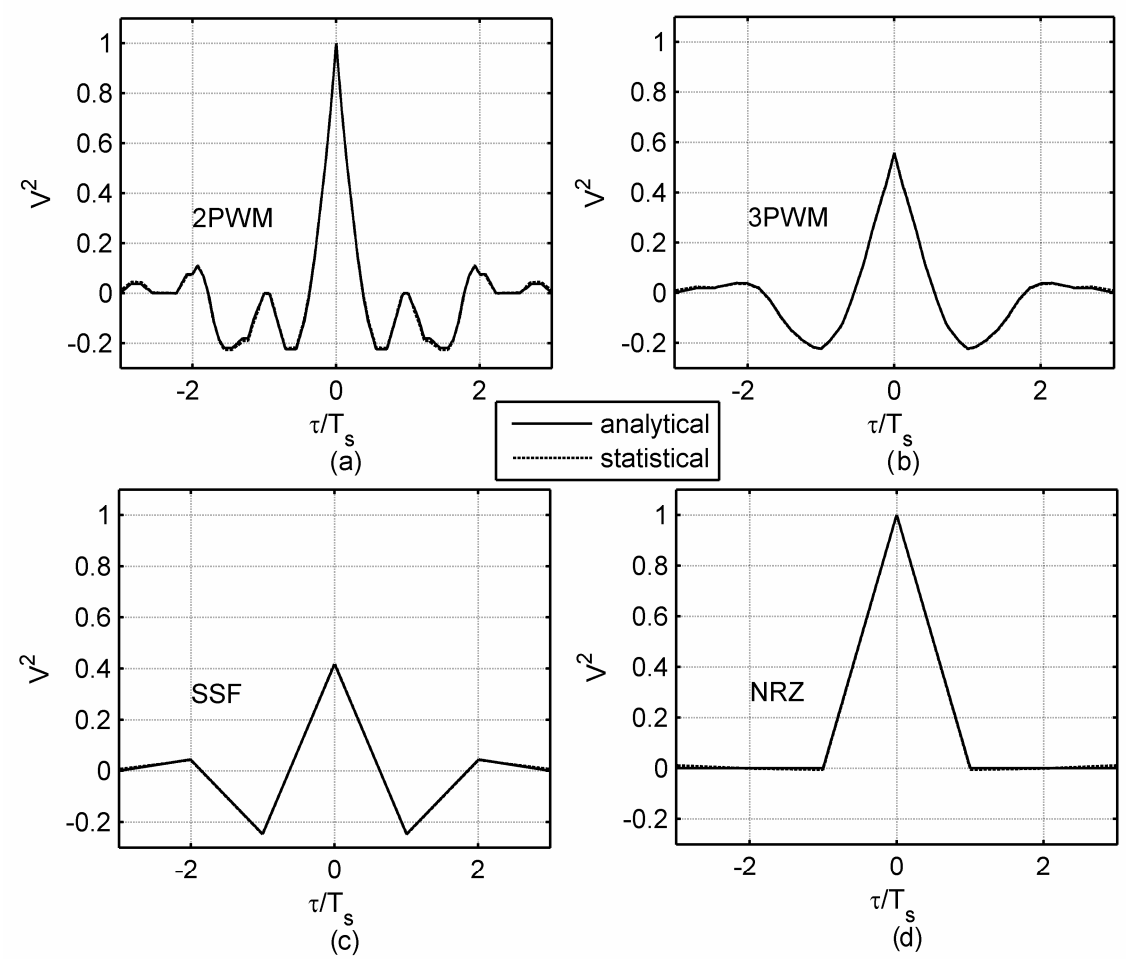

Fig. 11. Autocorrelations (from analytical and from statistical simulations). In some plots only the 'analytical' line can be seen because the 'statistical' line is right below it. (a) 2PWM (2-level, 3-tap). (b) 3PWM (3-level, 3-tap). (c) 3-tap SSF. (d) Polar NRZ.

The results are shown in Fig. 11. The individual plots show the following autocorrelations:

a) 2PWM filter,
b) 3PWM filter,
c) 3-tap SSF filter,
d) polar NRZ.

The signals are normalized to a +/-1V supply voltage headroom. We see that 3PWM has an autocorrelation that very closely matches that of 3-tap SSF. The 2PWM autocorrelation shows more high-frequency components, due to the higher switching frequency.

5.3.2. PSD functions

In this subsection we describe the calculation of the PSD functions for the multitap PWM filters. We now have all the information to do this, using the autocorrelations calculated above.

As described above, we use the Fourier transform to obtain the PSDs from the autocorrelations. To check the analytical calculations, we again compare the calculated PSDs to the results from a statistical simulation with 10,000 random symbols.

Fig 12 shows the PSDs for the following schemes:
(a) 2PWM filter
(b) 3PWM filter
(c) 3-tap SSF filter
(d) polar NRZ (2PAM, unequalized).

A smoothing filter was used on the statistically calculated PSDs. Small differences between the statistical simulations and the analytically calculated results are due to the statistical nature of the simulations: a non-infinite number of symbols need to be used to limit the simulation time.

The baseband part of the PSD (up to the Nyquist frequency f_N, which is at 0.5 on the x-axis) is very similar to that of the SSF filter. As expected, the 2PWM scheme has a higher power spectral density at frequencies above the baseband than the 3PWM scheme, due to its higher switching frequency. This is a well-known side effect of PWM. As long as the channel has a low-pass transfer function, the high-frequency part of the spectrum will be filtered out by the channel. Also a low-pass filter inside the transmitter could be used to filter out those components. A system designer should weigh up this side-effect of PWM against the advantage (for high-speed low-voltage CMOS technologies) of using a transmission equalizer that only needs to switch between two voltages.

