Chapter 1

Introduction

1.1. High speed data-communication over wires

Information and communication technologies are becoming ever more pervasive in our society. The growth rate of the Internet is exponential, with traffic doubling every year [Odlyzko]. There is a wealth of devices available that can be connected to a network or to each other.

These information and communication technologies are largely based on electronics, storing, processing and transmitting information in the form of bits (digital information units that can have a value of ‘1’ or ‘0’) and bytes (=8 bits). After the invention of the transistor, and subsequently the integrated circuit (IC), semiconductor technology also started to follow an exponential growth curve, known as ‘Moore’s law’. As is well known, IC scaling according to Moore’s Law means that every new IC process will have more transistors per chip (and therefore smaller transistors) than the previous generation. Since the electrical capacitance of the individual transistors decreases, the switching speed increases, and so ever more bits per second can be processed in electronic circuits. This is the ‘engine’ driving the growth in the electronics, computer and Internet industries over recent decades.

At some point, bits and bytes need to be transferred from device to device using either cables (optical or electrical) or wireless connections. Although wireless connections can sometimes be more convenient, the maximum achievable bit rate that they offer is lower than for wired connections. This is because using wired connections, we are better able to focus and contain the electromagnetic field energy.

Several applications of wired connections which drive the need for ever higher communication speeds are given below.

In datacenters where information is stored and then distributed over the Internet, interconnections need to be made between servers, switches and routers. With the growth of the Internet, ever more information needs to be processed, and as a result it is essential to increase the bit rates on these interconnections.

Another demand for higher communication speeds is imposed by the home computer user, who (for example) wants to share a video with his family or friends over the Internet. He or she then needs to make a connection between the video camera and the PC to be able to upload the video. The quality of video available to home users is ever increasing (for example, in terms of resolution). Recently, HDTV quality has become available to the home user. The number of bits needed to encode an hour of video is much higher than with older video standards. The home computer user wants to wait seconds rather than minutes for the video to transfer to their PC. This means that the speed of such short interconnects between devices and the PC needs to be increased.

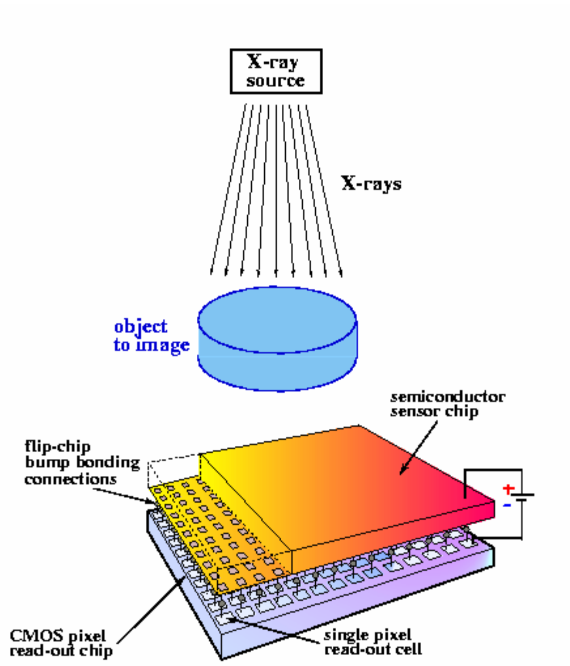

Fig 1. The Medipix2 setup [Llopart].

Inside the PC itself there are also many wired connections. For example, between the hard-disk and the motherboard there is a high-speed interconnect. A current interconnect standard, Serial ATA (SATA), offers the advantages of thinner cables compared to old parallel ATA (PATA) interface. As a result, the ventilation though the PC casing is improved. SATA cables are also cheaper than cables for PATA. The decrease in cable thickness is achieved by serializing the data. This means that the number of channels is reduced (and the speed per channel is increased). Furthermore, the hard-disk read/write speeds are also increasing, and therefore the data rates over the interconnections between the hard-disk and the motherboard are also increasing.

Another interconnect inside the PC is that from the memory to the microprocessor. This connection is made using a printed circuit board (the motherboard). Processor speeds and bus clock rates are increasing so that ever more data can be processed per second. The speed at which this data needs to be fetched from and read into the memory is therefore also increasing.

This PhD project originally started by addressing the need for a high-speed interconnect that would connect the Medipix CMOS X-ray detector to a data storage system. This Medipix detector is a CMOS X-ray imaging chip [Llopart]. It is bump-bonded to a detector chip, such as high resistivity silicon, which is used to convert X-rays into hole-electron pairs. This is illustrated in Fig. 1. In future digital X-ray imaging systems, very high data rates will be needed. For some applications, e.g. protein crystallography at synchrotron beams, data rates in the order of gigabits per second are anticipated. The present trend for such systems is to move from a parallel data bus towards a high-speed serial link. Therefore, more bits per second need to be transmitted from the chip to a storage system.

All these wired interconnections can be formed by copper cables or glass fibers. Glass fiber offers a higher bandwidth-distance product, and therefore outperforms copper over long distances at high data rates. However, a high speed laser driver is more expensive than a straightforward connection of copper to an electronic circuit. Also the coupling of the light into a high speed fiber demands high precision. With the dominance of microchips as the primary means of signal processing, all the signals are produced in the electrical domain. Using copper interconnects offers the advantage that that no electro-optical conversion is needed before transmission. Copper can provide a transmission speed of the same order and at high reliability for interconnect lengths up to approximately 100m, provided that high quality cables are used.

Our area of study is limited to electronic circuits for high speed serial links over copper interconnects. In this thesis, we analyze, implement and test techniques to increase the data rate of communication over copper channels.

This Introduction is organized as follows. In section 1.2, an overview is provided of the different types of copper interconnects and of their behavior in the frequency domain. Next, in section 1.3, it is shown how the problem of intersymbol interference limits the achievable bit rate. In section 1.4 it is explained how an equalizer can be used to increase the bit rate. After that, in section 1.5, the advantages and disadvantages of specific equalizer implementations and modulation schemes from the literature are discussed. At the end of this chapter (section 1.6) the research challenges are listed and an overview of the thesis is provided.

1.2. Types of copper interconnects and their behavior in the frequency domain

In this section we discuss some of the most frequently used types of copper interconnections available (subssection 1.2.1), and look at their electrical behavior and capacity limit from a frequency domain point of view (subsection 1.2.2).

