UNIVERSITÀ
DEGLI STUDI
DI UDINE

Dipartimento Politecnico di Ingegneria e Architettura

Corso di Dottorato in Ingegneria dell'Informazione

PhD Thesis

Modellizzazione, Progetto e Caratterizzazione di un Link Seriale a 10Gbps

Modeling, Design and Characterization of a 10Gbps Serial Link

Supervisor: Prof. Ing. Pierpaolo Palestri PhD Candidate: Andrea Bandiziol

---

Anno Accademico 2016 – 2017

Abstract

The goal of this PhD has been to model, design and characterize a 10Gbps serial interface suitable for automotive Electronic Control Units (ECU). The work has been carried out in collaboration with Infineon Technology.

High speed serial interfaces are a hot topic both in the academic and industrial world. Due to the stringent safety requirements and the extremely harsh environment in which the link must be able to correctly operate, the automotive sector lags some years behind the consumer market. Thus, the main goal of this work is to bridge the gap between the consumer electronic and the automotive electronic unit world, understanding which techniques are suitable for our working conditions among the ones that are already well established in the academic world and translating and improving these solutions to possibly make them more stable and less power consuming. This goal implies a deep understanding of a serial link both at system and transistor level, and the development of this thesis will follow this idea.

The first part of this work is dedicated to the transmitter: we will start from a system level analysis, creating a methodology to assess the equalization capability that has to be foreseen at transmitter side when dealing when channels typical of the automotive environment. The description of the transistor level design will follow, motivating design choices and supporting them with simulation results and comparison with the state of the art presented in literature. To conclude this first part of the work, measurements of the described transmitter will be presented and discussed.

The second part of the thesis is mainly focused on the receiver. As for the transmitter, we will start with a system level analysis, aimed at understanding the different equalization schemes proposed in the literature. With the help of a Simulink model, an architecture will be proposed. The transistor level analysis of the aforementioned architecture will follow and will be supported by transistor level simulations of the receiver alone and of the complete transceiver, along with the digital control part.

Finally, an experimental characterization of the full link will be presented, analyzing its performances with measurements performed in the design center of Infineon Technologies, Villach (A).

Contents

List of Publications 159

Acknowledgments 161

Chapter 1

Introduction

1.1 Why do we need faster High Speed Serial Interfaces?

Over the last decades, the innovation in semiconductor technology has been driven by Moore's Law, which states that the number of transistor in an integrated circuit doubles roughly every two year. A direct consequence of Moore's law is the continuous shrinkage of the feature size of the electronic devices and therefore a higher cut-off frequency for the devices themselves, which in turn enables a higher operation frequency for the integrated circuits and lower power consumption for each logic function to be performed. A natural consequence of all these points is an increased number of functionalities packed into a single processing unit and thus an increased amount of data to be stored and exchanged inside a single chip or among different chips.

There are two possible ways to achieve higher chip communication bandwidth: increase the number of serial interfaces or increase the transmission rate of the single interface. Increasing the number of I/O pins is in most cases not a viable options, both for padring space and cost (increasing the number of I/O pins means increasing the number of traces on the package and on the board, thus more metal and material to be placed on chip) reasons. In conclusion, increasing the data transmission rate of a single channel is a must for chip-to-chip communication, even if doing so is complicated. In fact, aside from the technological improvement of the devices, the bandwidth of the transmission channel is an intrinsic limit that has to be overcome, possibly containing at the same time the power budget of the link.

1.2 Evolution over the years

Over the past years, the need for exchanging data has led to the pervasive presence of high speed serial interfaces (HSSI) in many application fields [1], for instance:

- Telecommunication networks, e.g. IEEE Ethernet Standards [2]- [3]
- Computing units with wire-lined I/Os, e.g. Peripheral Component Interconnect Express (PCIe) [4], Hypertransport [5], or inserted in wireless networks, linking the radio equipment control and the radio equipment in wireless base stations, e.g. Common Public Radio Interface (CPRI) [6]
- Interfaces inside optical networks, e.g. Interlaken [7] and OIF-CEI [8]
- Chip-to-chip and board-to-board links
- Storage Applications, e.g. Serial Advanced Technology Attachment (SATA) [9]
- High-performance embedded processing, e.g.Serial Rapid IO (SRIO) [10]

The variety of fields in which high speed serial links are nowadays used had brought to the adoption of many standards [11], sometimes conflicting between each other. In order to compare different standards, common definitions both for channel specifications and signal parameters [12]- [13] have been created. The speed of serial links in these fields reflects their different stages of maturity reached, as depicted in Fig. 1.1. The common trait is that the communication speed keeps increasing for all different applications, lagging some years behind state-of-the-art academic publications (see Fig. 1.2). Along with the increase in speed, there has been a corresponding improvement in power efficiency as shown in Fig. 1.3, even though channel limitations have become more and more severe going towards higher transmission rates.

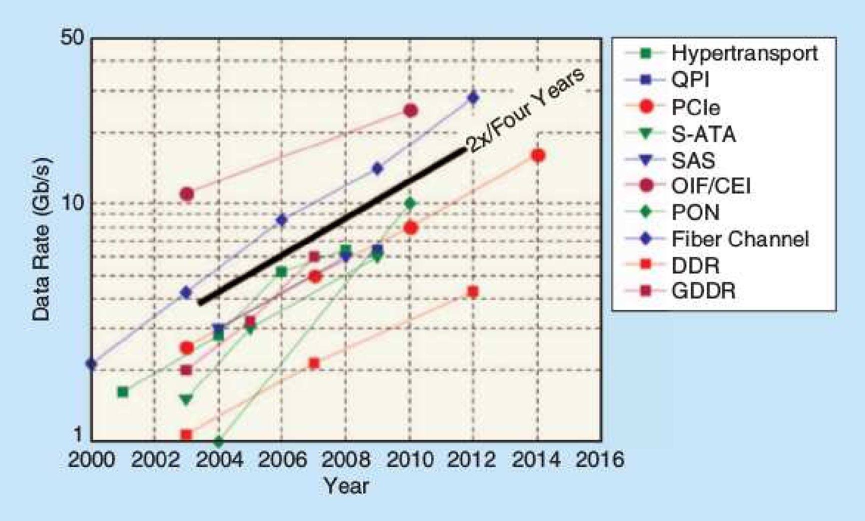

Figure 1.1: Trend in emerging I/O standards, showing data rates doubling every four years [14]. QPI: QuickPath Interconnect; PCIe: Peripheral Component Interconnect Express; S-ATA: Serial AT Attachment; SAS: Serial Attached Small Computer System Interface; OIF/CEI: Optical Internetworking Forum/-Common Electrical I/O; PON: passive optical network; DDR: double data rate memory; GDDR: graphics double data rate memory.

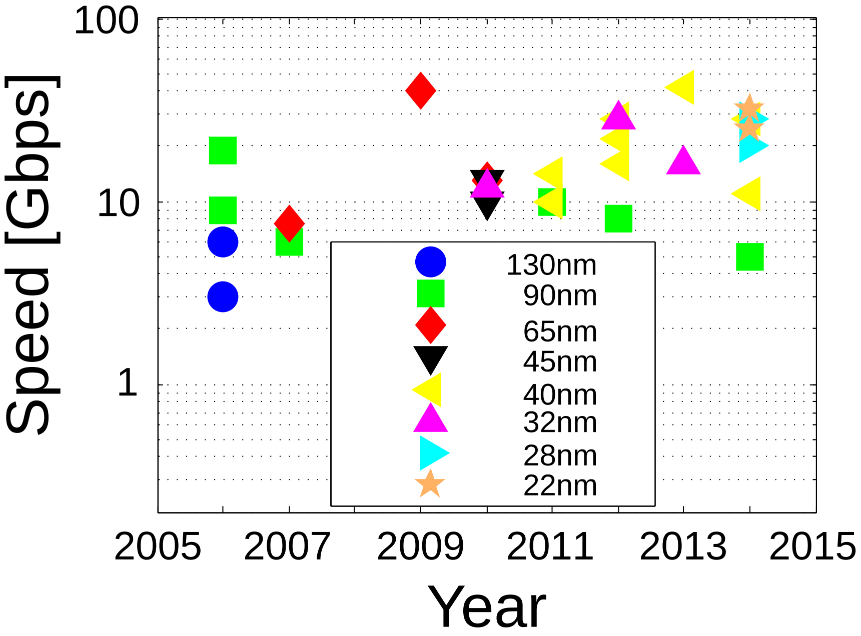

Figure 1.2: Trends in digital I/O publications at the International Solid-State Circuits Conference for different technology nodes, showing a clear move toward high I/O data rates in more advanced CMOS technologies over time. Elaboration from a figure present in [14].

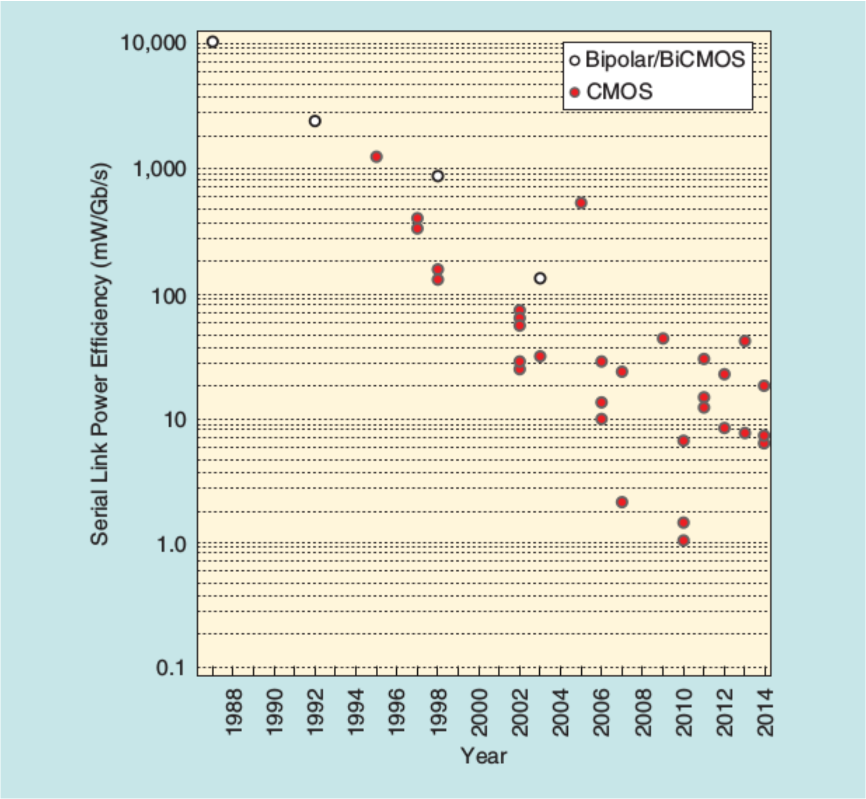

Figure 1.3: Serial Link Power Efficiency versus time both for hybrid BiCMOS/CMOS and CMOS technologies [15].

1.3 Architecture of a Serial Link: Main Problems and Proposed Solutions

Fig.1.4 shows the general structure of a serial interface [16], including transmitter, channel and receiver. The first block we encounter is the serializer, which takes the parallel data from the digital side and outputs them serially to the driver. The driver generates an output voltage swing on the channel while at the same time having an output impedance matched to the channel to avoid reflections. The timing of the serializer and the driver itself is obtained via a clock reference and a Phase Locked Loop (PLL), which takes the low frequency of the crystal and multiplies it accordingly to the data rate frequency and to the one needed by the digital (normally, ten to one hundred times smaller than the data rate). After the channel, at the input of the receiver there is a slicer, which samples the actual received voltage value and compares it with a threshold to decide whether it is a '1' or a '0'. After this, the data is regenerated to CMOS values and is passed to the deserializer, which takes as an input the serial data stream and outputs parallel data to the digital post-processing. The sampling position of the slicer is fixed via a Timing Recovery Circuit, which aligns the clock with the incoming data stream so to have it positioned where it's easier to correctly distinguish between a '1' and a '0' value. In the following, all three main blocks composing a transceiver, transmitter, channel and receiver, will be analyzed in detail.

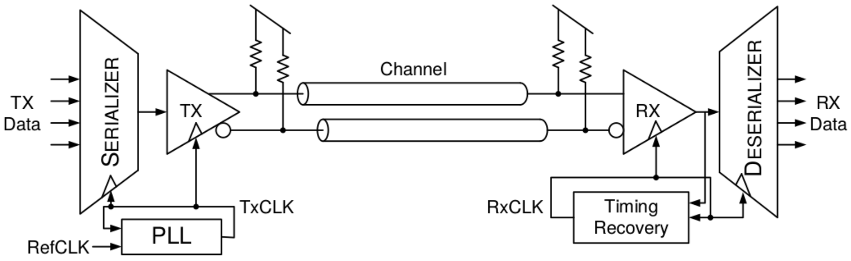

Figure 1.4: General Architecture of a high speed serial interface, including transmitter, channel and receiver [16].

1.3.1 Structure of the Transmitter

The general structure of a transmitter is reported in Fig. 1.5. Before transmitting the data as voltage levels via the driver, some digital preprocessing is usually performed by encoding the data, procedure that solves the problem of having a transmitter and a receiver block working on different DC levels [14]. In fact, this problem has to be faced every time two chips are communicating between each other, especially when they are fabricated in different technologies. The easiest solution to this problem is AC coupling, shown in Fig. 1.6. Anyway, using this structure, the ac coupling capacitance C_C and the receiver input termination resistance R_T form a high-pass filter, which in turn gives a slow loss of low-frequency components in the signal traveling along the channel. This loss results in a slowly drifting signal on the line whenever a long series of consecutive identical bits are to be transmitted: this issue is called baseline wander. This problem can be tackled in various manners: the easiest one is to use a very large AC capacitance so to have a very low cut-off frequency for the high-pass filter. Since this solution alone theoretically does not solve the problem, what can be done is to encode the data so to avoid the presence of extremely long sequences of identical bits, e.g. with the "8b10b encoding" [17]. This famous encoding scheme assures that no more than five identical bits in a row will be transmitted, but doing so it introduces a significant 25% overhead. Going towards higher data rates, more relaxed coding schemes introducing less overhead are now being used in many standards(e.g., 64b66b encoding in 10 Gigabit Ethernet [2]).

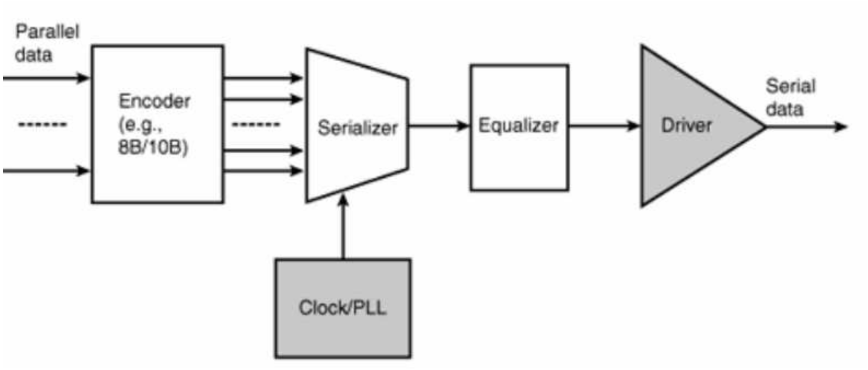

Figure 1.5: Typical structure of a high speed transmitter [18], including both digital preprocessing, analog transmission of the data and clocking concept.

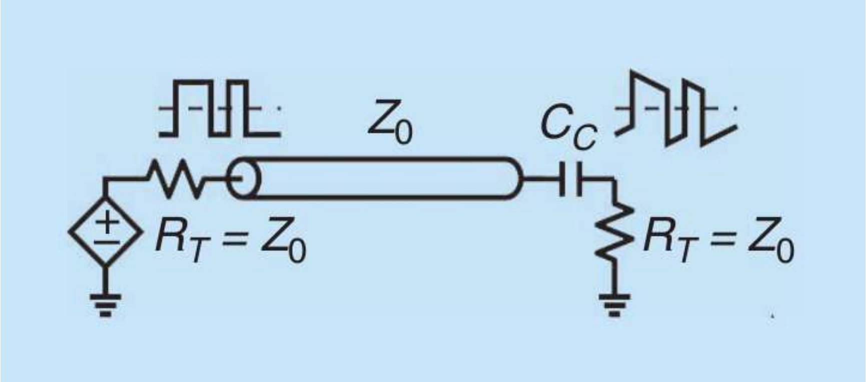

Figure 1.6: To simplify the communication between two chips ac coupling is common, introducing a baseline wander in the received waveform [14].