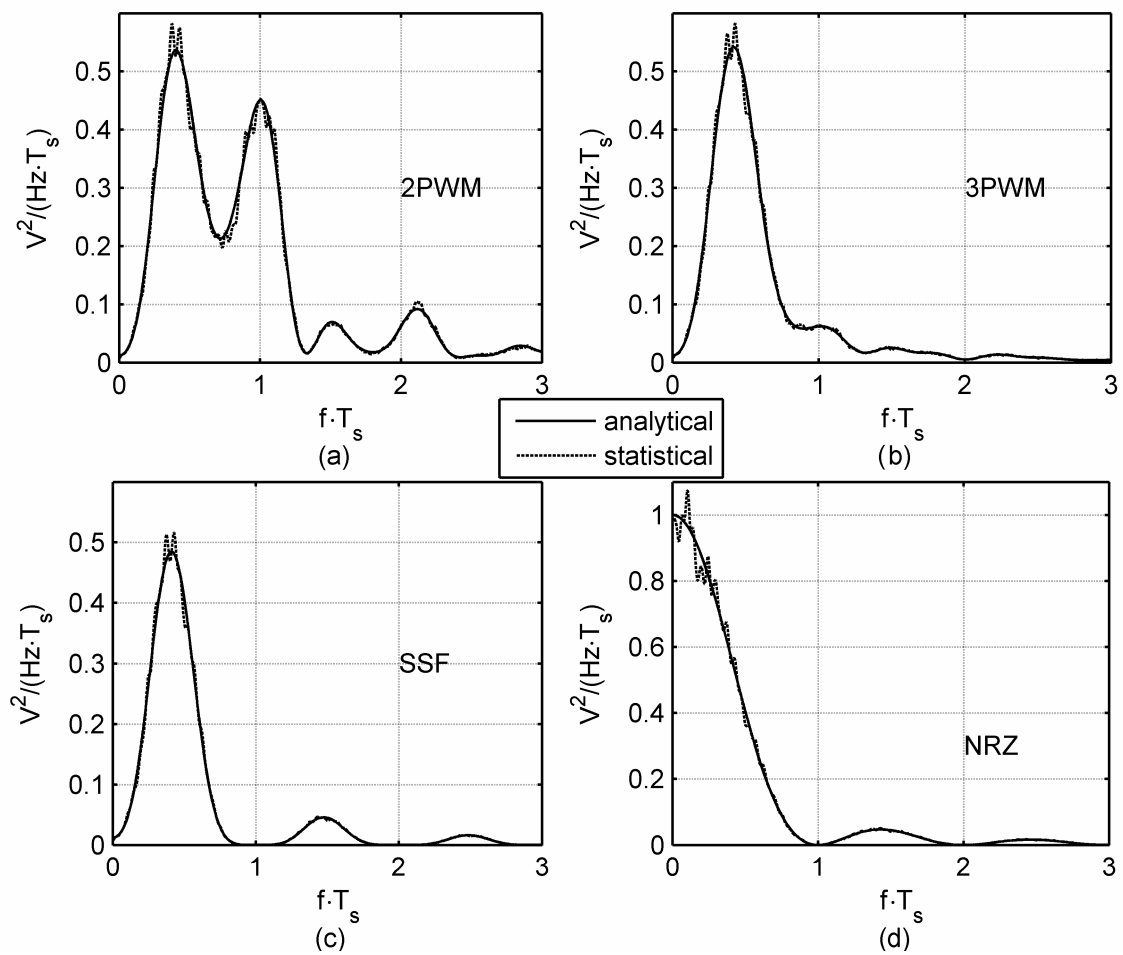

Fig. 12. PSDs (from analytical and from statistical simulations). Smooth lines are for 'analytical', jagged lines are for 'statistical'. (a) 2PWM (2 level, 3 tap). (b) 3PWM (3 level, 3 tap). (c) 3-tap SSF. (d) Polar NRZ.

5.4. Reducing the number of transitions in the 2PWM scheme

A possible disadvantage of the 2PWM scheme, presented in the previous sections, is its large number of transitions – and the resulting large width of the PSD. The output voltage switches twice in each singlet. There might be unnecessary transitions in the scheme, adding power in higher frequencies that are outside of the signal band. We want to see whether the amount of transitions can be decreased by moving to another scheme. In this final section, we decrease the number of transitions per second of the PWM multitap filter by adjusting the 2PWM scheme that was presented above. The goal is to keep the eye opening the same, while making the PSD narrower.

We still have an extra degree of freedom with the 2PWM scheme. This degree of freedom is the placing of the pulse in the singlet: whether the pulse is placed in the middle of the bit time, or at the left or right side. Up until now, our PWM scheme placed the pulse in the middle. However, looking back at the 1-tap PWM pulse shape from Chapter 4, we see that this pulse first has a positive voltage, and then a negative voltage. So it is placed on the left side (which is the best for a channel having predominantly post-cursor ISI). There is only one transition inside the singlet instead of two in the 3PWM and 2PWM schemes. To reduce the number of transitions, we can change our 2PWM scheme. We choose to move the pulse to the left of the singlet. This new scheme is illustrated in Fig. 13, dubbed the 2PWM-L scheme ('L' from 'left').