1.2.1. Types of copper interconnects

There are several types of copper cable and copper interconnect available for high speed links, as shown in Fig. 2. One type is the coaxial copper cable, which consists of a center conductor and an outer conductor, separated by an insulating dielectric, as shown in Fig. 2(a). This dielectric is typically polyethylene or Teflon, but can also be of another type, for example a low-loss dielectric. A plastic jacket is fitted to the outside. Another type, the twisted pair cable, consists of two twisted conductors, each surrounded by an insulating dielectric. Depending on the presence or absence of a metal shield, this cable is either called shielded twisted pair (STP) or unshielded twisted pair (UTP). The latter is commonly used for local area networks (LANs) in buildings. Again, a jacket is fitted to the outside, as shown in Fig. 2(b).

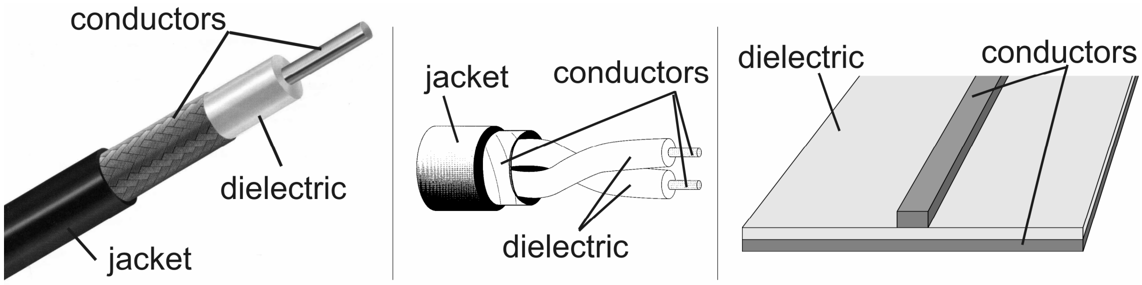

(a) (b) (c)

Fig. 2. Copper channels. (a) Coaxial copper cable. (b) Twisted pair copper cable. (c) Microstrip on printed circuit board.

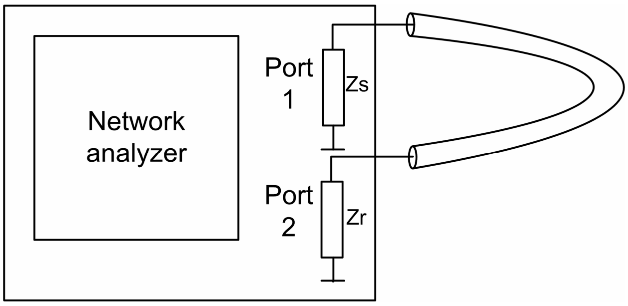

Fig. 3. Setup for loss measurement.

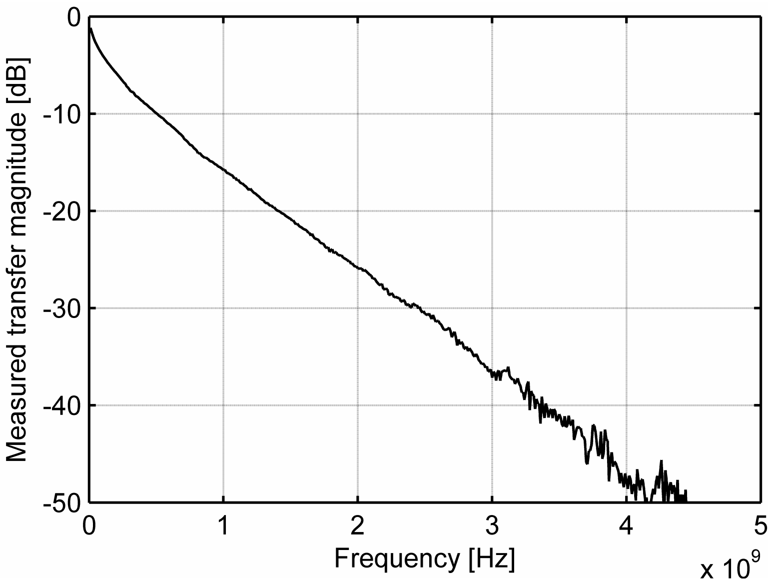

Fig. 4. Measured magnitude of transfer function of 25m RG-58CU coaxial cable.

A third type of copper interconnect, the printed circuit board (PCB), is commonly used to connect components. Examples are the motherboard inside a PC and the PCB inside a mobile phone. These boards are a sandwich of copper layers and dielectric layers. Fig. 2(c) shows a PCB with a single dielectric layer. The copper tracks can be etched using lithography. Several types of dielectric are available, the most common being the green-colored glass-epoxy FR4. More expensive, low-loss dielectrics are also available. See Fig. 2(c) for an illustration of a controlled-impedance microstrip over a groundplane.

1.2.2. Frequency domain behavior

We now look into the electrical behavior of these three copper channel types. How do they behave and what limits does this behavior impose on the achievable bit rate? Frequency domain measurements are shown, made using a network analyzer, to show the channel loss.

Clearly, practical conductors and dielectric materials fall short of the ideal. They exhibit undesirable attenuation ('loss'), because inside the conductor and dielectric, energy is converted from the electrical domain to other domains, for example to heat. They are also dispersive, which means that the signal delay between input and output varies with frequency.

To measure cable loss in the frequency domain, we use a network analyzer. Using the measurement setup shown in Fig. 3, we measure the channel loss (also known as magnitude of the transfer function, or the attenuation) as a function of the frequency for a 25m RG-58CU coaxial cable. Both ports of the network analyzer are terminated with the characteristic impedance of the cable (50Ω, Z_s at the source and Z_r at the receiving end). The magnitude of the transfer function S_{21} is shown in Fig. 4. We plot the loss against a linear frequency axis.

1.3. Factors limiting the bit rate – a time domain view

Clearly the cable shows a loss that increases with frequency. To understand how this limits the achievable bit rate on copper channels, we change our point of view and look at the problem in the time domain. This section describes intersymbol interference caused by channel loss and dispersion (subsection 1.3.1) and reflections (subsection 1.3.2).