Three are the main characteristics of the driver, which is the piece of the transmitter which outputs the analog levels on the channel: its mode (voltage or current, as we will discuss in extent in Chapter 2), its output swing and its output impedance. As previously said, it is of paramount importance to match the output impedance of the transmitter, especially going towards higher data rates. At these frequencies, the wavelength of the transmitted signal is comparable with the physical length of a chip-to-chip communication link, therefore it is common practice to design the transmitter to match 50Ω, which is the typical channel impedance for most serial links applications. Thus, given the fact that the output impedance to be matched is fixed, there is a strong interest in reducing the transmitter output signal swing in order to cut down the power dissipation on the termination resistance itself, given by $\frac{V^2}{R}$, where V is the voltage swing and R is the termination impedance. This has led to the development of various transmission standards based on Low Voltage Differential Signaling (LVDS), meaning by this that signals having low swing are traveling on the line.

1.3.2 Clocking Schemes

The clocking circuitry is an important part of a serial link, which varies depending on the adopted clocking strategy. Based on this, we can distinguish among four different classes of interfaces [14]:

- *Synchronous*, in which the clocks at the transmitter and at the receiver have exactly same frequency and same phase with respect to the data. This solution is almost never used at high data rates, as the wavelength at Nyquist frequency is comparable to the physical distance to go;

- *Mesochronous*, in which transmitter and receiver clock have exactly the same frequency, but different phase with respect to the data;

- *Plesiochronous*, in which transmitter and receiver clocks have almost the same frequency, but not precisely, and different phases with respect to the data;

- *Asynchronous*, where there is no relation whatsoever between the transmitter and the receiver clock. In links adopting this solution, normally the receiver does not even know about transmission rate.

A common solution to relax the clocking circuitry, whatever the adopted clocking scheme is, is to use so-called *half-rate* architectures for transmitter and/or receiver. In half-rate architectures, the data are processed at a frequency which is the half of the data rate, but they are usually transmitted/received at full-rate, meaning that they travel along the transmission line at full-rate. Using parts of the transceiver at half of the speed means that also a half-rate clock has to be generated, which is less critical to manage than a full-rate one. This is done in order to relax the timing of some critical blocks in the circuit (e.g., flip-flops) that would be otherwise difficult to design for full-rate frequency. Using half-rates architectures has also the big advantage to lower the power consumption, since all the clocking distribution circuitry operates at half frequency. In the next chapters we will dive more into this concept, both for transmitter and receiver architectures.

1.3.2.1 Clock and Data Recovery

It is easily understandable that, whatever the clocking scheme we are using is, at high frequencies the relation between data and clock at the transmitter side gets lost in the transmission, and it must be reconstructed at the receiver side. In fact, if the correct data-clock phase relation is not recreated, then it could happen that at the first sampling stage in the receiver wrong data are sampled, even though correct data have been transmitted, just because the sampling time is not the correct one. The circuit that takes care of aligning data and clock is the Clock and Data Recovery (CDR) [19], or as it is sometimes called Timing Recovery as in Fig. 1.4.

Since the job of the CDR is to align two phases, then a Phase Detector must always be present in it. As a particular study-case, we will now analyze the Alexander Phase Detector, one of the most commonly used Phase Detectors in High-Speed Interfaces. In order to extract phase informations, a transition in the data is needed: only when a transition occurs, the Phase Detector can take a decision (Fig. 1.7).

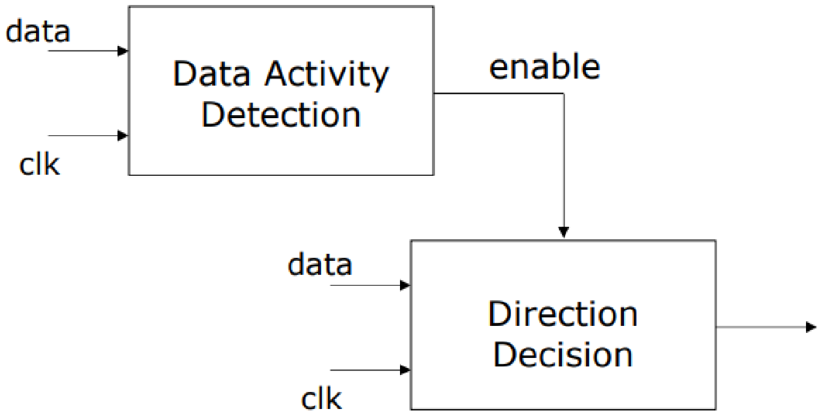

Figure 1.7: Generic scheme of a Phase Detector with Data Activity Detection.

The easiest and most common way to implement a Data Activity Detector is shown in Fig. 1.8: the data are fed into a sampler chain and two consecutive bits are sampled ($act[n-1]$ and $act[n]$); if the two are different, than a transition has occurred and the enable signal goes high.

Once the direction decision block has been enabled, the edge sample enters in the picture. All possible cases when the enable signal is high are reported in Fig. 1.9: $dir[n]$ is the edge sample, which is sampled at the falling edge of the clock, thus in anti-phase with the data. If the data transition occurs after the clock falling edge, then it means that the data have been sampled too early, therefore the decision will be to delay the next sampling point ($updw = +1$). If the data transition occurs before the clock falling edge, then it means that the data have been sampled too late, therefore the decision will be to anticipate the next sampling point ($updw = -1$).

The most straightforward way to implement the Direction Decision and to integrate it with the Data Transition Detection is shown in Fig. 1.10. The

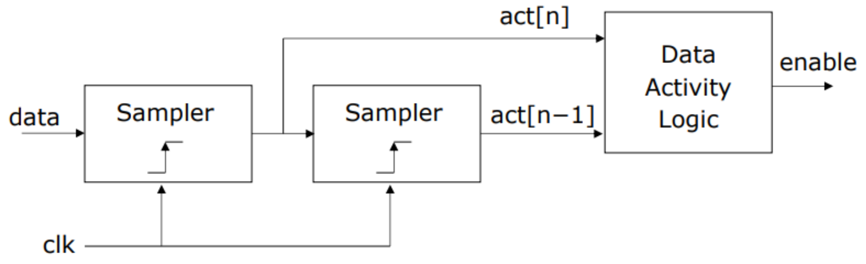

Figure 1.8: Architecture of a data activity detector as implemented in an Alexander Phase Detector with full-rate clock.

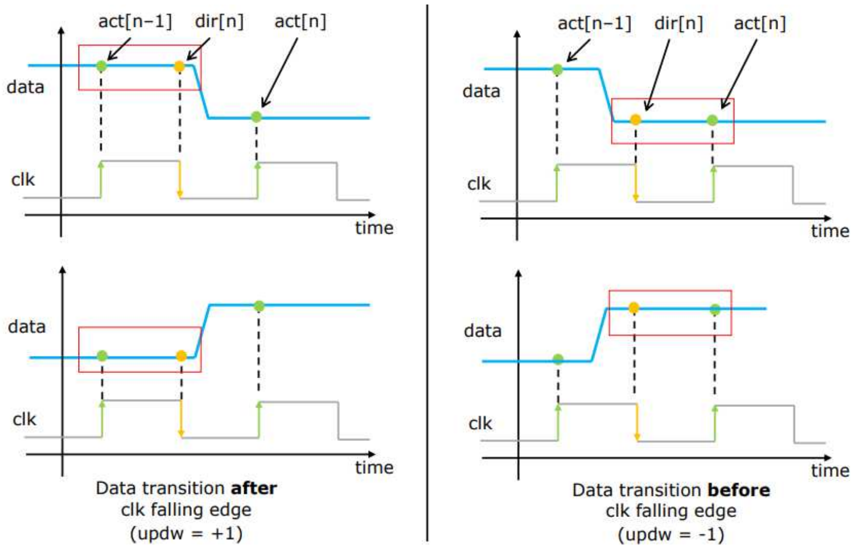
act[n−1]
dir[n]
act[n]
data
clk
time
act[n−1]
dir[n]
act[n]
data
clk
time
data
clk
time
Data transition **after** clk falling edge (updw = +1)
data
clk
time
Data transition **before** clk falling edge (updw = -1)

Figure 1.9: All possible cases of direction decisions when a data transition has been detected in an Alexander Phase Detector with full-rate clock.

sampler that works on the rising edge of the clock in the direction decision is necessary to make sure that the output $dir[n]$ is the sample of $data$ associated with the falling edge of clock that occurred before the n-th rising edge of clock.

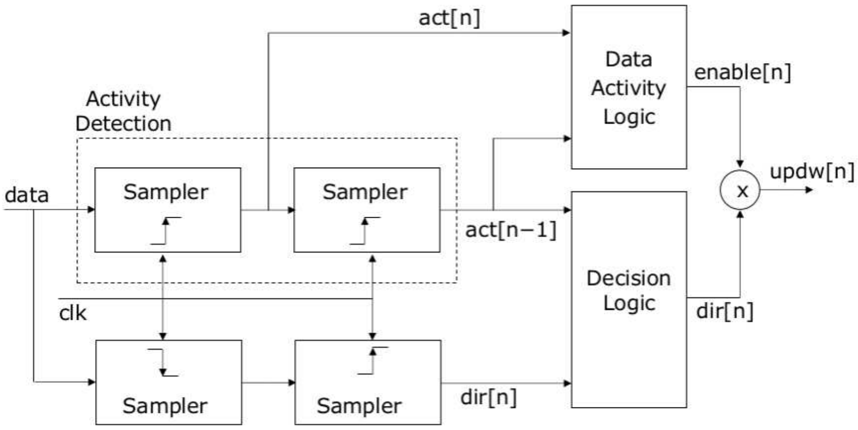

Figure 1.10: Block scheme of an Alexander Phase Detector, including both Data Activity Detection and Direction Decision circuitry.

The Alexander Phase Detector is a so-called $Bang - Bang$ phase detector, by this meaning that it will take a direction decision every time it recognizes that a data transition has occurred. In a Bang-Bang phase detector, a stable operating point is when the system is in a limit cycle at the highest possible frequency, so when the Bang-Bang output ($updw[n]$ in Fig. 1.10) swaps between -1 and 1 every clock cycle. In a CDR, the above limit cycle can thus occur only when the sampled data, $act[n]$, swaps between 0 and 1 every clock cycle, hence the loop can correct for a phase error every clock cycle and the edge sample, $dir[n]$, swaps between the value of $sct[n]$ and $act[n-1]$ every clock cycle.

Three are the most common solutions used in high speed links to embed the CDR in the transceiver: *oscillator-based recovery*, *phase interpolator-based recovery* and *forwarded clock architecture* [20].

1.3.2.2 Oscillator-Based Recovery

Fig.1.11 shows the scheme of a serial link with controlled oscillator(CO)-based clock recovery circuitry. This architecture falls under the asynchronous class, as the clock is generated thanks to a crystal reference and a PLL at the transmitter and then used for the data transmission timing. After the transmission, at the receiver side the Clock and Data Recovery (CDR) circuit reconstructs the clock from the timing of the received data by means of controlling a VCO (Voltage CO), both in frequency and phase, so to align the sampling clock with the data stream.

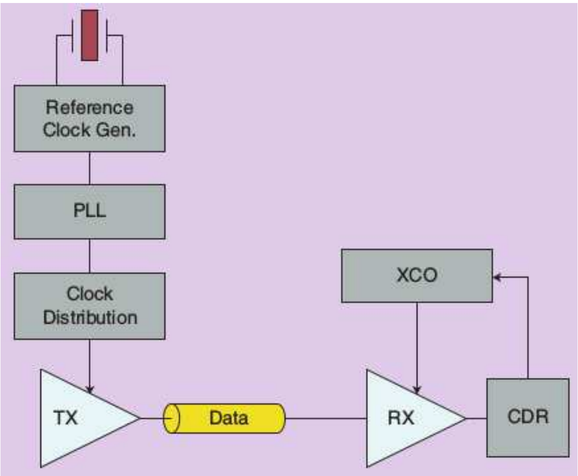

Figure 1.11: Scheme of a serial link with controlled oscillator-based clock recovery circuitry [20].

1.3.2.3 Phase Interpolator-Based Recovery

Fig.1.12 shows the architecture of a serial link with CDR based on phase interpolator. In this circuit, two clocks with nominally identical frequencies are generated at the transmitter and at the receiver thanks to two crystal references, which might differ inside the boundaries defined by their quality specifications, expressed in Parts Per Million (PPM). This means that this circuit falls under the category of plesiochronous systems. The CDR circuit at the receiver aligns the phase of the clock locally generated via a crystal and a PLL with the incoming data stream so to set an optimal sampling point. There are two main ways to do this, with a Delay Locked Loop (DLL) and with a clock divider and a Phase Interpolator (PI).

The DLL is an inverter chain that generates different delays with a separation step defined by the number of inverters inserted in the chain. If the separation step is fine enough, then one of the phases extracted from the DLL can be directly used to be the sampling clock.

In the second solution, the clock coming out from the PLL enters a clock divider, that outputs a clock with four different phases (0°, 90°, 180° and 270°, the so-called I-Q phases). These four phases then enter a PI, that works as a weighted summer between two of the four input phases [21]-[22], so that the clock at its output will have a phase somewhere in between these two.

This is done by means of a Delay Locked Loop (DLL) and a Phase Interpolator (PI). The DLL is an inverter chain that generates different delays with a separation step defined by the number of inverters inserted in the chain. Usually, four quadrature phases are extracted from the DLL, and then the PI interpolates among these four to obtain the correct phase to be applied to the clock.

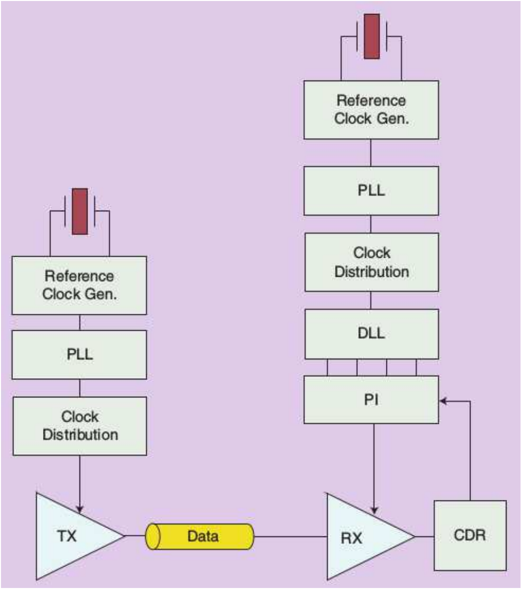

Figure 1.12: Scheme of a serial link with phase interpolator-based clock recovery circuitry [20].

1.3.2.4 Forwarded Clock Architecture

Fig.1.13 shows the scheme of a serial link with forwarded clock. In all previously analyzed schemes, the data was the only information transmitted to the receiver, whereas in this case both data and clock are sent over the channel. The clock is generated at the transmitter thanks to a crystal reference and a PLL, then sent with a different driver to the receiver along with the data. At the receiver side, the CDR then aligns the phase of the local clock with the incoming data stream, but with respect to the system described in Fig. 1.12 the difference is that the local clock is not generated by a crystal and a PLL, but it's the forwarded one. This circuit falls under the mesochronous systems category, but it's extremely impractical to be used at high data rates mostly because matched clock and data latency is not achievable. This in turn brings to a degradation of the CDR bandwidth, the highest jitter frequency the CDR can track. Moreover, using a different driver to transmit the clock requires more area and power, making this approach not so appealing.