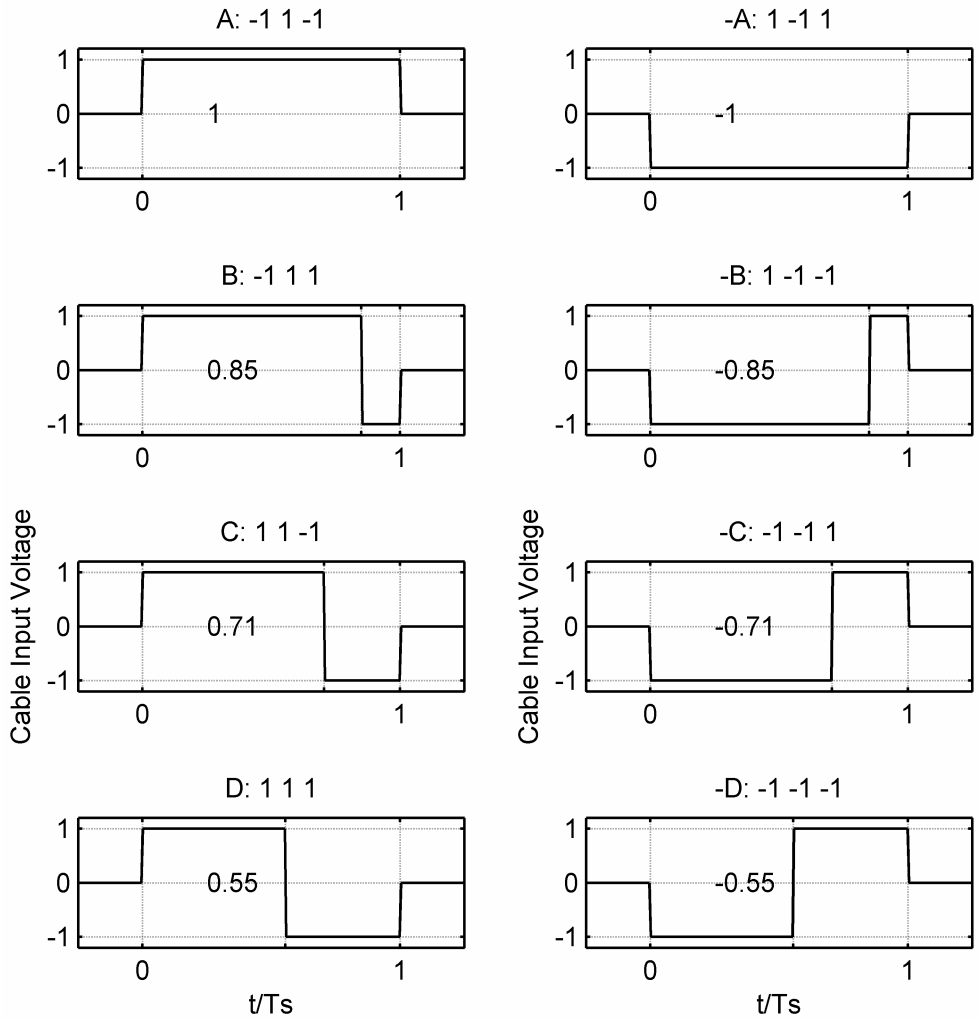

Fig. 13. Singlet shapes for 2PWM-L scheme.

The 2PWM-L singlet is defined as follows:

$$y_{2PWM-L}(t) = \begin{cases} 0, & t < 0, \\ sign(\psi), & t < |\psi| T_s, \\ -sign(\psi), & t > |\psi| T_s, \\ 0, & t > T_s, \end{cases} \tag{21}$$

where, at first, the same parameter values for ψ (denoting pulse-width and polarity) as in the 2PWM scheme are used.

Judging by the transient simulation and eye diagram in Fig. 14, we can easily see that the 2PWM-L scheme does not fit well to the channel. (The same channel is used as in the simulations with the previously discussed schemes.) Shifting the pulses to the left side of the singlets has changed the behavior of the filter in such a way that it is no longer enough to derive the pulse width simply from the SSF singlet height. The phase relations between the singlets are changed compared to the previous schemes, because the length of time that they are moved to the left depends on their width. We need to compensate for this by using different pulse-widths.

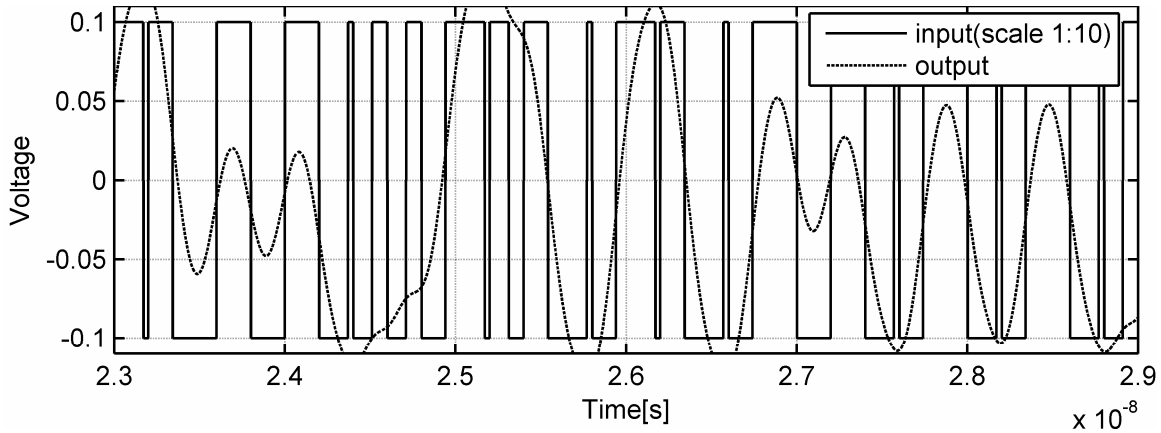

(a)