1.3.1. Intersymbol interference

The attenuation and dispersion effects can be seen in the time domain as intersymbol interference (ISI). In this subsection we illustrate ISI with the commonly known and widely used 2PAM modulation method for high-speed baseband transmission over copper channels [Dally], [Lee], [Kudoh].

The 2PAM modulation method simply uses one voltage to transmit a binary '1' and another voltage to transmit a binary '0'. When these levels are opposite this is called polar Non-Return to Zero (polar NRZ). This modulation method represents the most straightforward, 'natural', off-chip output of the on-chip digital signals, after serialization. (Later, in Fig. 8, a 2PAM signal is shown).

We look at the ISI problem using the pulse response (1.3.1.1), the transient response (1.3.1.2) and the eye diagram (1.3.1.3). Important definitions of the Nyquist frequency and the bit error rate (BER), which are used throughout the thesis, are given along the way.

1.3.1.1. Pulse response

To understand ISI, we can look at the cable response to a single PAM pulse. The width of this pulse is equal to the bit length T_s (which in turn is equal to the inverse of the bit rate R). In Fig. 6, our measurement setup is shown, consisting of a pattern/pulse generator, the channel, and an oscilloscope. The cable is again terminated with characteristic (50Ω) impedances Z_s and Z_r at the transmit and receive sides respectively.

The measured response of 25m RG-58CU coaxial cable to a single 2PAM pulse of length T_s=200ps (bit rate 5Gb/s) is shown in Fig. 7. The cable output is time-shifted to compensate for the propagation delay. In the figure, the sampling point of the received pulse (the cursor) is indicated with a triangle and the sampling points of the next and previous symbols are indicated with circles. The circles with a nonzero value add to the ISI. A long tail on the right side can be seen, causing a high level of "post-cursor" ISI. There is also one "pre-cursor" ISI point. There is a clear asymmetry in the pulse response, which is indicative of dispersion.

Fig. 6. Time domain measurement setup.

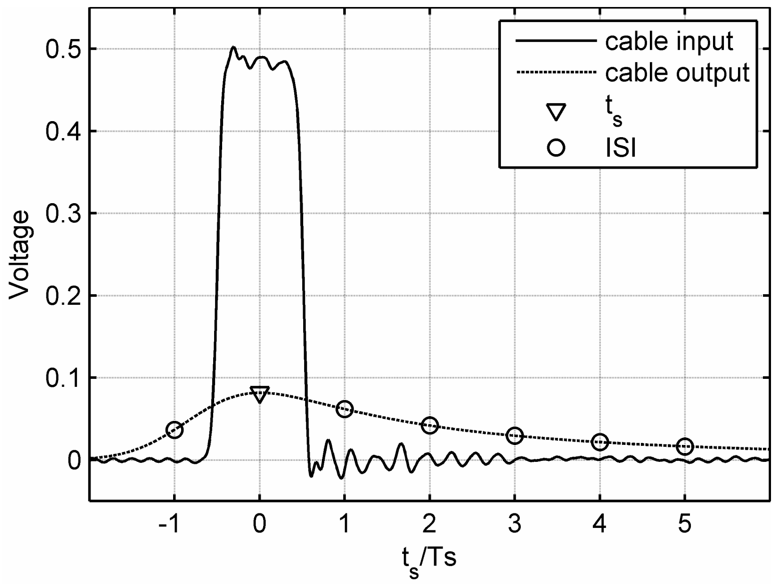

Fig. 7. Measured pulse response of 25m RG58CU (time shifted), and measured input pulse.

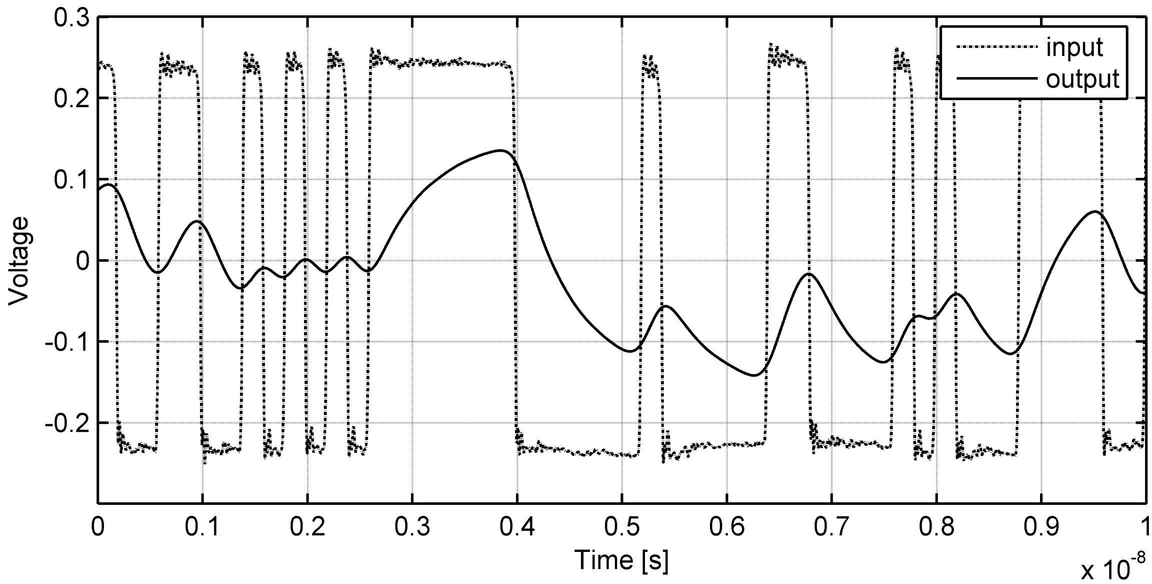

Fig. 8. Measured transient response of 25 m RG-58CU coaxial cable. 5 Gb/s data input and output.

1.3.1.2. Transient response

When a string of 2PAM modulated bits is transmitted over the cable using the pattern generator, we obtain the result shown in Fig. 8. Clearly, the channel response is severely distorted by intersymbol interference (ISI). All the tails of the separate pulse responses are summed up. It is no longer possible to simply place a threshold at 0V and decide whether the measured signal is above or below it. The channel has a low-pass transfer function, and therefore the high frequency (HF) components are attenuated more than the low frequency (LF) components. As a result, a 'baseline wandering' effect can be seen in Fig. 8.