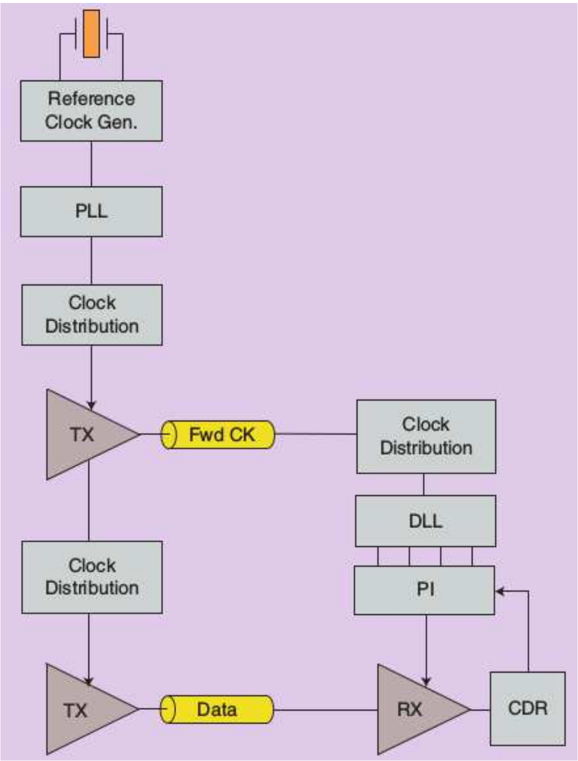

Figure 1.13: Scheme of a serial link with forwarded clock architecture [20].

1.3.3 Noise and Interference Sources in Serial Links

Fig.1.14 shows the cross-section of a complete backplane link and all elements that can bring interference and noise into the system. We have already introduced the bandwidth limitations given by the channel characteristics, but there is much more than that. A first rough distinction can be made dividing all disturbances into noise and dispersion characteristic of the link.

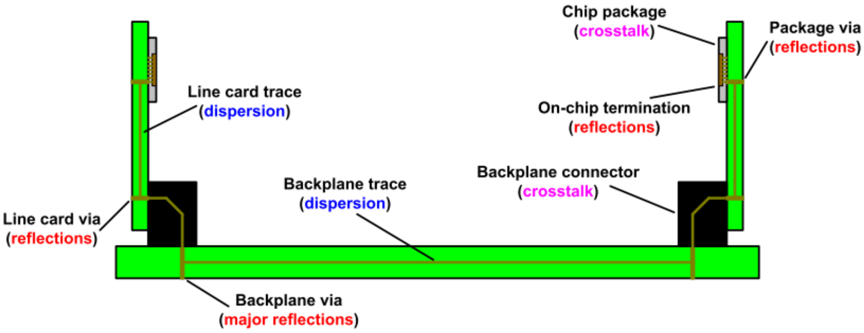

Figure 1.14: The cross-section of a complete backplane link, in which all elements of possible disturbances (reflections, crosstalk and dispersions) are highlighted [16].

1.3.3.1 Noise Disturbances

Noise sources can be roughly divided in two main category, *crosstalk* and *random signal fluctuations*. *Crosstalk* is a mechanism of coupling between two neighboring lines that can occur both via capacitance and mutual inductance, as can be seen in Fig. 1.15. In crosstalk two parts can always be identified, an *aggressor* and a *victim*. Depending on the position of these two, we refer to Far-End Crosstalk (FEXT) or Near-End Crosstalk (NEXT). The latter is the most dangerous one, because the aggressor is on the same chip as the victim and near to it, therefore it is a direct coupling mechanism. In FEXT, the energy of the aggressor travels along the channel and then couples at the receiver's side with the victim, being therefore already attenuated by the medium. *Random signal fluctuations* is always present in the system due to thermal and shot noise of active and passive devices and results in timing deviations of the signals, causing crossing edges to be anticipated or postponed. These deviations fall under the definition of *jitter*, which we will often talk about in the following.

1.3.3.2 Dispersion of the Transmission Medium

In this subsection we will analyze the dispersion and loss mechanisms inherent to the transmission medium, namely skin effect, dielectric loss and impedance mismatch and discontinuities.

Skin effect is a form of channel insertion loss, given by the fact that at high frequencies the effective cross-section of the wire in which the current flows

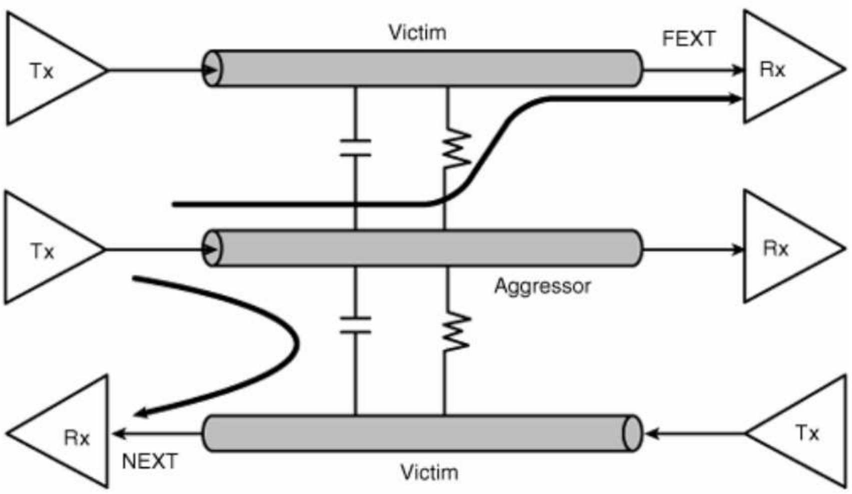

Figure 1.15: Two different crosstalk mechanisms, Near-End Crosstalk (NEXT) and Far-End Crosstalk (FEXT) [18].

reduces and therefore current is flowing just in the proximity of the surface. This effect can be modeled into a frequency dependence of the impedance per unit length of the channel that reads as

$$R(f) = R_0 + \sqrt{f} \cdot (1 + j)\, R_S \tag{1.1}$$

where R_0 is the channel DC loss per unit length and R_S is a skin effect constant. The dependence upon $\sqrt{f}$ means that the skin effect causes a loss of 10dB/decade.

Dielectric loss is given by the absorption of energy by the dielectric medium from the traveling wave, which is in turn transformed into heat. This effect can be summarized into an imaginary part of the permittivity ϵ of the dielectric medium that reads

$$\epsilon\,(f) = \epsilon^{'}\,(f) - j\epsilon^{''}\,(f) \tag{1.2}$$

where j is the imaginary unit. Eq. 1.3 can be rewritten in the form of an admittance to ground as

$$Y_{\perp}\,(\omega) = j\omega C + \omega C \tan\,(\delta) \tag{1.3}$$

where C is the transmission line's capacitance at low frequencies and δ is the loss angle. The dielectric loss results in a loss dependence of 20 dB/decade at high frequencies, where $Y_{\perp}\,(\omega)$ gets high.

The third loss mechanism to analyze is impedance mismatch and discontinuities. Packaging, sockets, connection to ac-coupling capacitors, PCB vias and connectors of different nature all introduce discontinuities in the channel's impedance [23], that in turn generate signal reflections that show up as notches in the pulse response of the system. Another effect of impedance mismatch is *mode conversion* [24]: part of the energy of a differential signal is transformed into a common mode component, that in turn may cause reflections.

In a generic channel, all these effects are coupled together and show up in the pulse response of the channel. In order to identify which loss mechanism is the dominant one in a channel, one might look either at the magnitude of the transfer function (also called S_{21}) or at the pulse response of the system, since they are strictly related to each other: both these metrics are shown in Fig.1.16. To better understand these pictures, we have to introduce the concept of InterSymbol Interference (from now on, ISI). ISI is a form of distortion of a transmitted bit in which one bit interferes with the subsequent and/or the previous ones. The amount of ISI of a channel can be derived by looking at its pulse response (Fig. 1.16 (b)): all residuals in the UIs before and after the one at which the pulse is transmitted interfere with the following bits. All residuals in the UIs before the current bit are called *pre-cursors*, all the ones after the current bit are called *post-cursors*. The current bit is often referred to as *main-cursor*. ISI is the reason why usually defining just the loss at a certain frequency is not enough in order to decide for an equalization strategy, and this can be better understood by looking at Fig. 1.16: in fact, all S_{21} shown in the figure present a loss of 25dB at half the transmission frequency (the so called Nyquist frequency), but the rest of the spectrum may vary very much from one interference mechanism to the other. In Fig. 1.16 b), the normalized pulse response (i.e. a pulse response scaled so that the absolute sum of all pre-, post- and main-cursors is 1) of the S_{21} is shown.

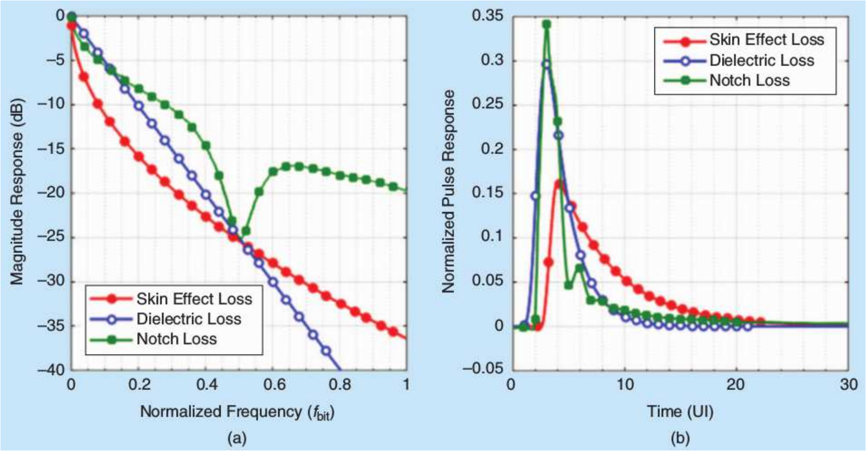

Figure 1.16: On the left, S_{21} (magnitude of the transfer function) versus normalized frequency ($\frac{f}{f_{bit}}$, where f_{bit} is the transmission frequency). On the right, normalized pulse response (i.e. the pulse response is scaled so that the absolute sum of all pre-, post- and main-cursors is 1) of the channel versus time, normalized in UI (1 UI equals to one bit period, thus 1UI=$\frac{1}{f_{bit}}$). Both graphs are shown for three different channels, one dominated by skin effect, one by dielectric loss and one by impedance discontinuities.

1.3.4 Equalization

As mentioned in the previous section, especially when moving towards extremely high data rates, the transmitting medium can completely kill the ability of a transceiver to establish even a Non-Return-to-Zero (NRZ) communication. What the designer can do is to foresee such detrimental effect of the channel and to compensate for it: this compensation is what in the literature is called equalization [25]- [26]. The main idea behind equalization is shown in Fig. 1.17: the "non-flatness" in the transfer function of the transmitting medium are compensated for by means of an inverse transfer function. Assuming the non-ideality in the transfer function to be proportional to the frequency (i.e., given by skin effect or dielectric loss), it is identical to apply the correction either at the transmitter or at the receiver. If the behavior of the non-ideality versus frequency is not merely proportional, but with various notches (i.e. coming from impedance discontinuities), then it's possible to correct for it either via digital techniques (e.g., with a Finite Impulse Response filter) or with non-linear ones (e.g., Decision Feedback Equalization). In the following subsections we will analyze the most used equalization techniques that can be found nowadays in serial links and are listed in Fig. 1.18.

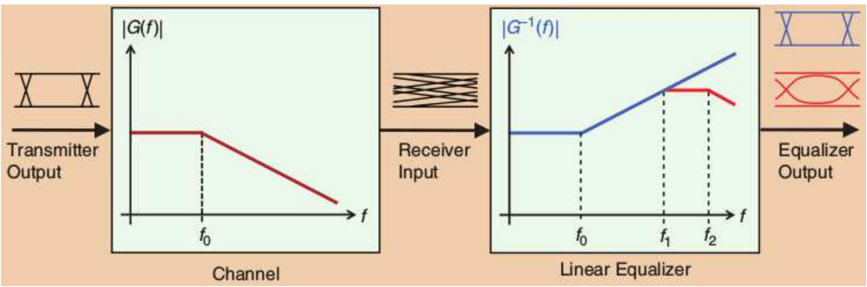

Figure 1.17: Transfer function of the channel, $G(f)$ and inverse transfer function, $G^{-1}(f)$, applied as an equalization technique, along with the eye diagram at the output of the transmitter, at the input of the receiver and after equalization [26].

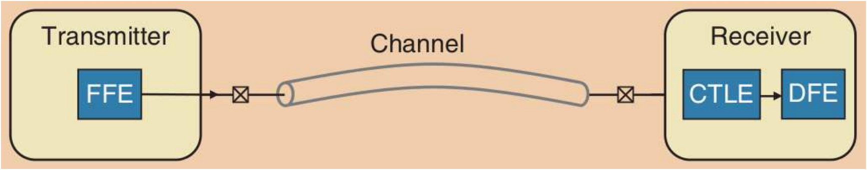

Figure 1.18: General structure of a serial link in which the most common equalization techniques have been highlighted [26].

1.3.4.1 Feed-Forward Equalization

Equalization at the transmitter is done via preconditioning of the signal before setting the voltage level at the input of the channel. In order to do so, one has two possible choices: amplifying the high-frequency components (*pre-emphasis*) or reducing the low-frequency ones (*de-emphasis*), since the wanted effect is anyway the one of a high-pass filter to counteract the low-pass characteristics of the channel. This can be easily done by multiplying delayed versions of the bit stream to be transmitted (the so called *taps*) by weighted coefficients. Delayed data can be easily created since the clocking circuitry is always present at transmitter side. This approach is sketched in Fig. 1.19 and will be vastly analyzed in the following chapters.

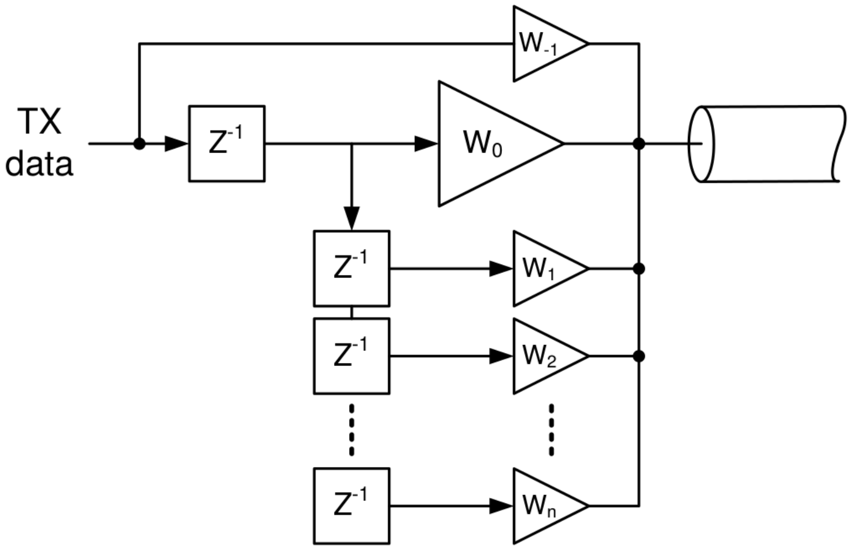

Figure 1.19: Conceptual scheme of a feed-forward equalization with one pre-cursor tap, the main tap and *n* taps, created via a flip-flop register chain [16].

As can be understood, this approach can be applied both at transmitter and at receiver side [27]-[28]-[29]- [30]. Anyway, usually such FIR filter is implemented at transmitter side (FFE) because at this point of the transmission chain the system is still working with digital bits, '1' or '0', coming directly from the digital encoding and thus easy to create delayed versions of just with a simple register chain, while the multiplication by suitable coefficients is realized by means of switching circuits. On the other hand using this approach at the receiver would imply working with analog signals. To avoid this, one could replace the first sampler at the receiver side with a very fast ADC, so to digitize the signal and then apply FFE [31], [32], [33]. As shown in the Fig. 1.19, all taps are summed together at the output of the transmitter, just before entering the channel.