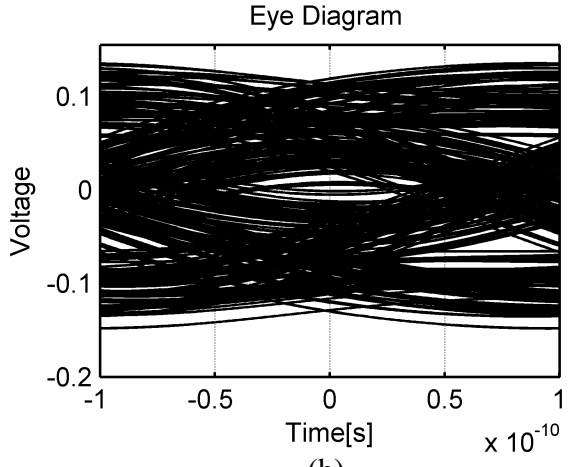

(b)

Fig. 14. Response of example channel to 2PWM-L scheme, T_s=200ps. (a) Time domain. (b) Eye diagram.

After some trial-and-error simulation runs, it was discovered that swapping singlets B&C (and –B&-C) provided a solution to the problem. This is equivalent to calculating the singlet pulse-widths from the SSF filter taps with taps 1 and 3 swapped. This new scheme incorporating the swap, is denoted as 2PWM-LBC. Mathematically, we define this scheme as follows. In Eq. 21, ψ is replaced by ψ_{alt}, where ψ_{alt} is given by:

$$|\psi_{alt}| = \left(|\alpha_{alt}| + 1\right)/2, \tag{22}$$

and α_{alt} is given by (compare Eq. 4):

$$\alpha_{alt} = w_1 \cdot b_{n-1} + w_2 \cdot b_n + w_3 \cdot b_{n+1}, \tag{23}$$

where b is the bit sequence, and w represents the tap weight values. This scheme is illustrated in Fig. 15.

In Fig. 16(b) the simulated time domain response of the channel to the 2PWM-LBC scheme is shown. For comparison, the response to the 2PWM scheme is repeated in Fig 16(a). Indeed, an open eye is seen and the number of transitions is lower (at 43) compared to the 2PWM scheme (at 51), just as we desired. Fig 17 compares the eye diagrams. The goal of keeping the eye opening the same is achieved.

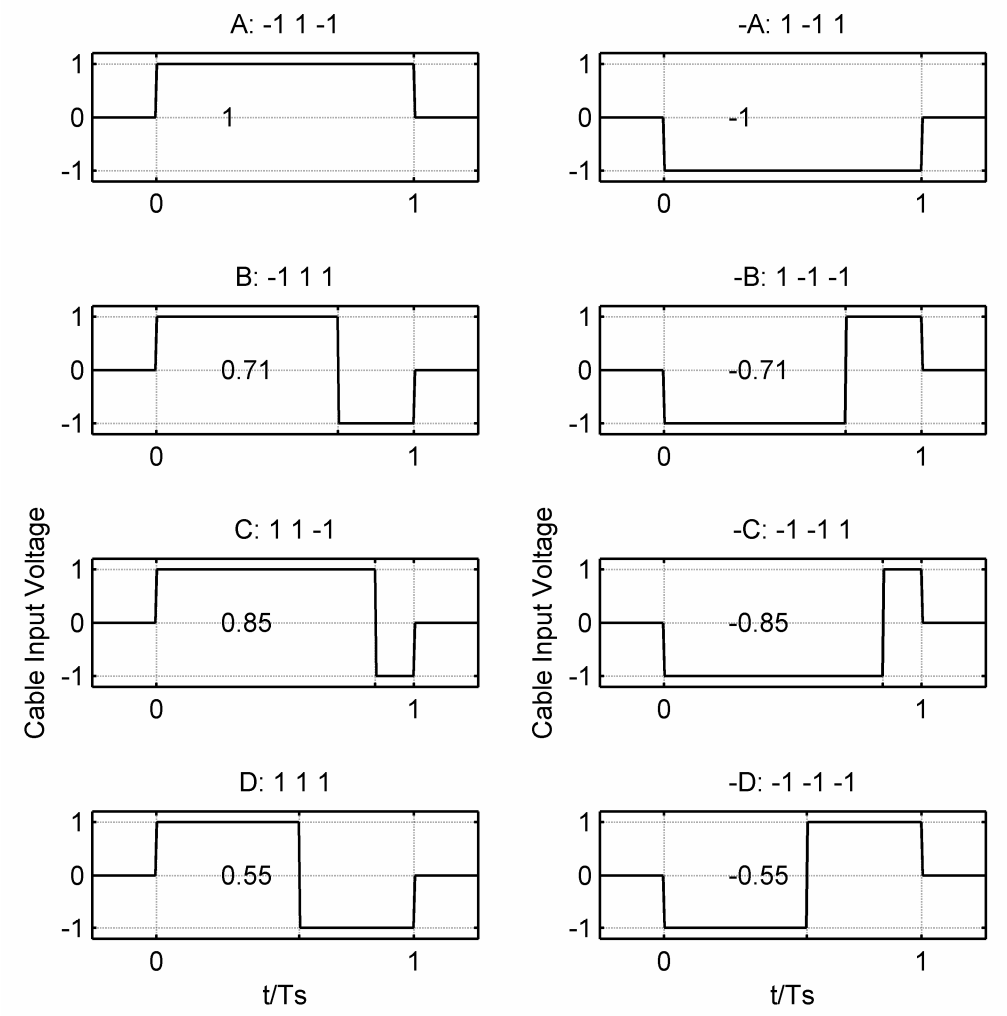

Fig. 15. Singlet shapes for the 2PWM-LBC scheme.