The Nyquist frequency

The Nyquist frequency f_N of a 2PAM signal is defined as half the bit rate. In our example, with a bit rate of 5Gb/s, the Nyquist frequency is 2.5GHz.

From the transient response shown in Fig. 8, we can roughly estimate the channel attenuation at f_N by looking at the fastest transitions, e.g. from 1.5ns to 2.5ns. Their amplitude is approximately 20mV. Seeing that the amplitude of the cable input signal is approximately 0.5V, the channel attenuation at f_N is 20m/0.5=0.04, which corresponds to 28dB loss. (Compare this with the frequency-domain measurement of 30dB in Fig. 4.)

The significance of the Nyquist frequency for the design of equalizers is that we can use it to define a figure-of-merit for equalizers, as is shown later.

1.3.1.3. Eye diagram

The quality of the received signal can be accurately assessed using an eye diagram. An eye diagram is made by setting the oscilloscope x-axis to one single bit time, so that all bits are plotted on top of each other. Effectively, the oscilloscope is now triggered by the bit clock. This results in an 'eye' shape, as illustrated in Fig. 9. Fig. 9(a) shows the input signal, with a clearly open eye, and Fig 9(b) shows the cable output, with a closed eye. The eye opening tells us a lot about the signal integrity. The eye opening and eye width can be used to estimate the voltage margin and timing margin respectively. This can be understood as follows. The receiver needs to sample each bit, so the sampling interval is T_s. The sampling position within the interval $(-T_s/2, T_s/2)$ is denoted as t_{sam}. There is a certain allowance for deviation of the sampling position from t_{sam}, which is given by the timing margin. Furthermore, the receiver needs to decide whether the transmitted bit was either a one or a zero, by measuring its value relative to a threshold v_{th}. The eye height tells us how large the voltage margin is. We need to have a certain margin because when noise (e.g. thermal noise) is added to the signal, there is a chance that the signal crosses the threshold and the bit will be detected wrongly, resulting in a bit error.

Jitter

Jitter (timing noise) in the receiver limits the symbol rate at which the receiver can reliably detect data. We can understand this using the eye diagram. When the sampling time t_{sam} varies due to jitter, the timing margin in the eye diagram determines the allowable variation for t_{sam} to prevent bit errors from occurring.

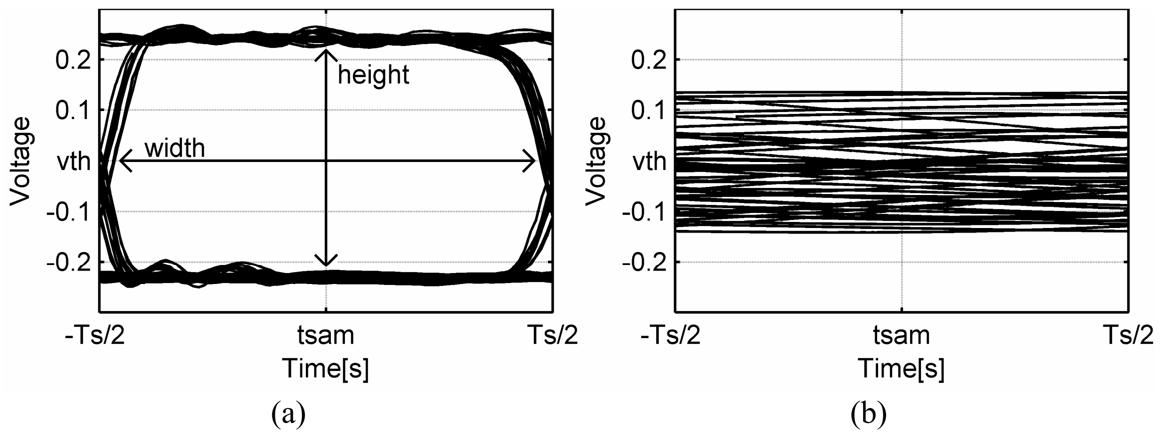

Fig. 9. Eye diagrams of measurements on 25m RG-58CU coaxial cable (R=5Gb/s, T_s=200ps). (a) Cable input. (b) Cable output: eye completely closed by ISI.

The origin of this jitter depends on the system implementation. Often, the sampling time t_{sam} is determined by a phase locked loop (PLL) via a clock and data recovery (CDR) loop. This PLL shows jitter, for instance because the thermal noise in the transistor channel translates into timing noise in the Voltage Controlled Oscillator (VCO). This happens for example when a noisy signal triggers a next stage in a ring oscillator or if the threshold of a comparator or logic gate is noisy due to thermal noise. The slew rate of the signal then determines the translation coefficient of the noise amplitude to timing noise [Abidi].

Bit error rate

When the receiver makes an incorrect decision about the received bit, a bit error occurs. This can be quantified using a statistical metric called the bit error rate (BER). This metric is defined as the ratio of the number of bits incorrectly received to the total number of received bits. For high-speed links, the target BER is typically in the order of 10^{-12} [Farjad-Rad].

1.3.2. Reflections

Another potential limit on the achievable bit rate comes from signal reflections.

At the connection of the transmission line to the circuit there are usually bondwires, and possibly a short section of another transmission line. At the boundaries between these lines there are often impedance discontinuities. Furthermore, connectors in the signal path might pose additional impedance discontinuities. When a signal traveling through a transmission line arrives at an impedance discontinuity, part of the signal will be reflected. This is undesirable, because reflections can distort the received bit pattern.

In the past, computer buses operating at low speeds would simply be designed to reflect 100% of the signal energy so as to double the received voltage. In Gb/s serial links, this is not the usual approach. At Gb/s speeds, the length of one bit traveling through a copper cable is only a few centimeters. For example, the length of one single bit in a 10Gb/s 2PAM signal traveling over a copper cable at a propagation speed of 2/3 times the light-speed is only 2cm. This means that a 10m long cable will typically have ~500 bits 'in transit'.