1.3.4.2 Continuous Time Linear Equalization

Continuous Time Linear Equalization (CTLE) is a linear technique that is implemented in the Analog Front-End (AFE) of the receiver. Fig. 1.20 show a typical implementation of this concept and explains what's its aim and how it is achieved. The goal of using CTLE is to compensate for the low-pass characteristics of the channel by introducing a zero at a frequency slightly smaller than the data rate. This is usually obtained with a source-degenerated differential pair, which has a high-pass filter behavior if the zero frequency is designed to be much lower than the dominant pole, which should be the one given by the output parasitics capacitance C_P. This technique works equally well both for pre-cursors and post-cursors, since it directly acts on the transfer function of the system.

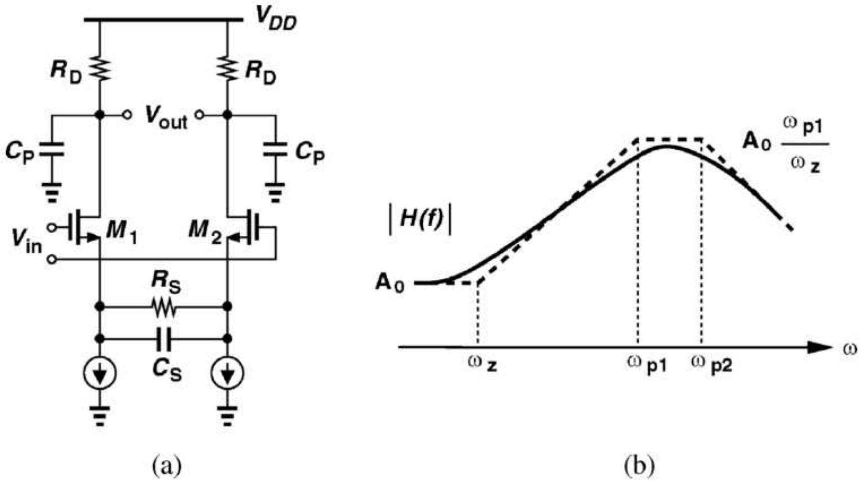

Figure 1.20: Source-degenerated differential pair used as CTLE: a) circuit implementation and b) transfer characteristic, both in its asymptotic approximation (dashed line) and its real version (full line) [34].

The main problem of CTLE is that with a single-stage implementation as the one shown in Fig. 1.20 a) is hard to obtain a high peaking gain. If one wants to have high-gain it should use multi-stage architectures, as many cascaded source-degenerated differential pair, but then it becomes extremely difficult to tune all parasitic poles present in such a structure over broad PVT variation ranges typical of the automotive environment. For this reason, CTLE normally acts together with other equalization schemes and is used as a mean of facilitating the work of other equalization techniques (e.g., DFE).

1.3.4.3 Decision Feedback Equalization

Decision Feedback Equalization (DFE) is a non-linear technique that is applied at the receiver side, usually after a CTLE stage. The idea behind DFE is to directly address Inter-Symbol Interference acting on the incoming analog voltage levels based on the already-detected bits, as shown in Fig. 5.1. The last sentence implies that it's not possible to address any pre-cursor ISI, and therefore DFE is normally used in collaboration with FFE or/and CTLE.

The biggest advantage of DFE is that, contrarily to FFE and CTLE, it does not amplify high-frequency noise, since it works on previously sampled digital bits: for this reason, over the last years researcher started to heavily prefer DFE to FFE. It can be understood also that this equalization technique is particularly effective when it comes to equalize for impedance discontinuities, which as we have seen manifest themselves as notches in the S_{21} curve, as it precisely counteracts for ISI and it is thus able to correct those notches [35]. On the contrary, the main drawback of this equalization technique is error propagation: if the noise at the input of the slicer is so big that a wrong detection occurs, this wrong decision affects in a detrimental way the equalization of the following bits, until a stream of correct bits are recognized and the error is pipe-cleaned out of the shift register composed by N latches for an N-tap DFE. So, while DFE is generally useful to decrease the Bit Error Rate (BER), from time to time it may occur that bursts of errors appear and need to be corrected. For this reason, it is quite common to use forward error correction codes (e.g., Reed-Solomon) in addition to it, to correct for these error bursts.

Another issue in implementing DFE is the stringent timing constraints that it imposes: in fact, the feedback loop corresponding to the first tap has to come in one bit period. While this might not be an issue while transmitting at low data rates, it means that for a receiver that works at 10Gbps the time for closing the feedback loop is 100ps. In order to address this problem, a variety of techniques have been used: the most common solution is to mix decision look-ahead schemes (the so called *loop unrolling* techniques) with half rate ones, to increase even by four times the time for closing the feedback loop, see Fig. 1.22.

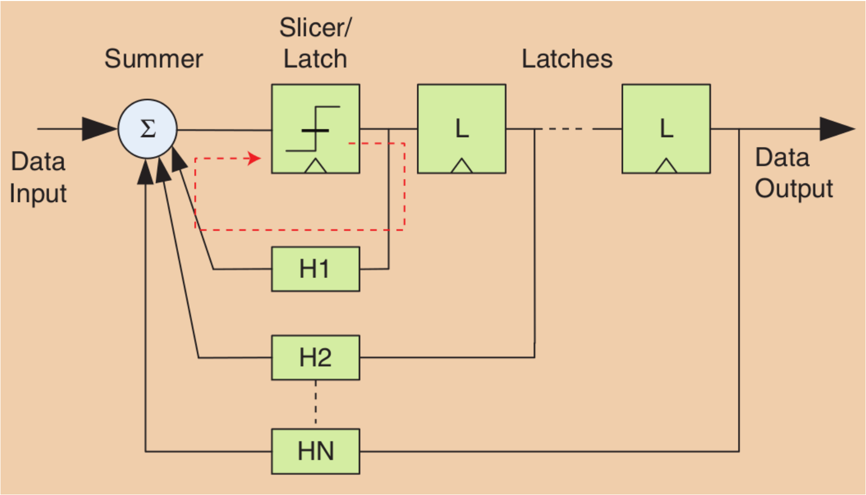

Figure 1.21: Conceptual scheme of a receiver with an N-tap DFE [26]. Highlighted in red there is the critical path whose timing has to be guaranteed at very high data rates.

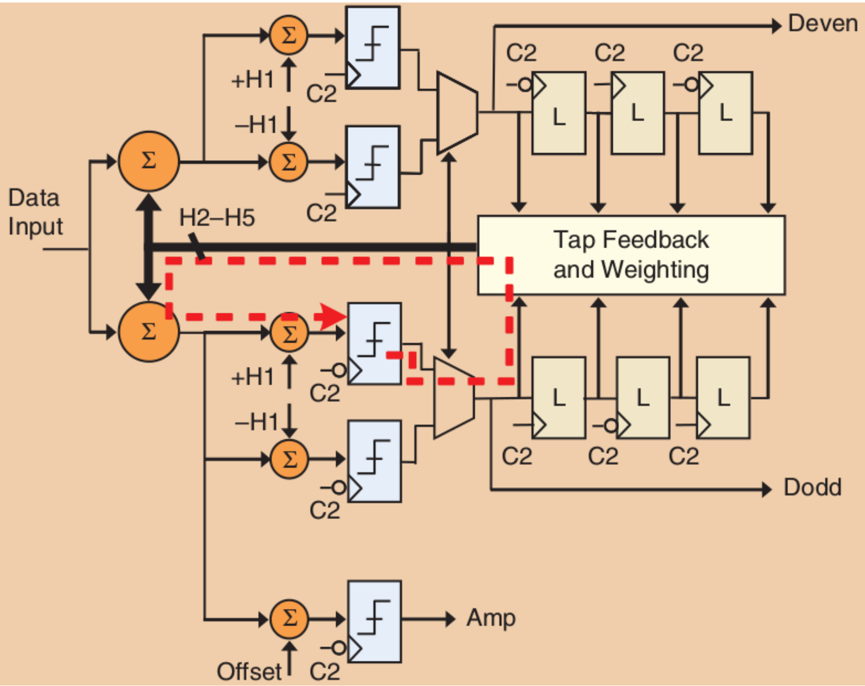

Figure 1.22: Conceptual scheme of a half-rate receiver featuring 5-tap DFE with a 1-tap loop unrolling [26]. The critical path, shown in red, is now two times longer than the one of Fig. 5.1. In the bottom branch of the receiver there is the so called *eye monitoring*, which checks the effective improvement of the eye diagram after DFE is applied.

1.4 High Speed Serial Interfaces in the Automotive Environment

Nowadays, pervasive and smart electronics is present in every aspect of our life, by this meaning that the environment in which we are living is capable to react appropriately in case of foreseeable events thanks to the presence of sensors and actuators. Vehicles have followed this same path, driven by development of actuators, control and monitoring systems for different applications inside the car, ranging from Safety and Infotainment to Drive Assistance, Power-train and Comfort, as briefly summarized in Fig. 1.23.

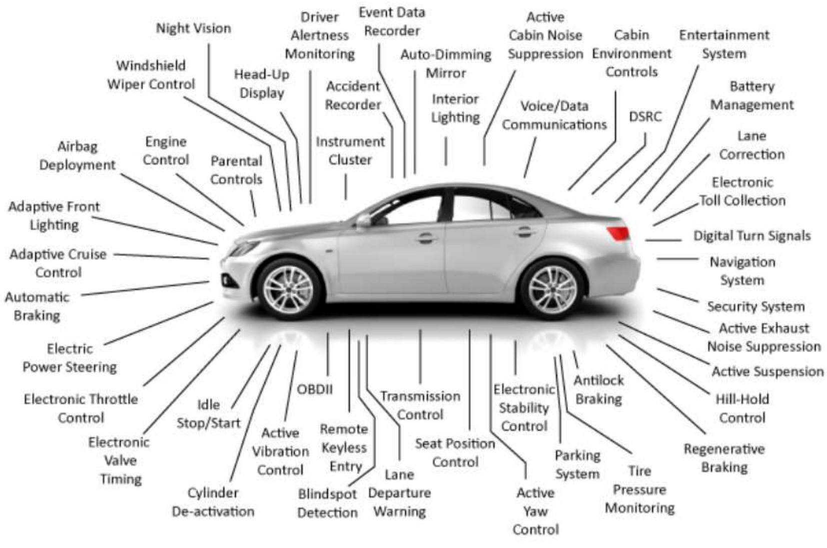

Figure 1.23: Partial summary of electronic-aided functions that can be found in a car nowadays.

The increase of electronics in automotive environment seems unavoidable, and the market to be shared is very likely to continue to grow in the next decades [36]- [37]- [38], especially with the advent of Advanced Driver Assistance Systems (ADAS) and autonomous driving in the upcoming years, not to mention the already ongoing explosion of the infotainment sector. The amount of functions to be performed in vehicles is increasing everyday, thus also the number of Electronic Control Units (ECUs) present in a car is going to further raise. Increasing the number of computational nodes implies increasing the network connection among these nodes, therefore increasing the amount of data to be exchanged and thus strongly enhancing the current data transmission rate.

Nowadays, the most diffused communication standards inside the vehicles are Controller Area Network (CAN), Local Interconnect Network (LIN) and Flexray, but all these standards range from few kb/s to 10Mb/s. It is clear that the applications cited above require much higher data-rate, and the adoption of 1Gbps Ethernet as a standard for automotive goes exactly in this direction and demonstrates a trend that is going to be confirmed in the next years. In

Table 1.1: Requirements of automotive electronics [40].

Parameter	Consumer	Automotive
Temperature	0°C - 40°C	-40°C - 175°C
Voltage	3.3V	80V
Operation Time	1-3 years	up to 25 years
Humidity	Low	0% to 100%
Tolerated Field Failure Rate	<1000ppm	Target: zero failure
ESD	4-8kV	8-15kV

fact, as shown in Fig. 1.24, the increase in data transmission rate in automotive environment follows the one in consumer electronics, but with 10-15 years of delay. So it's quite fair to assume that in few years having serial connections at 10Gbps will be the state-of-the-art for automotive.

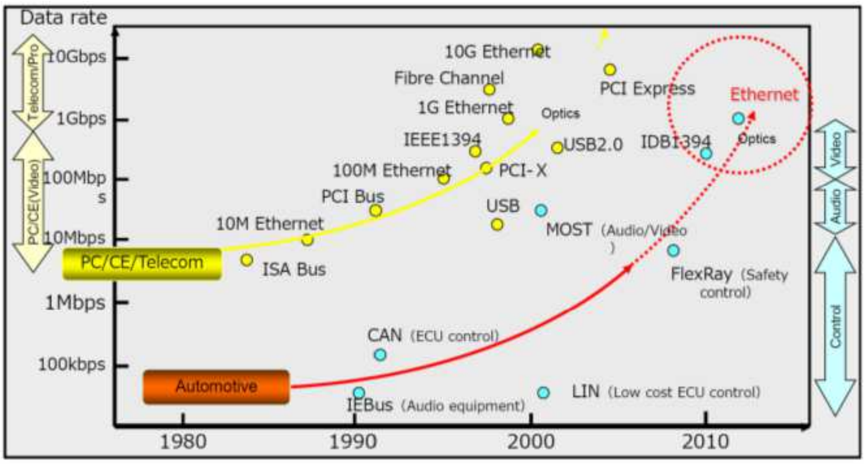

Figure 1.24: Overview of data rates over the last decades for both consumer and automotive electronics [39]. From the graph it can be seen that automotive electronics follows the improvement of the consumer market, but lags behind it by 10-15 years.

Nevertheless, a direct technology transfer from consumer to automotive environment is not possible, due to the serious challenges that the harsh automotive standards pose to the circuit designers, especially to HSSI designers. In fact, requirements in automotive environment are much more stringent than for consumer applications, as summarized in Tab. 1.1.

The requirements that pose the highest challenge to the HSSI designer are for sure the extremely broad temperature range, in which very shrunk devices may see their threshold voltage V_T vary even by 200mV, and ESD protection, which works as a low-pass filter for the transmitter output and receiver input and may therefore be a killing aspect. From what we have stated above, it appears quite clear that designing HSSIs in an automotive environment is a path full of obstacles, even more than in the other markets.

1.5 Motivation of the Work and Thesis Organization

From this introduction, it appears clear that there is a strong interest surrounding the world of high speed serial interfaces, both in the academic and industrial world. At the same time, it is pretty understandable that the automotive sector lags some years behind the consumer market, especially due to the stringent safety requirements and the extremely harsh environment in which the link must be able to correctly operate. Thus, the main goal of this work is to bridge the gap between the consumer electronic and the automotive electronic unit world, understanding which techniques are suitable for our work conditions among the ones that are already well established in the academic world and translating and improving these solutions to possibly make them more stable and less power consuming. This goal implies a deep understanding of a serial link both at system and transistor level, and the development of this thesis will follow this idea. After this introduction, the thesis is divided into chapters, which will be here briefly presented:

- Chapter 2 focuses on the system level design of the transmitter, creating a procedure to assess the equalization strategy at the transmitter side and defining its main parameters by evaluating performance metrics as pulse response and eye diagram;

- Chapter 3 is centered around the transistor level design of the transmitter. Here design choices stemming from literature will be analyzed, justified and explained. The design of a 10Gbps transmitter will be detailed along with the results deriving from schematic and post-layout simulations. After this, a schematic analysis of the effect of parasitic inductance on FFE will be carried out.

- Chapter 4 talks about the experimental characterization of the transmitter above, showing the results of the measurements of the first test-chip and demonstrating the improvements in BER obtained thanks to FFE.

- Chapter 5 deepens on the system and transistor level analysis of the receiver. Here an extensive review of the existing literature will go by with a complete Simulink representation of a half rate receiver with DFE and loop unrolling, so to guide the design choices. Along with this, our CDR implementation will be presented. The transistor level design will be confirmed with post-layout simulations of single receiver blocks (CDR, DFE+CTLE) and schematic simulations of the whole 10Gbps system including transmitter and receiver. Finally, mixed-signal simulation will help to verify the correctness of the CDR and its algorithm.

- Chapter 6 regards the characterization of the receiver, showing the results of the measurements performed in the lab on the second test-chip.

- Chapter 7 wraps up all work done so far, presenting the results of the whole transceiver and illustrating all possible equalization and CDR solutions implemented in our design.

- Finally, an appendix about Electromagnetic Interference (EMI) in High-Speed Serial Interfaces, topic analyzed in the first months of work, will be added in the final version of the thesis.