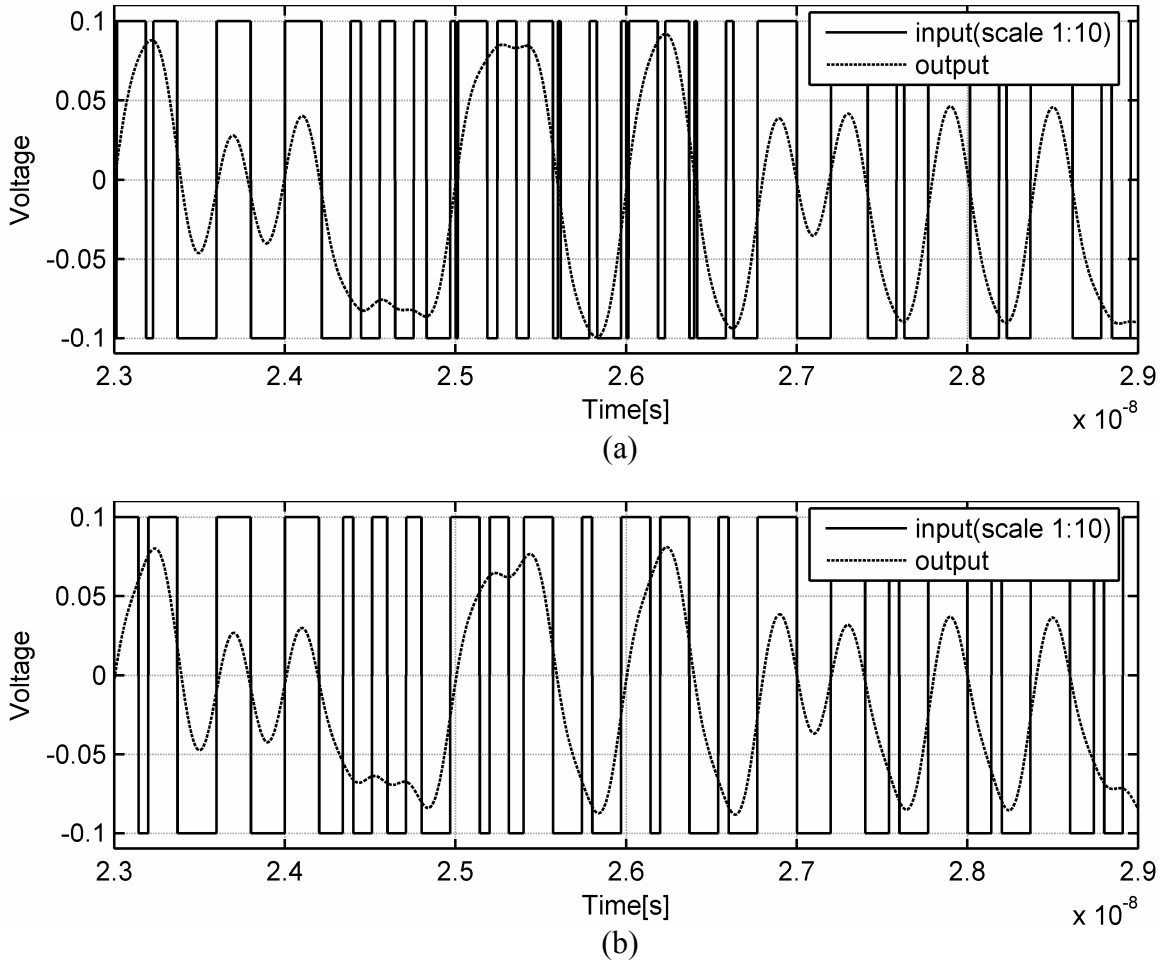

Fig. 16. Simulated transient response of example channel to 2PWM and to 2PWM-LBC, and transmitter outputs (=channel inputs), T_s=200ps. Both PWM filters have three taps and two output voltage levels. (a) 2PWM (51 transitions). (b) 2PWM-LBC (43 transitions).