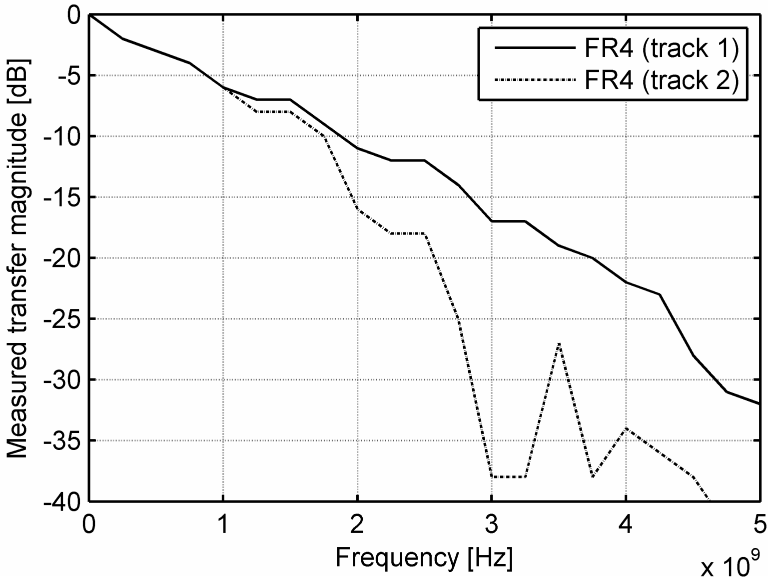

Fig. 10. Measured magnitude of transfer function of two 20” FR4 backplane tracks.

By using double termination, both at the transmitter and receiver end, the maximum signal integrity is achieved. This is called incident-wave signaling [Dally-2]. For good quality cables, reflections can be decreased to below 10%, with well-controlled characteristic impedance of the cable and proper matching of the termination impedances – plus well designed connectors. Even so, a 10% reflection added to the received signal will limit the effective SNR to 20dB.

For printed circuit boards, the problem of impedance discontinuities is harder to avoid than for cables. There are usually vias (connections through the dielectric) along the track which cause impedance discontinuities. Also connectors pose additional impedance discontinuities. This can lead to severe notches (spectral nulls) in the channel transfer function, as illustrated in Fig. 10.

1.4. Equalization

The previously discussed problems of ISI and reflections can be partially solved by using equalization. In this section, we explain the concept of equalization (subsection 1.4.1). A figure-of-merit (FoM) for equalizers is given next: the loss compensation (subsection 1.4.2).

1.4.1. Concept of equalization

Provided that the SNR is high enough, it is possible to reconstruct a PAM signal that is low-pass filtered by the channel. To achieve this, an equalizing filter is placed in series with the channel. This filter can be placed either at the transmitter side, where it is called a TX pre-/de-emphasis filter, or at the receiver side, where the technique is called RX equalization. (Illustration in Fig. 11.)

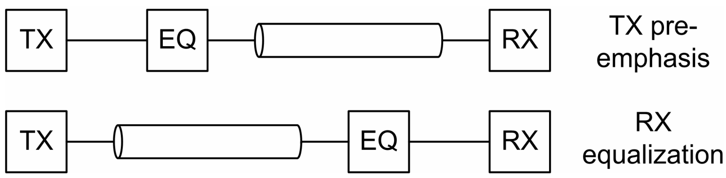

Fig. 11. Transmit (TX) pre-emphasis versus receiver (RX) equalization.

This filter's transfer function is the inverse of that of the channel, so that the effects of attenuation and dispersion are compensated. Below, we illustrate the concept of an equalizer from a frequency domain point of view. First, we assume that the channel to be equalized has a first order low-pass characteristic, given by

$$H_{1st}(j\omega) = \frac{1}{1 + \frac{j\omega}{\omega_{ch}}}. \tag{1}$$

By placing an equalizing filter in series with the channel, which has the transfer function

$$H_{eq}(j\omega) = \frac{1 + \frac{j\omega}{\omega_{ch}}}{1 + \frac{j\omega}{\omega_{p,eq}}}, \tag{2}$$

we can equalize the channel and the equalized channel has the following transfer function:

$$H_{tot} = H_{1st} \cdot H_{eq}(j\omega) = \frac{1}{1 + \frac{j\omega}{\omega_{p,eq}}}. \tag{3}$$

Thus we can extend the channel bandwidth, provided that $\omega_{p,eq}$ is larger than ω_{ch}. Before equalization, the channel was band-limited to ω_{ch} rad/s. After equalization, the -3dB bandwidth is increased to $\omega_{p,eq}$ rad/s. In the time domain, this has the effect of resolving the ISI. This is discussed in more detail in Chapter 4.

The effect of an equalizer is illustrated in Fig. 12. In Fig. 12(a), the equalizer gain plot is given. In Fig. 12(b), the unequalized and equalized 1^{st} order low pass channel are plotted. In the example, we have chosen the values ω_{ch}=150MHz and $\omega_{p,eq}$=2.5GHz. Clearly, the equalized channel is flat over a much wider frequency range, so that the -3dB bandwidth is increased.

For comparison, the measured loss of the 25m RG-58CU cable is also plotted. (Note that the frequency axis is now logarithmic instead of linear as before.) The -3dB bandwidth of the cable is only 60 MHz. The cable has many poles in the high frequency (HF) range.

When the transfer function of the channel that we want to equalize exhibits deep notches in the signal band (as in Fig. 10), equalization is less useful. An equalizer for such a channel would need to have a very high gain near the notches. It is then unavoidable that the equalizer also amplifies the noise within that frequency range to a high level. A possible solution is to use modulation schemes that avoid putting information at those locations, as we discuss later in this chapter.