Chapter 2

System Level Design of the Transmitter

This chapter will focus on the system level design of the transmitter, creating a procedure to assess the equalization strategy at the transmitter side and defining its main parameters by evaluating performance metrics as pulse response and eye diagram.

In the introduction, we have talked about the rise of high-speed links and their increase in speed over the last years. Since the baud rates at which these links work today are extremely high, inter-symbol interference (ISI) has become the most limiting factor [41]. In order to mitigate ISI, the most effective solution is channel equalization [26]. We have seen that equalization can be done either at the transmitter side just before the channel (and it is called Feed-Forward Equalization, FFE) or at the receiver (for example in the form of Decision-Feedback Equalization, DFE). In this chapter we will focus on the system level design of the transmitter: our goal is to create a procedure to analyze the amount of equalization needed at the transmitter side depending on the channel, calculate the weights for such equalization scheme and assessing the impact of tap quantization in a real-world scenario. After this, the correctness of the applied FFE will be evaluated by checking the pulse response and some eye diagram parameters, such as eye height and eye width.

Fig. 2.1 shows the general structure of a transmitter with feed-forward equalization: delayed versions of the bit stream are fed to drivers with different strengths, thus implementing an FIR filter.

2.1 Choice of the Driver Topology

There are two possible ways to implement a high-speed driver: using a current-mode driver (Fig. 2.2(a)) or a current-mode driver (Fig. 2.2(b)). Each of the two implementations presents advantages and drawbacks. Current-mode drivers support higher data rates than voltage mode ones and have a low susceptibility to power supply noise, but they consume four times more power than the voltage-mode ones for the same output swing when a differential termination is used.

In order to demonstrate the last sentence, we will assume that the MOSFETs in Fig. 2.2 have an output resistance $r_o = Z_0$, where Z_0 is the characteristic

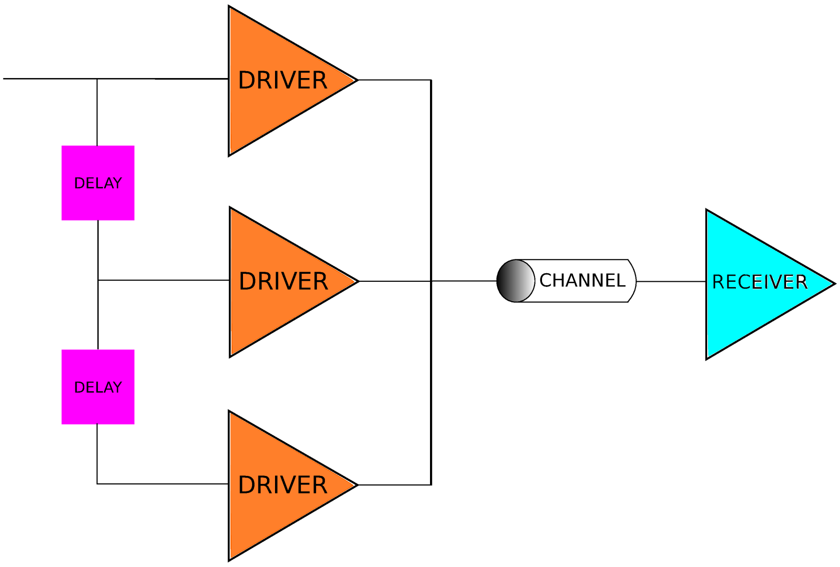

Figure 2.1: General scheme of a serial link implementing FFE using driver slices.

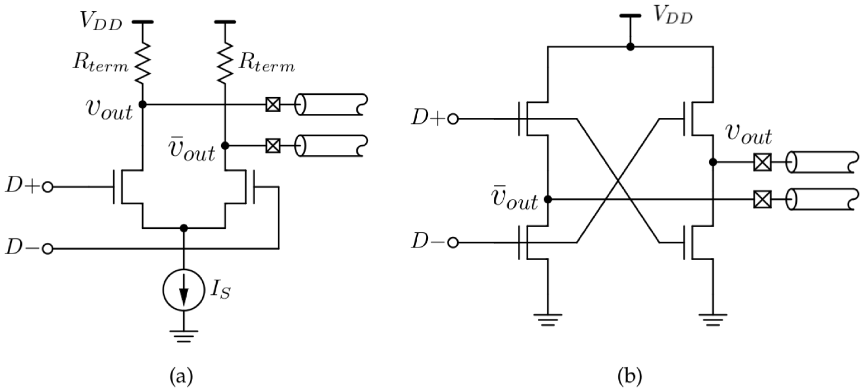

Figure 2.2: (a) Differential current mode driver and (b) differential voltage mode driver topologies.

impedance of the channel (usually 50Ohms). An alternative to this assumption (which would in fact imply using very small MOSFETs for the driver) is to insert a termination resistance $R_{term} = Z_0$ between the MOSFETs and the channel and using MOSFETs big enough so that their resistance can be assumed to be negligible when they are on. If one of this two assumptions is met, when using a differential termination $2R = 2Z_0$, for a current-mode driver one finds a differential output voltage that reads

$$\begin{aligned} V_{diff,out} &= V_{out} - \bar{V}_{out} \\ &= I_S(Z_0||3Z_0)\frac{2Z_0}{3Z_0} - I_S(Z_0||3Z_0)\frac{2Z_0}{3Z_0} \\ &= I_S \cdot Z_0 \end{aligned}$$

from which it follows that the power dissipation for a current-mode driver is

$$P_D = V_{DD} \cdot I_S = \frac{2V_{DD}V_{diff,out}}{Z_0} \tag{2.1}$$

whereas under the same hypothesis for a voltage-mode driver one finds

$$\begin{aligned} I_{DD} &= \frac{V_{DD}}{4Z_0} \\ V_{diff,out} &= \frac{V_{DD}}{2} \end{aligned} \tag{2.2}$$

from which it follows that the power dissipation is

$$P_D = I_{DD}V_{DD} = \frac{V_{DD}V_{diff,out}}{2Z_0} \tag{2.3}$$

Thus, from Eqs. 2.1,2.3 it follows that in order to achieve the same output voltage swing $V_{diff,out}$ a current-mode driver consumes four times more power then a voltage-mode one [15].

2.2 Driver Architecture

For the transmitter, a voltage mode architecture has been chosen in our project, because it consumes less power than a current mode for the same output swing, as stated in the previous section. Fig. 2.3 shows the general structure of a voltage-mode differential driver: two inverters are driven by bit b_i and to produce a differential voltage v_o. If the MOSFETs are large enough to have a negligible voltage drop when on, therefore negligible resistance, then impedance matching is implemented via the resistances R_{D_i}.

Fig. 2.3 is also the starting point to explain the basic principle of FFE applied to voltage mode drivers. A single driver can be split into many slices, and the same bit b_i of the serial data input can drive many of these slices, i.e. the different drivers in Fig. 2.1 are each formed by many slices with the schematics of Fig. 2.3. This can be better understood in Fig. 2.4. The index i of a bit indicates the position in the bit stream.

When dividing the driver into slices, one must still guarantee impedance matching. In fact, when FFE is not implemented, only one slice might be used, and its output resistance should match R_0=50Ω. If more slices are used, it has to be considered that the total resistance of all the slices in parallel should be 50Ω, leading therefore to bigger values for R_{D_i} of each slice. This results in the following equation

$$\sum_i \left(\frac{1}{R_{D_i}}\right) = \frac{1}{R_0} \tag{2.4}$$

Under this assumption, one can find a closed-form expression for the output voltage v_o starting from an equivalent circuit for the half-driver as depicted in Fig. 2.5. In fact, it holds that

$$\begin{aligned} v_{eq} &= \sum_i v_i \left(\frac{1/\sum_{k\neq i}\frac{1}{R_{Dk}}}{R_{Di}+1/\sum_{k\neq i}\frac{1}{R_{Dk}}}\right) \\ &= \sum_i v_i \left(\frac{1}{1+R_{Di}\sum_{k\neq i}\frac{1}{R_{Dk}}}\right) \\ &= \sum_i v_i \left(\frac{1}{1+R_{Di}\left(\sum_k \frac{1}{R_{Dk}}-\frac{1}{R_{Di}}\right)}\right) \end{aligned} \tag{2.5}$$

from where, recalling Eq. 2.4 one gets

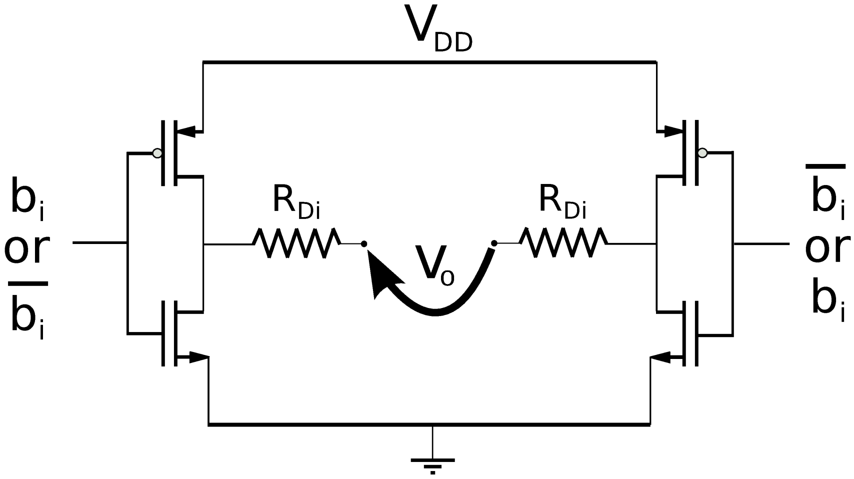

Figure 2.3: General scheme of a voltage-mode transmitter.

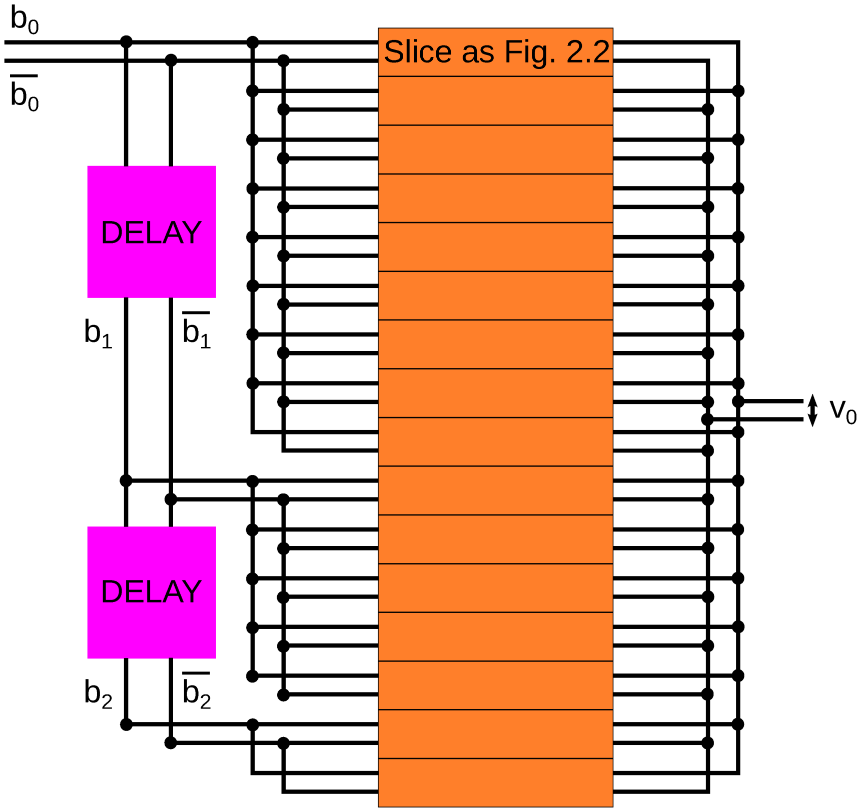

Figure 2.4: Example of FFE implementation using slices (with schematics as in Fig. 2.3) put in parallel.

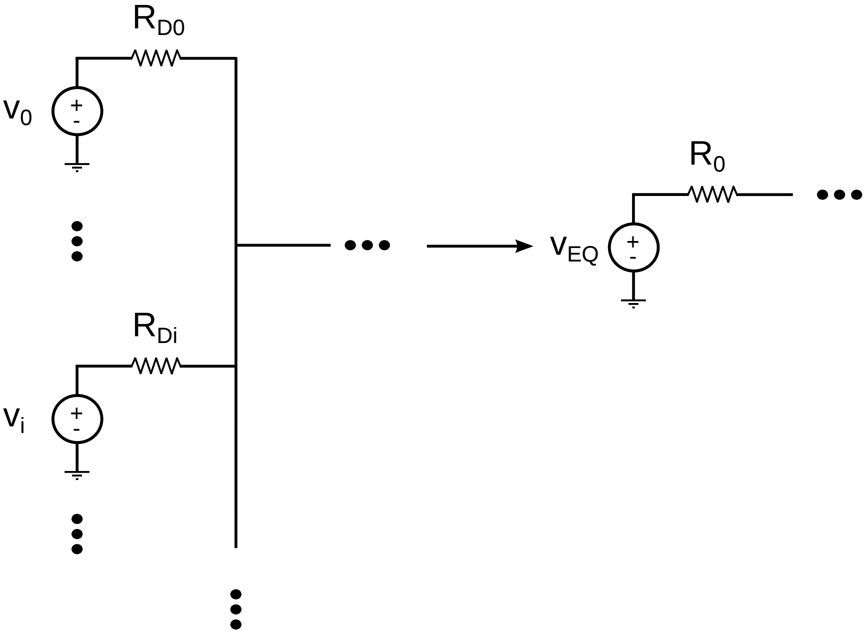

Figure 2.5: Thevenin equivalent for a half driver.

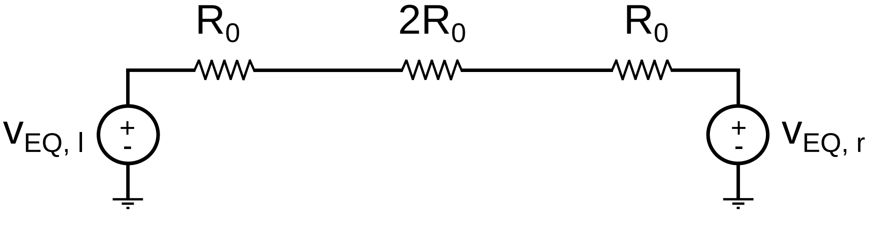

Figure 2.6: Thevenin equivalent for the whole driver with differential termination.

$$v_{eq} = \sum_i v_i \frac{R_0}{R_{Di}} \tag{2.6}$$

Considering that v_i can only be V_{DD} or 0 depending on the value of b_i (1 or 0) and on the connection with the driver, Eq. 2.6 can be rewritten as

$$v_{eq} = \sum_i \left(\frac{V_{DD}}{2} \pm \frac{V_{DD}}{2} b_i \right) \frac{R_0}{R_{Di}} \tag{2.7}$$

which holds true for both $v_{eq,l}$ and $v_{eq,r}$ just by swapping the ± sign in front of the b_i term, since when $v_{eq,l}$ is 1 $v_{eq,r}$ is 0 and vice-versa. Finally, one can write a close expression for v_0 that reads

$$\begin{aligned} v_o &= \frac{v_{eq,l} - v_{eq,r}}{2} \\ &= \sum_i \left(\pm \frac{V_{DD}}{2} b_i \right) \frac{R_0}{R_{Di}} \\ &= \frac{V_{DD}}{2} \cdot \sum_i \left(\frac{n_i}{M} \right) \cdot b_i \cdot sgn_i \end{aligned} \tag{2.8}$$

where (see Fig. 2.4) n_i is the number of slides connected to the bit b_i, M is the number of total slices and sgn_i is the sign of b_i (i.e. putting b_i at the left or right part of Fig. 2.3 making the driver inverting or non-inverting). For the b_i, '1' corresponds to b_i =1 and '0' means b_i =-1.