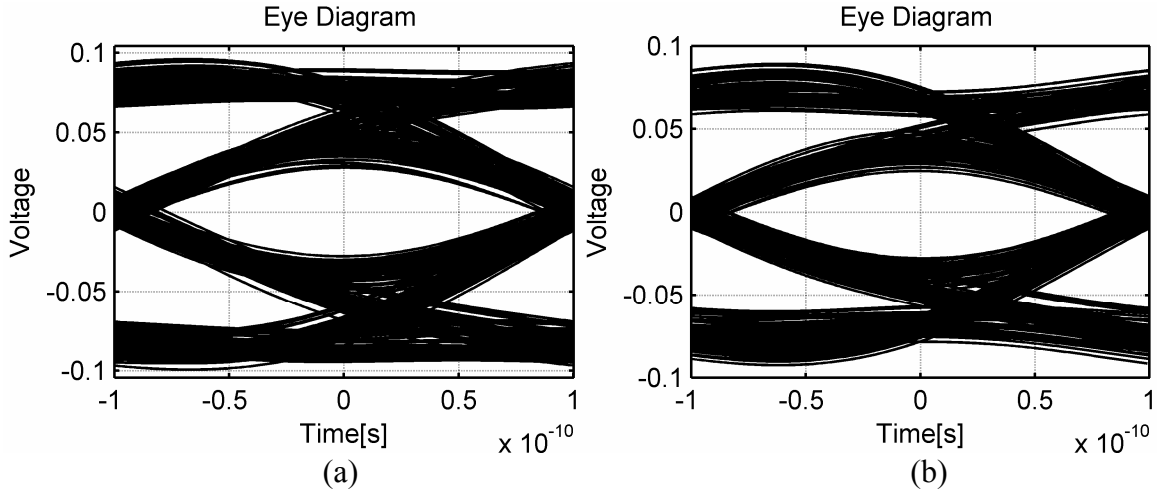

Fig. 17. Eye diagrams of response of example channel to 2PWM input and to 2PWM-LBC input, T_s=200ps. (a) 2PWM. (b) 2PWM-LBC.

Next, we would like to know whether the PSD of the 2PWM-LBC scheme is narrower than that of the 2PWM scheme. Fig. 18 shows the autocorrelations and PSDs of the 2PWM-L and 2PWM-LBC schemes described above. In Fig. 19, the PSDs are compared to those of the 3-tap SSF and the 2PWM scheme. In Fig. 19, we can clearly see that the PSD for 2PWM-L is indeed very different from the others (in the baseband). The PSD for the 2PWM-LBC scheme (with B&C swapped) again matches the PSD of 3-tap SSF more closely. The second lobe for the 2PWM-LBC scheme has moved to the left a little bit (the PSD is narrower) compared to the 2PWM scheme, indicating that the switching frequency is lower. However, the power spectral density around $0.7\text{Hz}\cdot T_s$ has increased.

As an aside, we note that the PWM scheme described in Chapter 4 is in fact a subset of the 2PWM-L and 2PWM-LBC schemes. We can obtain the previous PWM scheme by making the pulse-widths for singlets {A,B,C,D} in the 2PWM-L scheme equal to $\{d,d,d,d\}$, where d is the duty-cycle of the PWM scheme from Chapter 4.

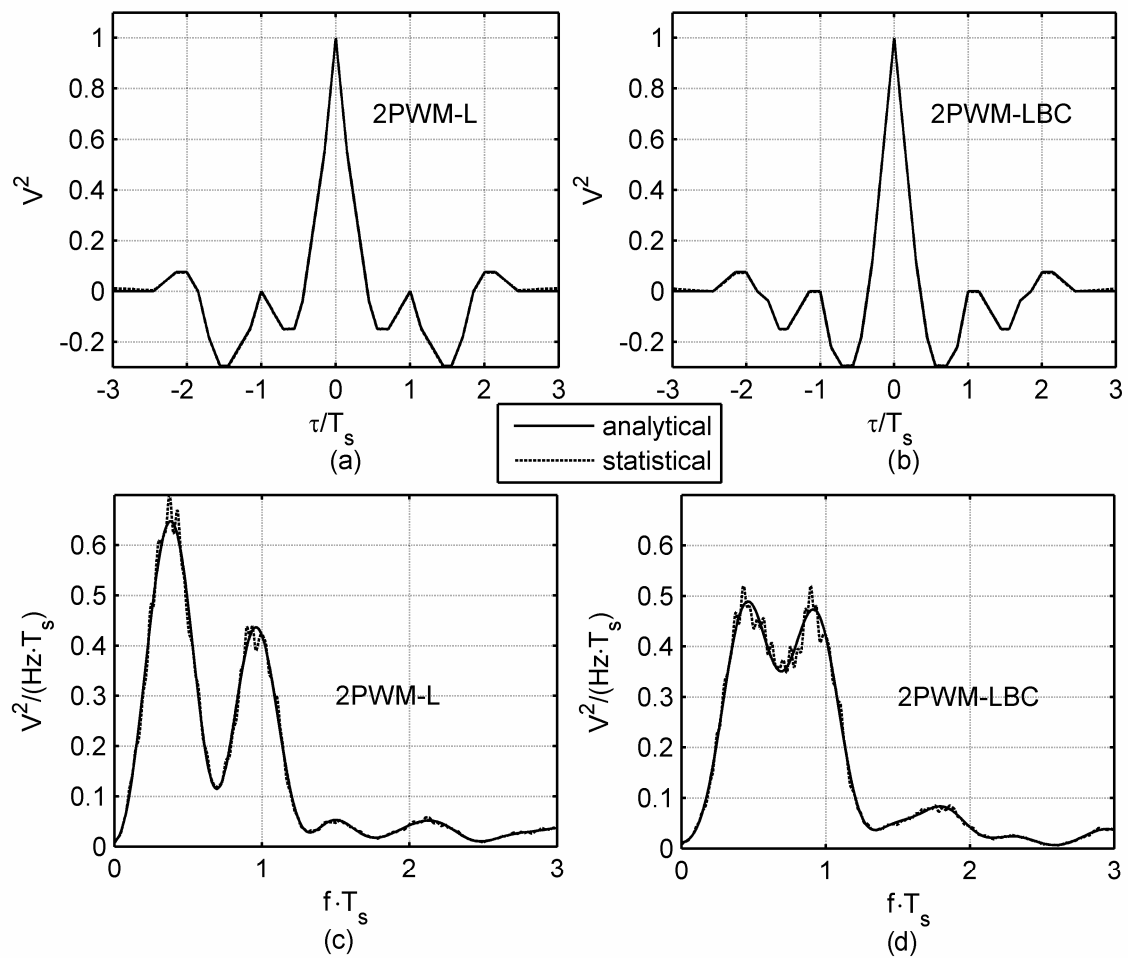

Fig. 18. Autocorrelations and PSDs for 2PWM-L and 2PWM-LBC schemes (from analytical and from statistical simulations). In some plots only the 'analytical' line can be seen because the 'statistical' line is right below it. (a) Autocorrelation for 2PWM-L. (b) Autocorrelation for 2PWM-LBC. (c) PSD for 2PWM-L. (d) PSD for 2PWM -LBC.