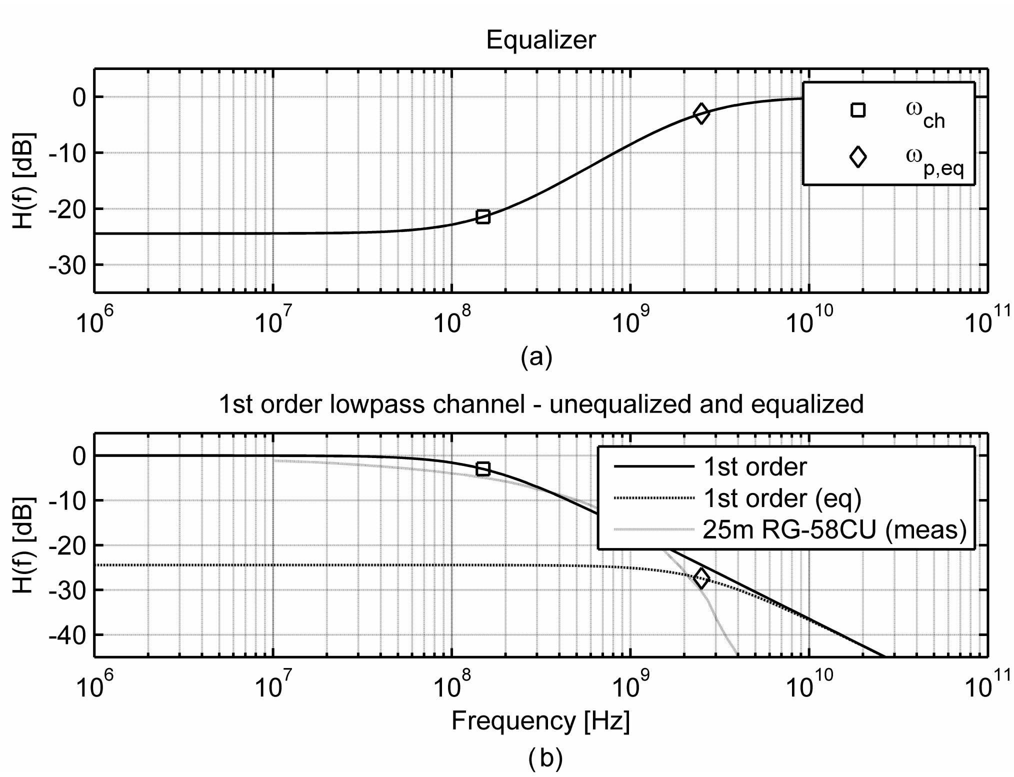

Fig. 12. Effect of equalizer. (a) Equalizer transfer function. (b) Transfer function of unequalized channel and equalized channel, which is flat over a wider frequency range. Measured loss of a 25m RG-58CU cable is also shown.

1.4.2. Equalizer gain plot and loss compensation

From the equalizer gain plot we can read the level of gain (in dB) that the equalizer provides at the highest frequency, relative to its DC gain. In the example in Fig. 12(a), the equalizer gain is approximately 24dB. This gain is often adjustable so that the equalizer can be adapted to the channel.

A figure-of-merit for equalizers that we use throughout this thesis is the *loss compensation*. It is closely related to the equalizer gain. The loss compensation of the equalizer is equal to the measured channel loss at the Nyquist frequency, under the following conditions. We take a channel of a certain type and length such that the equalizer needs to be set to its maximum gain, while a still acceptable BER is measured at the receive side. This BER is, of course, dependent on the eye opening. A certain eye closure is acceptable as long as the eye closure does not degrade the BER below the acceptable level. (As mentioned above, in general the acceptable BER for high speed links is set very low at $<10^{-12}$).

The loss compensation figure is not dependent on the bit rate that can be achieved by using an equalizer, because of its definition as the channel loss at the Nyquist frequency. We could, for example, have two different equalizers: one for high bit rates and one for low bit rates, with completely different pole and zero locations yet with the same loss compensation. It would only mean that the equalizer for the slower bit rate could transmit information over a longer cable (or over a cable type with more loss per meter). For this reason, it is also important to specify the bit rate that an equalizer can handle.

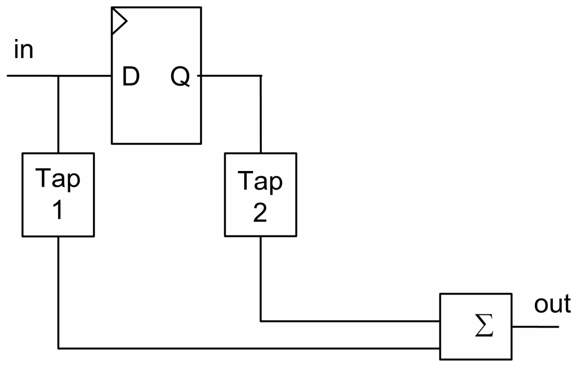

(a)

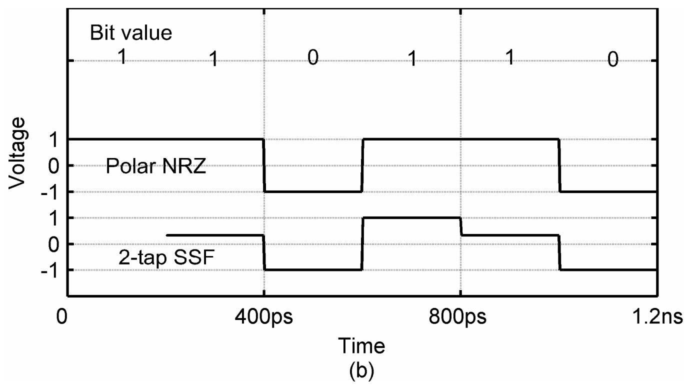

Fig. 13. A 2-tap symbol spaced FIR pre-emphasis filter and its output. (a) 2-tap SSF filter architecture. (b) Time domain signals for 2PAM (polar NRZ), and for the output of the 2-tap SSF filter.

1.5. Solutions from the literature

In this section, we discuss solutions from the literature to the attenuation, dispersion and reflection problems introduced above. Equalization schemes (subsection 1.5.1) and modulation schemes (subsection 1.5.2) are discussed that allow us to approach the theoretical capacity limits of the copper channels.

1.5.1. Equalizers

In high-speed serial PAM systems, equalizers are often used to increase the bit rate. In this subsection we discuss the main advantages and disadvantages of proposed equalizer solutions from the literature, to place the work of this thesis in context. There are currently several methods available for achieving a desired equalizer function. First, in 1.5.1.1, the commonly used transmit-side FIR pre-emphasis (PE) is discussed. Next, in 1.5.1.2, receive-side equalizers are discussed.