Eq. 2.8 can be rewritten as

$$v_o\left[j\right] = \frac{V_{DD}}{2} \sum_i \left(w_i \cdot b_i\left[j\right]\right) = \frac{V_{DD}}{2} \sum_i \left(w_i \cdot data\left[j-i\right]\right) \tag{2.9}$$

where j indicates a bit period and w_i is the strength of the slices connected to the i-th bit normalized to the full driver strength. From now on, we will refer to w_i as to the "weight" for the i-th tap. From Eq. 2.9 we see that the structure realizes an FIR filter, which implements a convolution between the bit stream ($data\left[j\right]$) and the tap vector w_i.

It should be noted that in order to obtain fine impedance tuning and compensate PVT (Process, Voltage and Temperature) variations, one does not strictly follow Eq. 2.4: the driver (divided in slices) is sized to obtain a resistance much larger than 50Ω; then many replicas of the structure are duplicated and the number of such replicas put in parallel is adjusted to match the 50Ω target [42].

2.3 Choice of the Equalization Taps

In this paragraph, we will cover the basic steps that are required to determine the weights w_i for FFE. The approach that we follow in our equalization process is the zero-forcing method [43]. Here the discussion will be at a tutorial level to bridge the gap between the theory in [43] and the actual hardware implementation (Figs. 2.3, 2.4). The goal is to minimize the distance between the desired response of the transmitter+channel (i.e. a signal without ISI) and that actually received; this is done via a Least Squares Minimization problem that reads

$$\min_{w_{ZFE}} ||z_{DES} - H_{CH}w_{ZFE}||^2 \tag{2.10}$$

where w_{ZFE} is the weights vector $[w_0, w_1, \ldots, w_i, \ldots w_N]$ to be determined by the minimization problem and z_{DES} is the desired output response. In other words, if we consider a bit stream '10000', the transmitter will generate a sequence of pulses of height $[w_0, w_1, \ldots, w_i, \ldots w_N]$, each one stimulating the channel. We want to set w_i in such a way that the receiver samples the original sequence '10000' without ISI, meaning that z_{DES} = [1,0,0,. . . ,0.]. H_{CH} is a matrix that rearranges the channel pulse response (h) in order to transform the convolution with the different pulses of height w_i into a matrix product. For example, if the channel pulse response is not null only for the first three samples and we decide to equalize with 5 post-cursor taps, then we have

$$H_{CH} = \begin{bmatrix} h_0 & 0 & 0 & 0 & 0 \\ h_1 & h_0 & 0 & 0 & 0 \\ h_2 & h_1 & h_0 & 0 & 0 \\ 0 & h_2 & h_1 & h_0 & 0 \\ 0 & 0 & h_2 & h_1 & h_0 \end{bmatrix} \tag{2.11}$$

If one wants to include also pre-cursor taps in the weights vector to be found, some slight changes have to be performed on both z_{DES} and H_{CH}. In fact, if one pre-cursor tap is to be considered, then z_{DES} becomes z_{DES} = [0,1,0,...,0], whereas H_{CH}, in a case where the channel pulse response is not null only for one pre-cursor tap and the first three post-cursor taps with one pre-cursor and five post-cursor taps of equalization, is

$$H_{CH} = \begin{bmatrix} h_{-1} & 0 & 0 & 0 & 0 & 0 \\ h_0 & h_{-1} & 0 & 0 & 0 & 0 \\ h_1 & h_0 & h_{-1} & 0 & 0 & 0 \\ h_2 & h_1 & h_0 & h_{-1} & 0 & 0 \\ 0 & h_2 & h_1 & h_0 & h_{-1} & 0 \\ 0 & 0 & h_2 & h_1 & h_0 & h_{-1} \end{bmatrix} \tag{2.12}$$

This approach can be extended to the number of pre-cursor and post-cursor equalization taps needed. The solution of the minimization problem introduced by Eq. 2.10 requires extracting the channel pulse response. In order to do this, the simulation setup shown in Fig. 2.7 is used. This setup is implemented in Ansys Electronic Desktop [44] and consists of a differential link with a pulse generator (with ideally steep rise and fall times) and an S-parameter block. This block changes from system to system and represents all the elements that compose the system after the transmitter.

The solution of the Eq. 2.10 reads

$$w_{ZFE} = \left(H_{CH}^{T} \cdot H_{CH}\right)^{-1} \cdot H_{CH}^{T} \cdot z_{DES} \tag{2.13}$$

which is the well-known solution for the LS objective function of Eq. 2.10. The effect of a correct choice of equalization tap weights is sketched in Fig. 2.8.

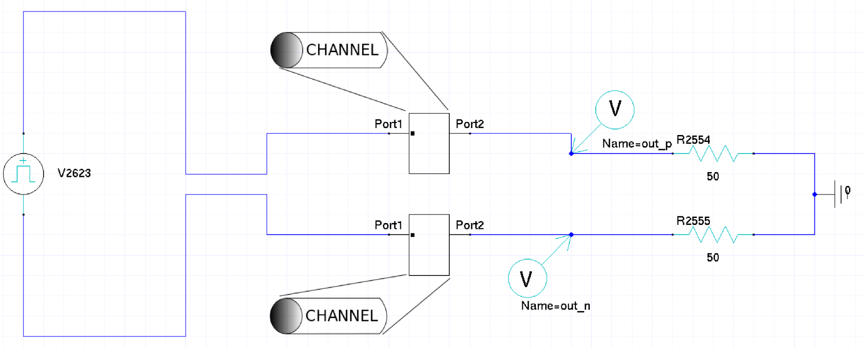

Figure 2.7: Setup implemented in Ansys Electronic Desktop in order to extract the channel pulse response.

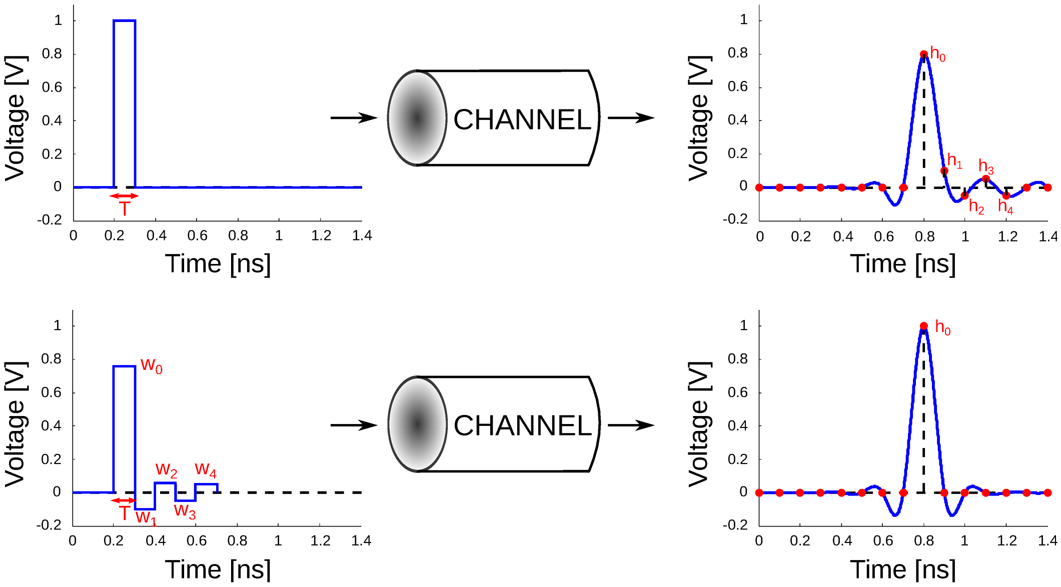

Figure 2.8: Pulse response when no feed-forward equalization is applied (top) and when ideal weights w_i are applied (bottom).

2.4 Example with Realistic Channels

In this section, as an example, we analyze two different systems, which we will refer as “BGA system” and “Leadframe system”. The Ball Grid Array (BGA) system is composed by a via, approximately 5 mm long, which connects the transmitter output signal to the package output, a BGA package, a Printed Circuit Board (PCB) and a cable (10 cm long). At the receiver end, the impedance is matched at 100Ω differential and the differential output is measured via voltage probes. The Leadframe System is similar to the BGA one, but instead of a BGA package it uses a leadframe one. Figs. 2.9 and 2.10 show the S_{21} of both systems obtained via quasi-3D electromagnetic simulations with Ansys SIWave [45]: Fig. 2.11 shows how a BGA package looks like when analyzed with Ansys SIWave. As stated in the previous chapter, looking at S_{21} is a good way to understand which interference mechanism is the dominant one in the transmitting medium.

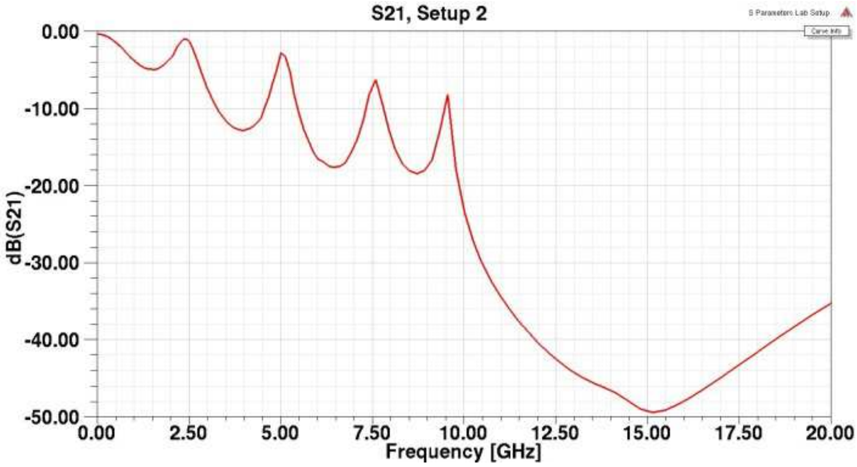

Figure 2.9: S_{21} of the Leadframe system.

Once the optimal weights have been found with Eq. 2.13, another verification about the correctness of the w_i is performed by simulating a structure equivalent to the system composed by transmitter, package, board and channel and evaluating the improvements obtained in the eye diagram. The software used to this end is again Ansys Electronic Desktop [44]. Fig. 2.12 illustrates a typical simulation setup: the only difference with respect to the structure presented in Fig. 2.7 is the use of a PRBS (Pseudo Random Bit Sequence) generator instead of a pulse generator. The PRBS generator offers the possibility to adapt its output based on equalization weights inserted by the user: in this case the vector w_{ZFE} obtained with Eq. 2.13. The low and the high voltage levels have been set to -450mV and 450 mV respectively, in order to analyze a mimic a possible scenario with a V_{DD} power domain of 900mV that directly supplies the driver slices. Thermal noise typical of MOSFETs in the transmitter has not been modeled and the channel is terminated with a 50 Ohm resistance.

One can evaluate how close the overall transmitter+channel response fits the wanted z_{DES} by checking it against the product $H_{CH} \cdot w_{ZFE}$ (i.e. the overall response of the FFE+channel). This is shown in Fig. 2.13, which compares the

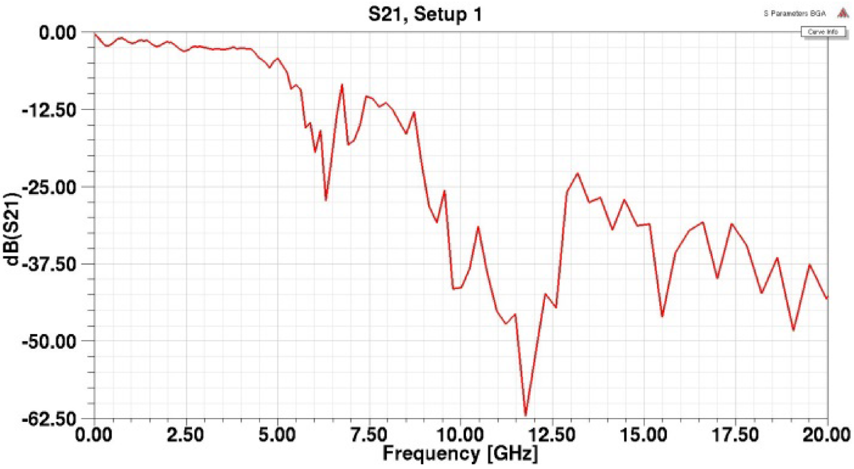

Figure 2.10: S_{21} of the BGA system.

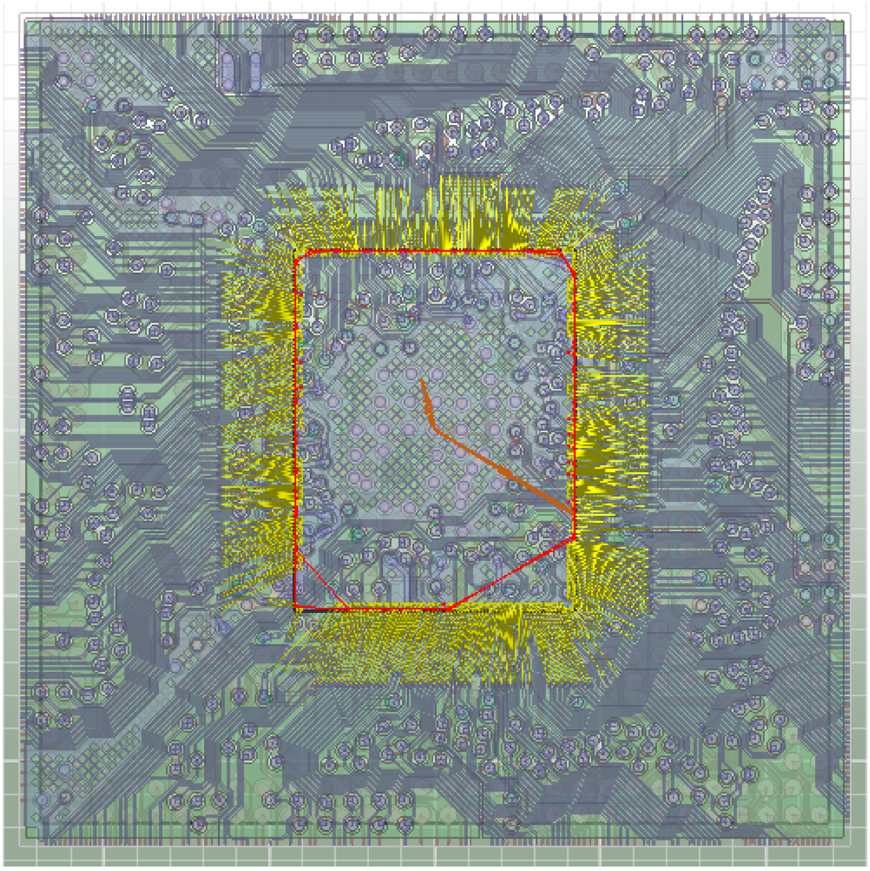

Figure 2.11: Overview of a BGA516 package as imported in Ansys SIWave [45].

effect that equalizations with different number of taps has on the overall pulse response of the channel of Fig. 2.10. The weights used to obtain the curves in Fig. 2.13 are reported in tab. 2.1 In this particular case (consistent with the description above), only post-cursor taps were used, therefore the desired impulse channel response would be a 1 followed by a number of zeros equal to the post-cursor taps used. So, for a 4-taps equalization (main and three post-cursors) the desired channel pulse response would be z_{DES} = [1,0,0,0]. As said before, if pre-cursor taps are used, then the z_{DES} elements are shifted to the right by a number of places equal to the number of pre-cursor taps.

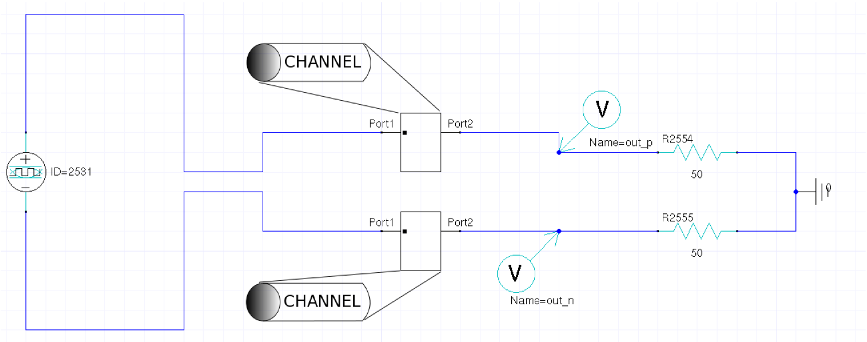

Figure 2.12: Setup implemented in Ansys Electronic Desktop in order to extract the obtain the eye diagram of the system. The only difference with respect to the structure presented in Fig. 2.7 is the use of a PRBS (Pseudo Random Bit Sequence) generator instead of a pulse generator.