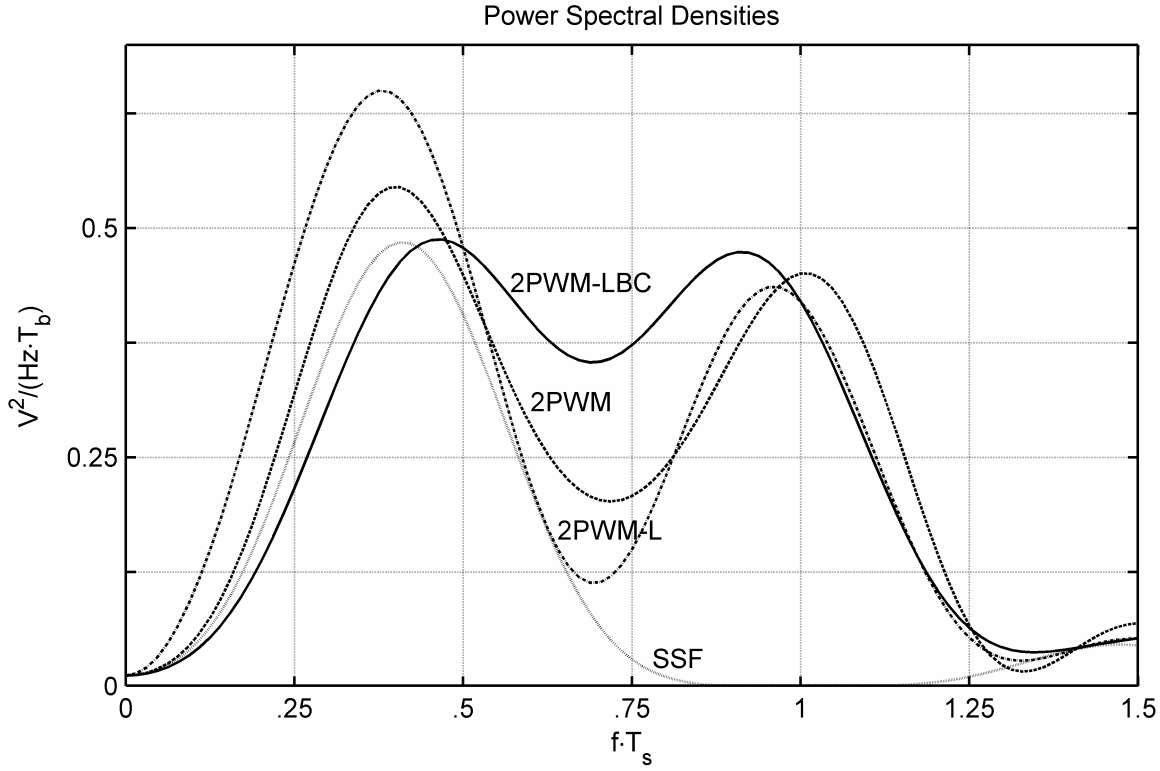

Fig. 19. PSDs for 3-tap SSF, 2PWM, 2PWM-L and 2PWM-LBC schemes.

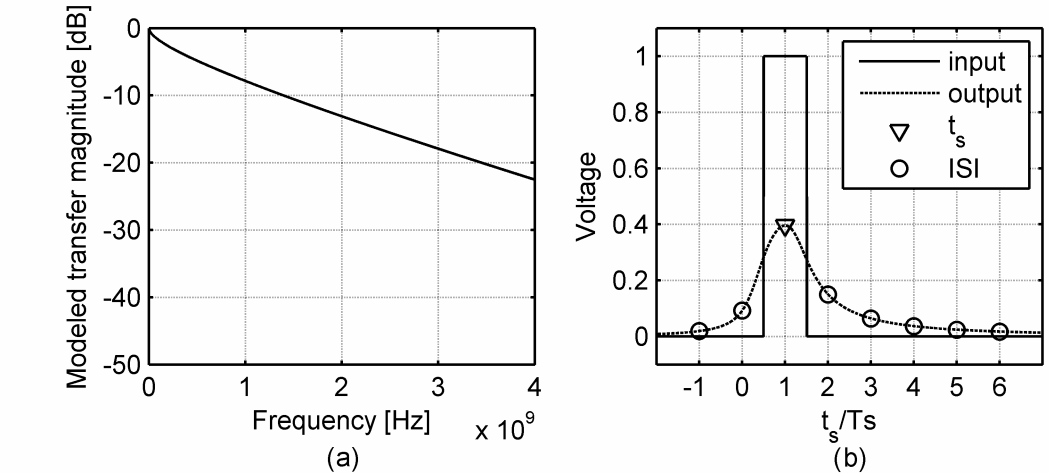

Fig. 20. Alternative (theoretical) channel used for time domain simulations. (a) Channel loss. (b) pulse response (T_s=200ps).