1.5.1.1. FIR PE

Finite impulse response (FIR) filter-based transmit-side equalization can be implemented in a straightforward way by a FIR-DAC (digital-to-analog converter). Such a FIR-DAC has a digital input and an analog output. In Fig. 13(a), a typical FIR-DAC implementation of a 2-tap symbol-spaced FIR (SSF) pre-emphasis filter is shown. A flip-flop is used to delay the signal with one bit time, and the input and output of this flip-flop are weighed and added to form the output. The output is shown in Fig. 13(b), together with the input 2PAM signal. It can be seen that the filter gives a maximum output when there is a transition in the bit stream, either from ‘0’ to ‘1’ or vice versa, and when there are no transitions the output is attenuated. This has the effect of attenuating the lower frequencies. Because the transitions are emphasized, this filter is termed a pre-emphasis filter. Actually, *de*-emphasis is a better word because the equalizer attenuates the lower frequencies to make the transfer function flatter.

Methods described in the literature for equalization of coaxial cables, twisted pair cables and PCBs are commonly based on this FIR-DAC implementation [Lee], [Farjad-Rad], [Dally], [Gai], [Balan], [Stonick], [Zerbe]. Usually the filters are 2-tap symbol-spaced FIR (SSF) filters. However these 2-tap filters typically offer only 10dB of loss compensation, e.g. [Lee] at 4Gb/s (2PAM), up to a maximum of approximately 18dB, e.g. [Kudoh] at 5Gb/s (2PAM).

In [Gai] a 5-tap symbol-spaced equalizer is described that can handle 30dB of channel loss at the Nyquist frequency, at 3.125Gb/s (2PAM). Thus higher loss compensation is possible with 5 manually tuned FIR taps. However, the 5-tap implementation imposes strict demands on transistor matching. Deterministic jitter can arise as a result of mismatches. Also, the large number of taps introduces many degrees of freedom that could make it harder for an adaptation algorithm such as LMS to adjust the equalizer to the channel. (For example in [Stonick], 60% of the 1W power budget is spent on the digital blocks that implement the adaptation algorithm.) Furthermore, the number of voltage levels at the output increases by a factor of two with every tap added and linearity demands on the output stage become progressively stricter.

We conclude that 2-tap SSF pre-emphasis, while straightforward to implement, can not provide more than ~20dB of loss compensation. Furthermore, we conclude that 30dB loss compensation is only achievable using more complex circuits, but which have more parameters that need to be tuned to match the equalizer to the channel.

1.5.1.2. RX side Equalizers

Receiver side equalizers have the advantage that the information to adapt the filter coefficients is readily available. The feedback loop is short and a quick adaptation can be made to changing channel conditions if necessary. We categorize the receiver-side equalizers into two groups: continuous-time equalizers and discrete-time equalizers.

Continuous-time equalizers

A continuous-time equalizer is – unlike discrete time equalizers – not based on sampling but implements the desired transfer function using an analog time-continuous circuit. We discuss two types of continuous-time equalizers: second-order derivative equalizers and equalizers using RC-degenerated stages.

The second order derivative equalizer works by adding three processed versions of the input [Gai]:

1) the input itself;
2) the derivative of the input;
3) the second order derivative of the input.

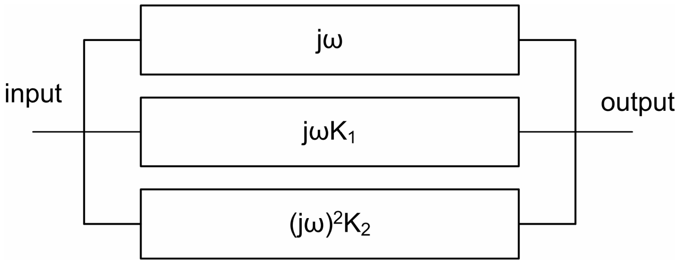

Fig. 14. Second order derivative equalizer.

All these signals are scaled with constants, such that the total output signal can be described as:

$$H(j\omega) = K_0 + j\omega K_1 + (j\omega)^2 K_2 .$$

By choosing the optimum constants, a transfer function can be made that closely approximates the inverse of that of the copper interconnect. In each path there are variable gain amplifiers (VGAs) to provide the scaling *K*. Fig. 14 is a conceptual illustration of such a second order derivative equalizer.

Using a second order derivative equalizer, up to 30dB of loss compensation is achieved at 3.125Gb/s in [Gai]. However, matching the latencies through the three paths is not straightforward and the circuitry is complex. Furthermore, care has to be taken that the input stage does not become saturated, because in that case the received input signal would show clipping and information would be lost.

An implementation which avoids the use of separate paths is the 'RC-degenerated' implementation [Choi]. Such an architecture uses an RC impedance for source degeneration. Fig. 15 shows a degenerated differential pair. Z_2 forms the source degeneration. The impedance Z_1 is usually just a resistor, but Z_2 incorporates a larger impedance network, which contains one or more poles and sometimes also a zero.

By cascading multiple stages of such RC-degenerated differential pairs, in [Maillard] 30dB of equalizer gain was achieved in 0.18µm CMOS, at a speed of 3Gb/s. To obtain this result five stages were needed. Each stage has a Z_2 impedance that can be switched, to enable adaptation of the equalizer to the channel transfer function. However, this involves a large number of carefully matched poles and zeros (of which a certain number needs to be switched on depending on the channel parameters). Also, the switching order needs to be taken into account. Saturation of the input needs to be prevented, as in the second order derivative equalizer. (This is a general point for receiver equalizers.)

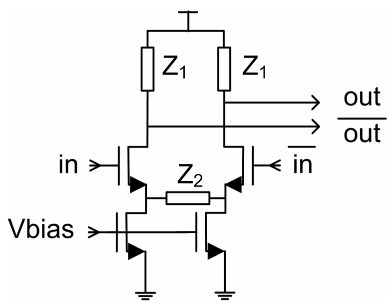

Fig. 15. Single RC-degenerated stage.

In conclusion, a loss compensation of 30dB has been demonstrated at the receiver side, but only with a second order derivative equalizer and with a 5-stage RC degenerated equalizer.

Discrete-time equalizer: decision feedback equalization

A discrete-time receive-side equalizer architecture is the decision feedback equalizer (DFE). The concept is illustrated in Fig. 16. In a DFE, decisions of previously detected symbols are subtracted from the input to remove ISI in the current symbol. This has the advantage that the DFE rejects high-frequency noise, because of the 'threshold' decision process [Proakis]. This makes the DFE a good candidate for the noisy environment of 'backplane' PCBs (e.g. those suffering from crosstalk) [Stojanović], [Balan], [Zerbe].