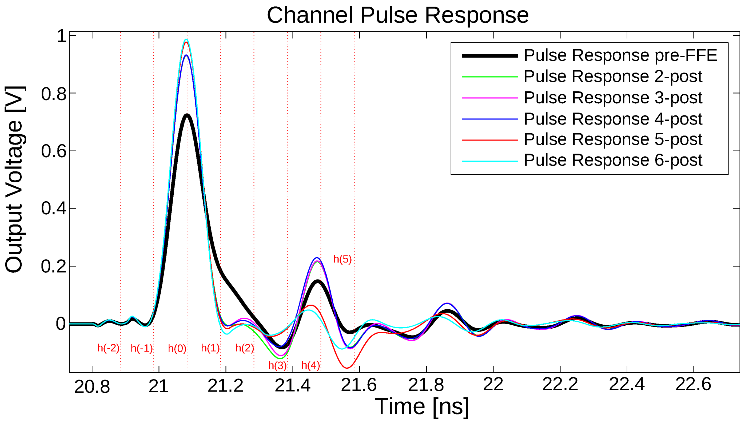

Figure 2.13: Response of channel of Fig. 2.9 operating at 10Gbps along with responses after equalization for various numbers of taps, obtained with the simulation setup of Fig. 2.7.

In order to confirm the validity of our simplified approach, in Figs. 2.14 and 2.15 two eye diagrams of the same system (at 2.5 Gbps) are shown. One is obtained with a PRBS generator and the other one with a transistor level model of the transmitter presented in [46], and their eye diagram parameters, eye width, eye height and Signal-to-Noise Ratio, are very similar, provided

Table 2.1: Tap weights for FFE for the curves reported in Fig. 2.13.

	2-post	3-post	4-post	5-post	6-post
w(0)	0.8158	0.7885	0.7714	0.6374	0.577
w(1)	-0.1842	-0.1849	-0.1795	-0.1545	-0.1526
w(2)		0.0266	0.0176	0.0101	0.0148
w(3)			0.0314	0.0622	0.0579
w(4)				-0.1359	-0.1392
w(5)					0.0585

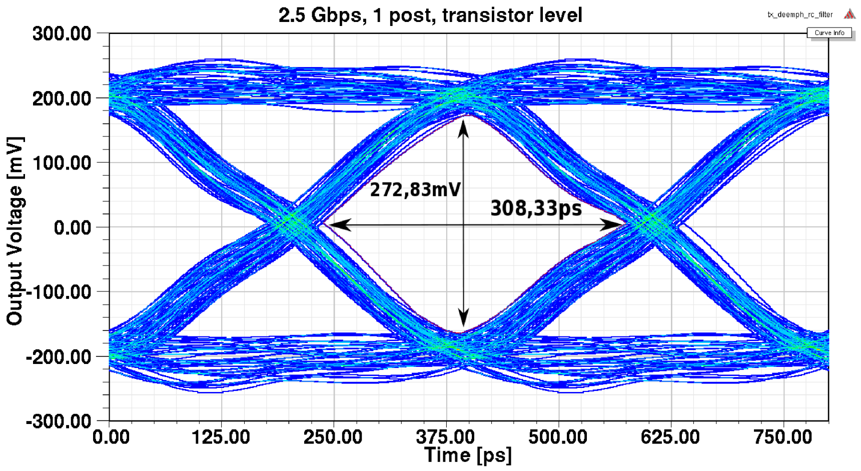

Figure 2.14: Eye diagram at 2.5 Gbps for the BGA System obtained with a transistor level simulation and one tap de-emphasis. The transmitter is described in [46].

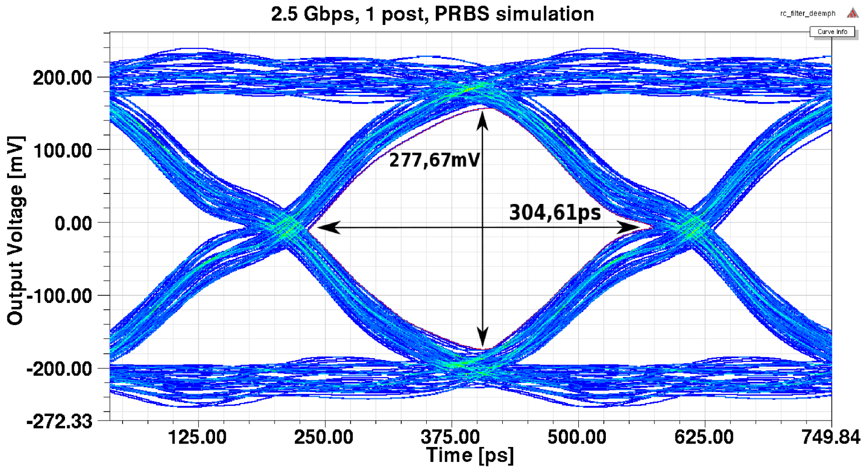

Figure 2.15: Eye diagram at 2.5 Gbps for the BGA System obtained with a PRBS generator and one tap de-emphasis.

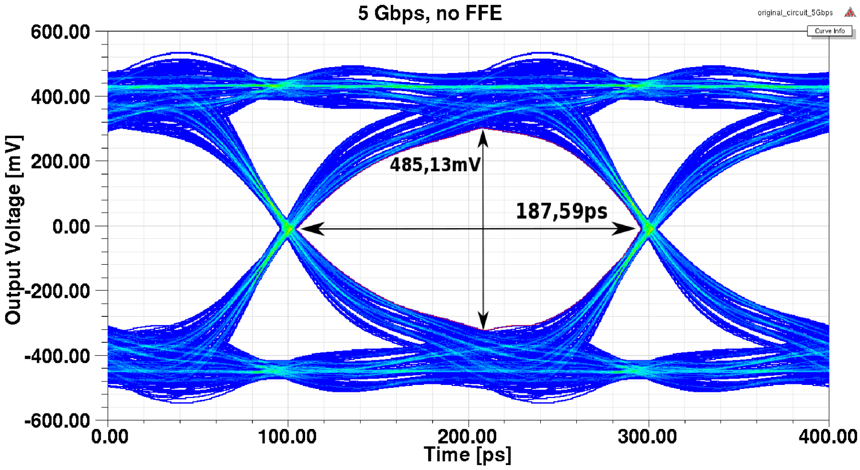

Figure 2.16: Eye diagram at 5 Gbps for the Leadframe System without Feed-Forward Equalization.

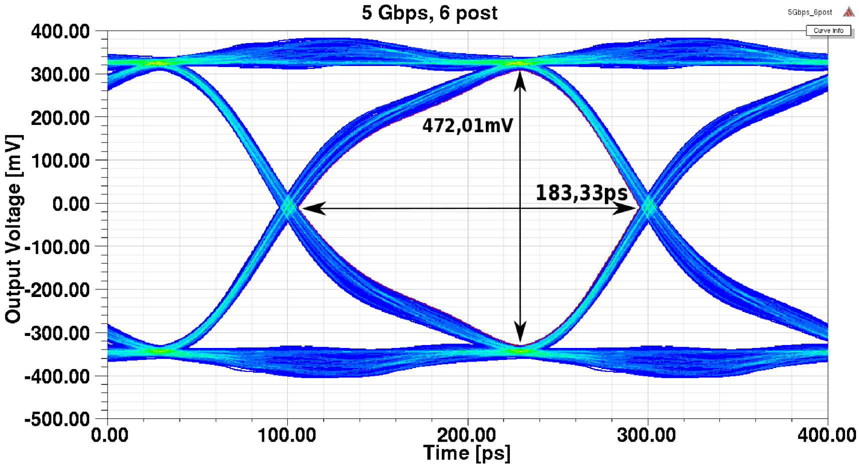

Figure 2.17: Eye diagram at 5 Gbps for the Leadframe System with six post-cursor taps (main tap plus five post-cursors). The weights that generate this eye are w_0=0.8009, w_1=0.0403, w_2=-0.1018, w_3=0.0395, w_4=-0.0045, w_5=-0.013.

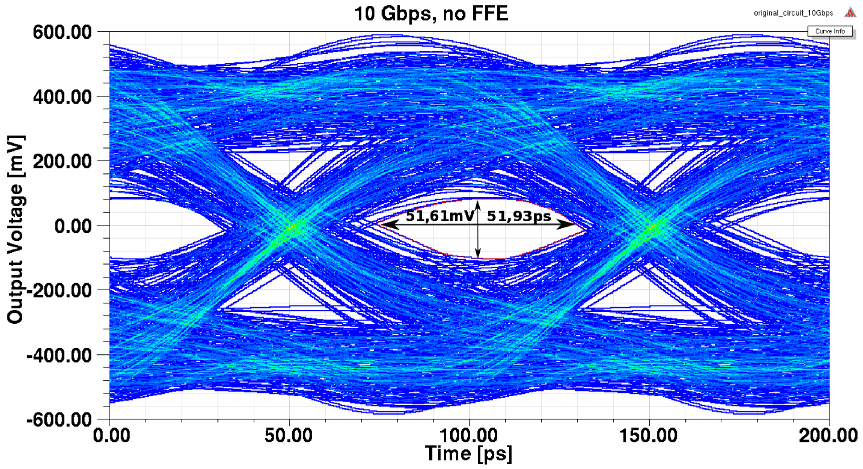

Figure 2.18: 15 Eye diagram at 10 Gbps for the Leadframe System without Feed-Forward Equalization.

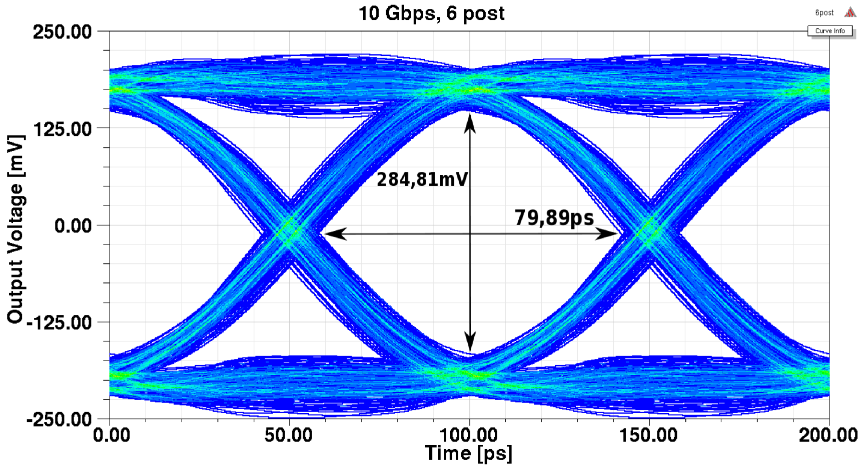

Figure 2.19: Eye diagram at 10 Gbps for the Leadframe System with six post-cursor taps (main tap plus five post-cursors). The weights that generate this eye are w_0=0.577, w_1=-0.1526, w_2=0.0148, w_3= 0.0579, w_4=-0.1392 and w_5=0.0585.

that the PRBS high and low levels are set to match the ones of the transistor level implementation. The eye width is the width of the horizontal histogram across the eye-crossing point, the eye height is the difference between high and low levels at the sampling time in which this difference is maximum and the Signal-to-Noise-Ratio is defined as difference between the average ’1’ and ’0’ levels divided by the sum of the standard deviations of the two levels at the center of the eye [44].

Since the system level design is performed before actually designing the transmitter at transistor level, in the following, for the 5 Gbps and the 10 Gbps cases only the PRBS source is used: we can follow this procedure since we have already proven the correctness of such a simulation approach In Figs. 2.14-2.15. Figs. 2.16, 2.17, 2.18 and 2.19 report the eye diagrams in such cases with and without equalization. For the 5 Gbps case, the improvement due to FFE is marginal, whereas to work at 10 Gbps with the Leadframe System, FFE is mandatory. Note that, at given V_{DD}, the inclusion of FFE lowers the high and low levels of the eye (e.g. 400mV vs. 300mV in Figs. 2.14 and 2.15): this is because when FFE is implemented, some slices will be driven by bits having opposite sign with respect to the main one, and this implies that the driver is not working at full strength. More precisely, the high and low levels are now shifted at $\pm \frac{V_{DD}}{2} \sum_i w_i$

2.5 Effect of Tap Quantization

Eq. 2.13 provides optimum tap weights, but one must also think at a real world implementation, which obviously implies quantizing these weights since each bit will be connected to a finite number of slices. This problem is peculiar to the voltage-mode transmitter divided in slices. In fact previous works already introduced equalization implemented with sliced drivers, but mainly in current mode logic [47]- [48], which makes equalization easier to implement with high granularity. In fact, as shown also in Fig. 2.20, the granularity of the weights is solely determined by the number of bits of the DAC controlling the current sources I_i of the taps.

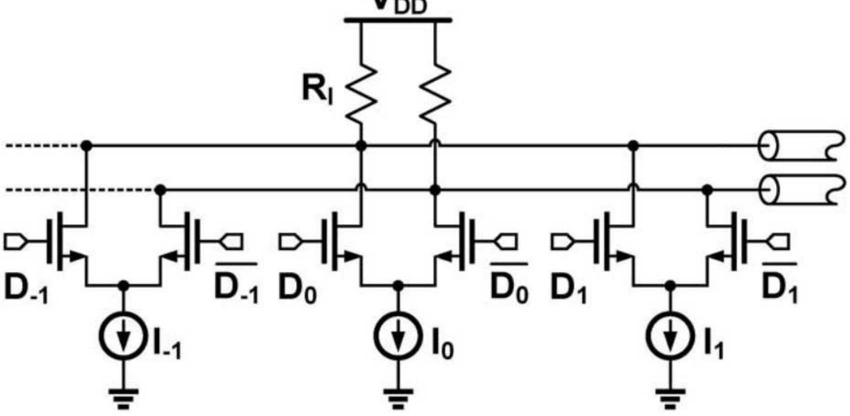

Figure 2.20: Architecture of a current mode transmitter with 2-taps FFE (one pre-cursor and one post-cursor) as implemented in [49].

Here we analyze the effect of quantization with two different granularities, 8 and 16 levels (i.e. M=8 or M=16 as in Figure 2.4). These two different granularities offer quantization steps of 0.125 and 0.0625 respectively. Figs. 2.21 and 2.22 show the effect of quantization on the operation at 10 Gbps of Leadframe System when equalized with 6 post-cursor, which without quantization has already been shown in Figure 2.19. With 16 slices we obtain eye parameters very close to what is obtained from Eq. 2.13, whereas with 8 slices there is a degradation of the eye.

If the transmission speed gets even higher, then a higher number of taps is needed and the effect of quantization becomes more and more relevant. The eye parameters with and without FFE and including different granularity in the tap quantization are summarized in Figs. 2.23-2.26. We include in these figures also a hypothetical 15 Gbps situation that requires 9 taps for equalization when considering the Leadframe System (7 for the BGA System). Obviously, at this high speed, the main wall to climb would be the transistor level design of some critical blocks in se, as we will see in the next chapter. In Figs. 2.23-2.26 we see that FFE improves the eye height and width, although part of the improvement is lost if a too coarse granularity is used for the w_i.