The question may arise of whether the 2PWM-LBC scheme also works for channels with lower loss levels. To evaluate this, time domain simulations were done with an alternative channel. The loss and pulse response for this channel are shown in Fig. 20. The channel has approximately 15dB loss at the Nyquist frequency of 2.5GHz. (The bit rate is again chosen to be 5Gb/s.)

In Fig. 21, the eye diagrams of the response of this alternative channel to the 2PWM-L scheme (Fig. 21(a)) and to the 2PWM-LBC scheme (Fig. 21(b)) are shown. Clearly, the channel response to the 2PWM-LBC scheme shows a larger eye width, indicating that for this alternative channel, the phase relations between the singlets in the 2PWM-LBC scheme are also a better fit than those in the 2PWM-L scheme.

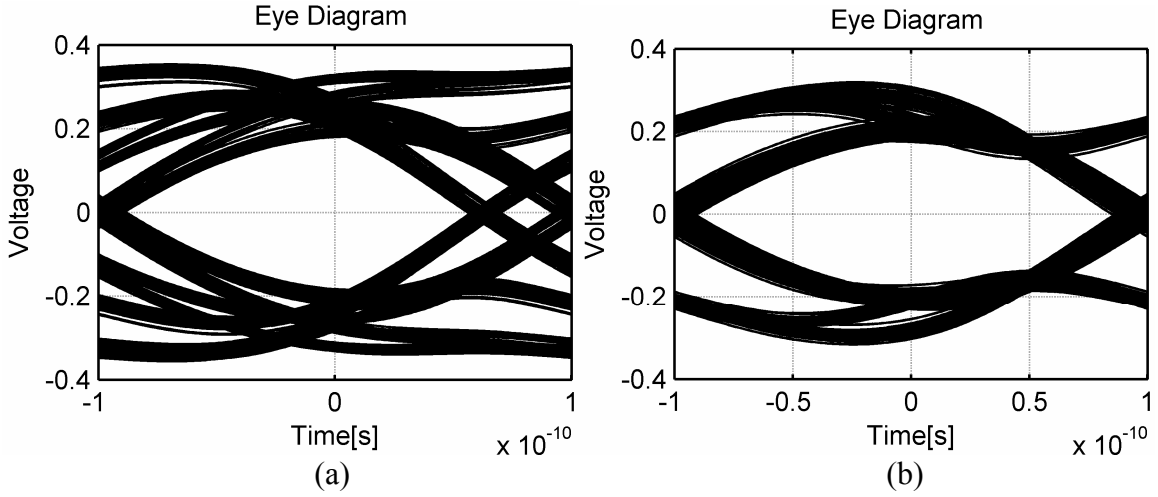

Fig. 21. Eye diagrams of response of alternative channel to 2PWM-L input and to 2PWM-LBC input, T_s=200ps. (a) 2PWM-L. (b) 2PWM-LBC.

In conclusion, we have found an alternative 3-tap, 2-voltage-level PWM scheme that a gives a comparable eye opening as the 2PWM scheme while making less transitions per second. The pulse-widths were found by calculating them from the SSF filter taps with taps 1 and 3 swapped.

5.5. Conclusions

We can extend the PWM pre-emphasis technique to a multitap version. The duty-cycle of this multitap PWM filter is a function of multiple bits instead of a function of only the current bit, as in Chapter 4. In this way, more complex filter transfer functions can be constructed, as with multitap FIR filters. We can use this for example to cancel reflections on backplane PCBs. The transmitter needs to switch between only two voltages, again allowing a transmitter implementation that is based on timing accuracy instead of on amplitude accuracy. The correct switch timing needs to be found as a function of multiple bits, in order to create the desired (effective) filter transfer function. This can be done by deriving the pulse-widths from the SSF filter taps by using a 'same-area' approach.

Three 3-tap PWM filters were discussed:

1) '3PWM': three output voltage levels (-1V, 0V and 1V),
2) '2PWM': two output voltage levels (-1V and 1V),
3) '2PWM-LBC': two output voltage levels (-1V and 1V) and a reduced number of transitions.

The '3PWM' filter transmits, during the time length of one bit, first 0V, then 1V, then 0V again. The '2PWM' filter transmits, during the time length of one bit, first -1V, then 1V and then -1V again. The '2PWM-LBC' filter transmits first 1V, then -1V (less transitions). The voltage levels are inverted when the equivalent FIR filter would otherwise output a negative voltage.

Time domain simulations show eye openings nearly equal to those of the equivalent 3-tap SSF filter.

The power spectral density functions of the 3-tap PWM filters are calculated using their autocorrelation functions. At frequencies around the Nyquist frequency, these filters have a higher power spectral density than the SSF filter due to harmonic signals, inherent to PWM.
A system designer should weigh up this side-effect of PWM against the abovementioned advantage of using a transmission equalizer that needs to switch between only two voltages. As long as the channel has a low-pass transfer function, the high-frequency part of the spectrum will be filtered out by the channel.

The '2PWM-LBC' filter reduces the number of transitions compared to the '2PWM' scheme. This reduces the HF energy transmitted, since the PSD is narrowed. However, shifting the pulses to the left disturbs their phase relations. Using transient simulations, it was found that the receiver eye diagram can be opened up again by using pulse-widths calculated from the SSF filter taps with taps 1 and 3 swapped. The resulting eye opening is nearly the same as for the '2PWM' scheme.