However, the DFE is sensitive to error propagation: when the value of a bit is incorrectly decided upon, it will influence future decisions, because of the feedback [Balan]. Furthermore, at very high bit rates it is challenging to feed back the decisions quickly enough to allow implementation of the first filter tap. Loop unfolding techniques can be used to mitigate this [Zerbe], but this unfortunately also introduces unwanted loading in the critical signal and clock paths, and therefore makes the clock and data recovery (CDR) circuit design more complicated [Stojanović].

A further limitation of the DFE architecture is that only post-cursor ISI can be cancelled, because the bit decisions have to be available to be fed back. For these reasons the achievable loss compensation with DFEs is generally low, and so they are often used in combination with FIR transmit equalizers.

1.5.2. Modulation: OFDM

As was shown in Fig. 10, an FR4 'backplane' PCB can exhibit deep spectral nulls in its transfer function, caused by impedance discontinuities. When using PAM for these copper channels, any equalizer would need to have a very high gain at the frequencies with destructive interference, also amplifying the noise to a high level in the process. The information at those frequencies is needed to restore the original signal.

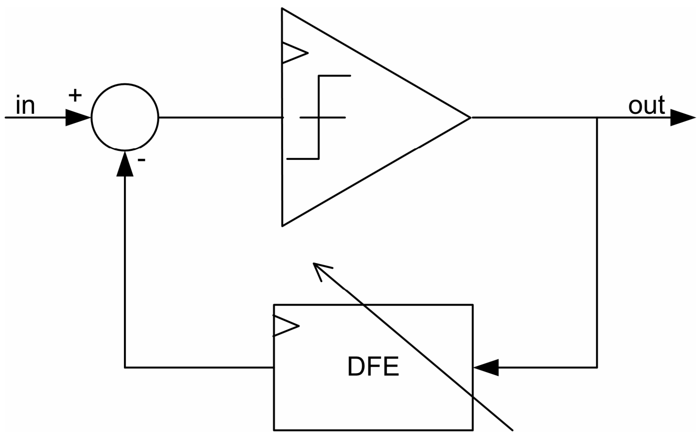

Fig. 16. Decision feedback equalization.

While we have seen above that DFEs can offer a solution to this ‘noise amplification’ problem, this is limited because a high enough noise level will eventually cause incorrect decisions in the DFE. In severe cases, it might actually be better to use a modulation method that avoids putting information at those frequencies where notches occur in the channel transfer function. Orthogonal Frequency Division Multiplexing (OFDM) [Weinstein], [Bingham] is, in principle, well suited to this task.

The required data rates in the backplane environment are much higher than those in commonly used and proven digital multi-tone (DMT) implementations of OFDM. For example, data rates achieved over digital subscriber lines (DSL) are around 10Mb/s. However, to achieve a data rate of 10Gb/s using a DMT implementation would require an extremely high bandwidth, high resolution ADC, and digital signal processing that is currently just not feasible. Therefore the experimental high-speed analog OFDM implementations need to use analog mixers and filters to combine and separate the subcarriers. Solutions of this form are being investigated, e.g. a so-called analog multi-tone (AMT) system [Amirkhany-1]. However, when the local oscillator signals for the mixers experience jitter and duty-cycle deviations, the orthogonality constraints of OFDM are violated. We analyze the impact of such timing non-idealities in this thesis.

1.6. Research challenges and thesis overview

In this section the research challenges are summarized (subsection 1.6.1) and an overview of the thesis is provided (subsection 1.6.2).

1.6.1. Research challenges

The main goals of this work are:

- to achieve an increase in the bit rate over lossy, dispersive copper channels, over the maximum currently possible,
- to maintain a low Bit Error Rate (BER),
- to implement the above using simple circuits that are feasible to implement in future high-speed CMOS processes.

We have concluded in this introduction that the main problems in communication over copper channels are attenuation, dispersion and reflections, which result in intersymbol interference in the time domain.

Concerning the reflections, we have seen that, for channels with many impedance discontinuities (which lead to deep spectral nulls in the transfer function) an analog OFDM system might offer a solution. However, the high bandwidth necessitates an analog implementation, and we need to investigate what the effect of this implementation is on the OFDM channel orthogonality. Special attention needs to be given to timing non-idealities.

Concerning the attenuation and dispersion arising from conductor and dielectric losses, we have seen that equalization offers a solution. However, equalizer implementations described in the literature that achieve a very high equalizer gain / loss compensation (up to 30dB) can typically only be implemented using complex circuits with multiple degrees of freedom to adjust the equalizer transfer function to the channel. We would like to develop a simpler equalizer that decreases this number of parameters while still providing a high loss compensation.

Finally, for effective numerical computer simulations of the new techniques, we need an accurate time domain model of the channels.

1.6.2. Thesis overview

In Chapter 2, we review existing channel models, with a focus on transient (time-domain) simulations using copper cables and PCBs. We look into recent publications to find a model for the dielectric loss which improves on those described in current textbooks (which typically give the frequency dependent loss tangent only at certain points). Special attention is given to causality, so as to provide an accurate impulse response for use in high speed transient simulations. The selected model is fitted with and compared to measurements.

In Chapter 3, we analyze whether OFDM is a useful modulation method for data transmission over channels that suffer from severe reflections. An analog high-bandwidth OFDM system is described and its performance is analyzed. We focus on the sensitivity to timing non-idealities such as jitter and duty-cycle deviations. It is expected that these timing non-idealities may cause serious inter-carrier interference (ICI) in OFDM systems, which limits the achievable bit rate.

In Chapter 4, we look at alternative equalization techniques to increase the bit rate over copper channels. It is shown that a simple pulse-width modulation (PWM) pre-emphasis technique works well on copper channels, achieving up to 30dB loss compensation. We analyze PWM pre-emphasis in the time and frequency domains and compare the loss compensation to that achieved using a conventional FIR pre-emphasis filter. The design of three proof-of-concept transmitter chips is described. Measurement results are given for both coaxial and differential cables, and for a printed circuit board.

In Chapter 5, the PWM pre-emphasis technique is extended to a multitap version. As with a multitap FIR filter this enables us to make more complex transfer functions. The power spectral density functions of the multitap PWM filters are calculated using their autocorrelation functions, and time domain simulation results are given.

In Chapter 6, the thesis is summarized, the original contributions are listed, and suggestions are given for future research topics.