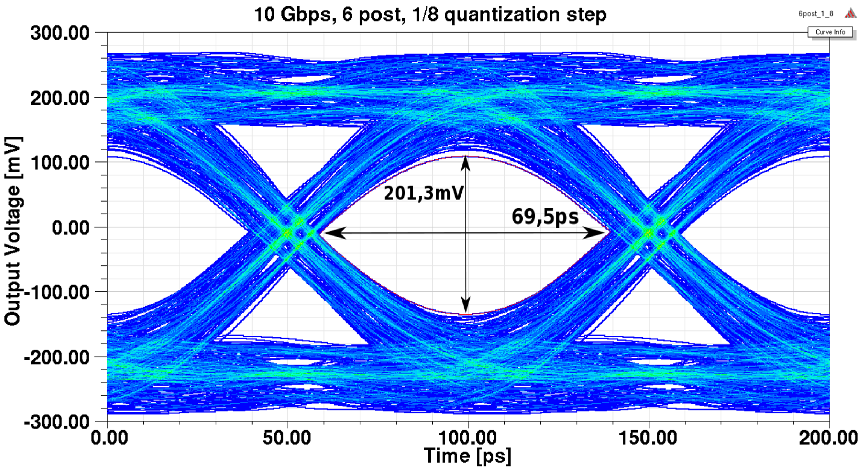

Figure 2.21: Eye diagram at 10 Gbps for the Leadframe System with 6-post cursor taps and quantization step of $\frac{1}{8}$. The weights are w_0=0.625, w_1=-0.125, w_2=0, w_3=0, w_4=-0.125 and w_5=0.125, which correspond to 5 slices connected to b_0 while b_1, b_4 and b_5 require one slice each.

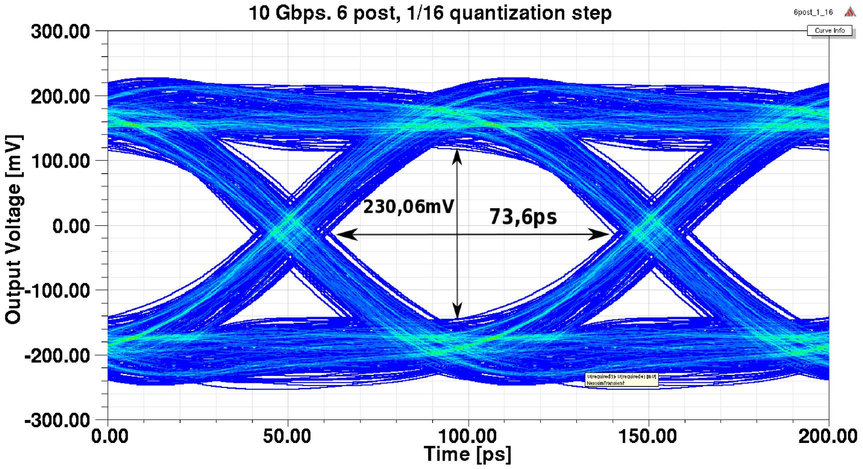

Figure 2.22: Eye diagram at 10 Gbps for the Leadframe System with 6-post cursor taps and quantization step of $\frac{1}{16}$. The weights are w_0=0.5625, w_1=-0.1875, w_2=0, w_3=0.0625, w_4=-0.125 and w_5=0.0625, which correspond to 9 slices connected to b_0, three connected to b_1, one slice connected to b_3, one to b_5 and two to b_4.

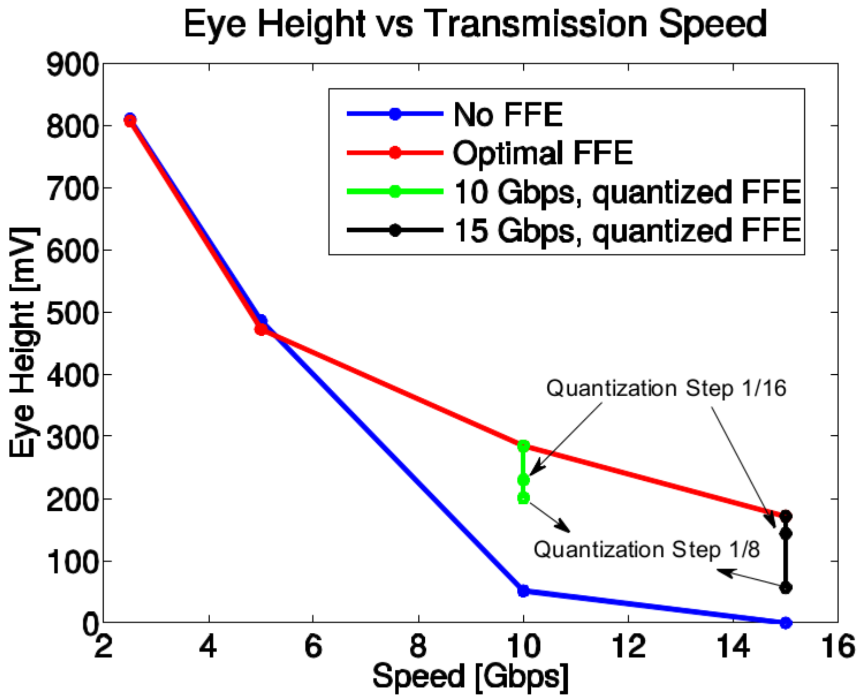

Figure 2.23: Eye height versus transmission speed for the Leadframe System when not equalized, optimally equalized (Eq. 2.13) and when quantization (8 or 16 steps) is applied to w_i.

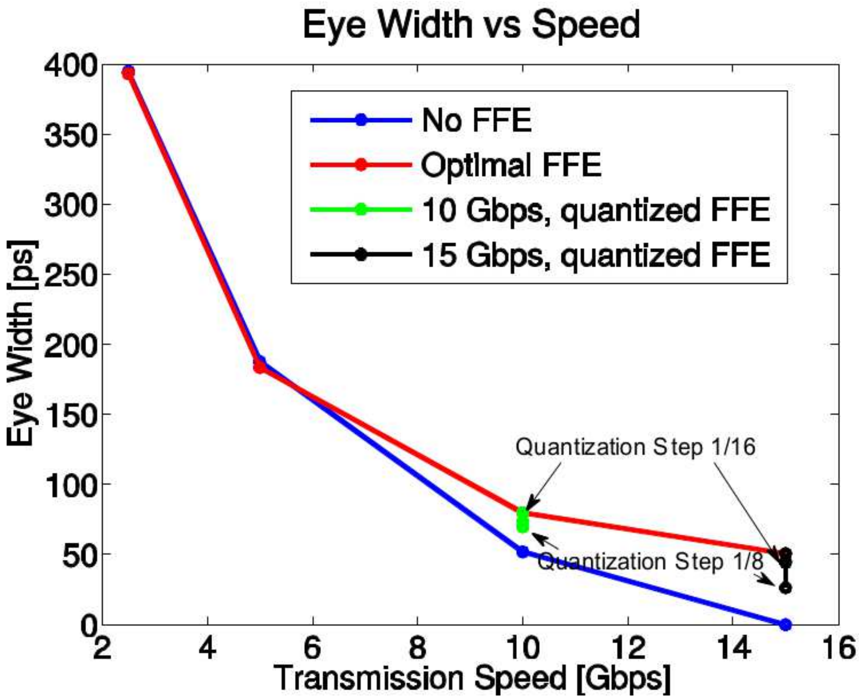

Figure 2.24: Same as for Fig. 2.23, but for the eye width.

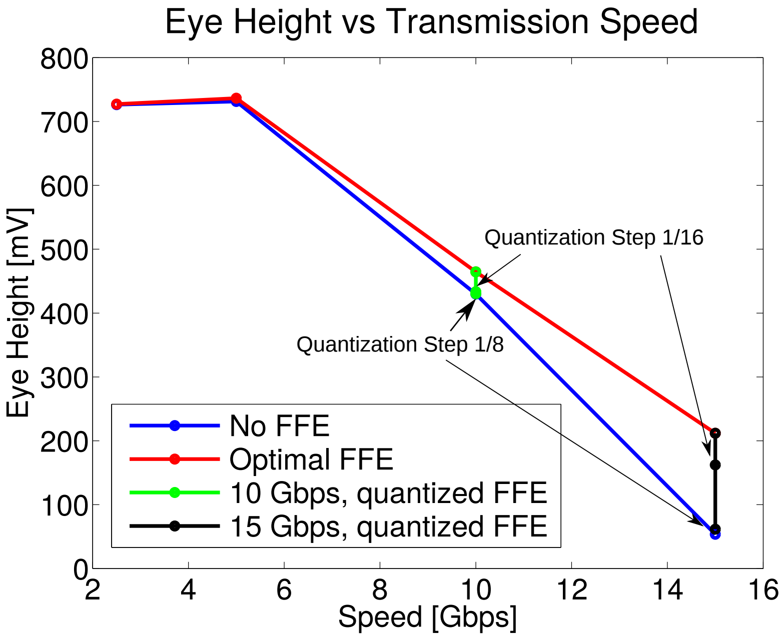

Figure 2.25: Same as for Fig. 2.23, but for the BGA system.

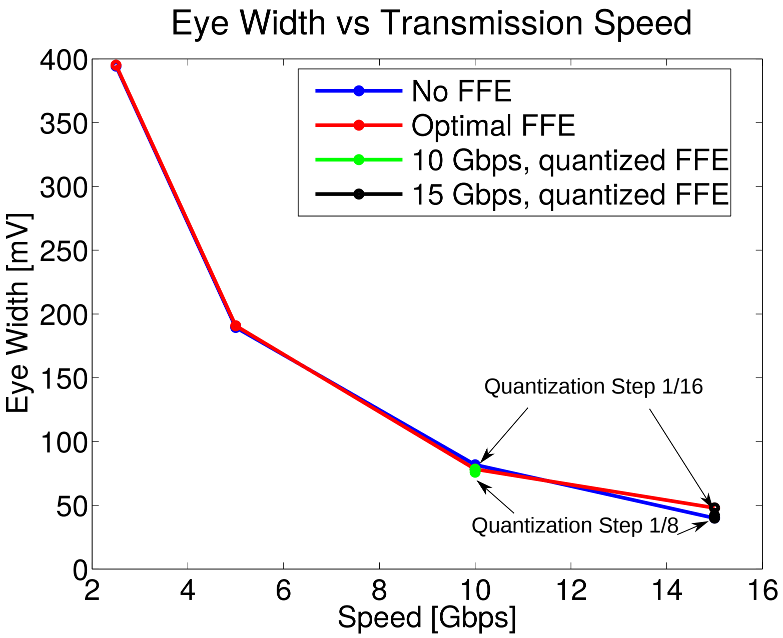

Figure 2.26: Same as for Fig. 2.25, but for the eye width.

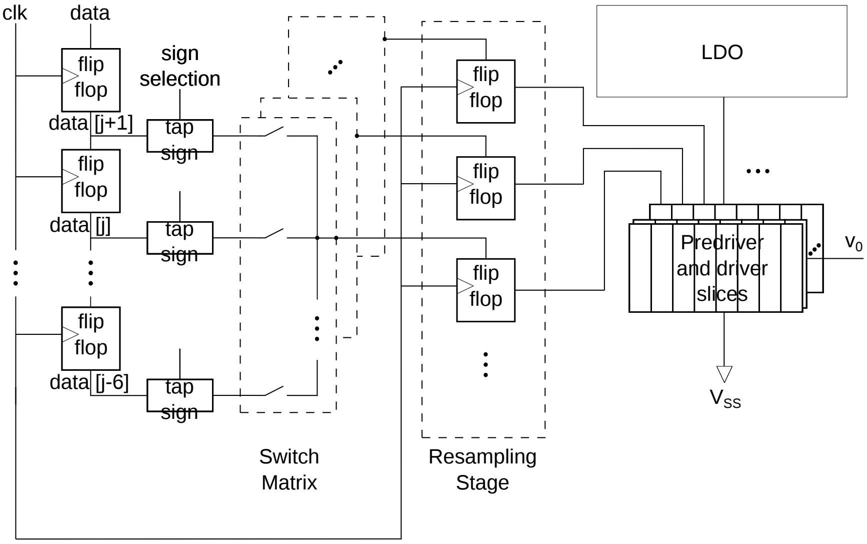

Figure 2.27: General architecture of the transmitter. Despite the fact that single data and clock line are shown in this picture, the architecture is fully differential and requires positive and negative versions of the clock and data signals. $data[j+1]$ to $data[j-6]$ are made available by the shift-register on the left side; the switch matrix is composed by 8x8 switches. We do not show explicitly the 8 slices of the driver that are hard-wired to the main tap but only the 8 that are connected to the switch matrix. The vector of 16 differential slices is duplicated K times to allow impedance tuning.

2.6 Architecture of the Transmitter

After the system level analysis performed in the previous sections, we can finally draw the architecture of the transmitter. The single blocks will be analyzed in detail in Chapter 3. The structure of the transmitter is sketched in Fig. 2.27. On the left, we see the shift-register that makes available to the driver (right side of the picture) the delayed version of the serial data (*data*) to be transmitted. We decided to have one pre-cursor tap because such a correction will not be possible with DFE at the receiver and one more post-cursor tap with respect to the analysis carried out in the previous chapter to prevent possible effects deriving from PVT corners not foreseen by such analysis. The main driver is composed by M=16 slices (whose structure is sketched in Fig. 2.28) and the switch matrix sketched at the center of Fig. 2.27 determines how many slices n_i are connected to the i-th delayed version of $data$. A "sign" (sgn_i) is associated with each tap. Although not shown in Fig. 2.27, 8 slices over 16 are hard-wired to the main tap (i=0) with positive sign. Since the driver (Fig. 2.28) and all signals are differential, the selection of the sign is implemented by just swapping the two lines composing the signal. The re-sampling stage in Fig. 2.27 realigns the various signals at the output of the switch matrix reducing deterministic jitter.

All the slices are connected in parallel to the output signal. The block of 16 slices is replicated K times in order to allow impedance matching, by activating

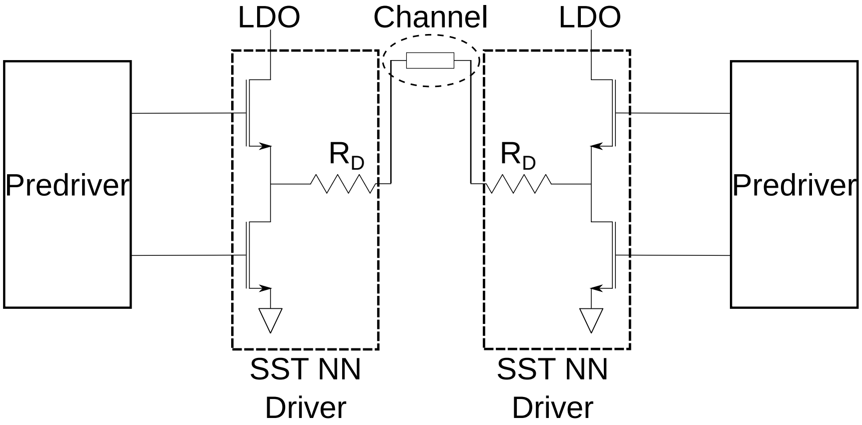

Figure 2.28: Scheme of the predriver and driver slices on the right of Fig. 2.27. The predriver slice is composed by two inverters with progressive widths. Each driver slice is supplied by the LDO output voltage (nominally,400 mV). The channel features a 100 Ω differential resistance.

only k replicas. The resulting output impedance is thus $R_{out} = R_D/(M \cdot k)$ regardless of the transmitted signal and of the choice of the tap weights, where (see Fig.2.28) R_D is the series impedance of a single slice.

Since we wanted to consume as less power as possible, we wanted to transmit signals with low swing, so to minimize the power consumption over the termination resistance. In order to do so, we have to reduce the supply that feeds the driver slices via an LDO, but this implies that it's almost impossible to use a conventional inverter structure as shown in Fig. 2.3, since there is no voltage room to switch on the pMOSFET. For this reason, we have decided to implement an NN slice as shown in Fig.2.28. We have selected the single resistor implementation for the slices because it guarantees lower parasitic caps at the output of the slice [29]. Obviously also the series resistance of the pull-up/pull-down devices are part of the impedance matching network, but having sufficiently wide transistors guarantees that their resistance is small compared to the big values of R_D. We set K=16 that is extremely helpful in solving impedance mismatches originating from spread over technology corners. Once k is properly selected to achieve $R_{out} = Z_{channel}$, we have to rework Eq. 2.9 since the slices are not directly supplied via V_{DD} anymore, so that the differential output voltage at clock period j is now given by:

$$v_0[j] = \frac{V_{LDO}}{2} \sum_{i=-1}^{6} w_i \cdot data[j-i] \tag{2.14}$$

where V_{LDO} is the output voltage of the LDO. Eq. 2.14 shows that the transmitter with FFE behaves as a FIR filter with up to 8 taps (one pre-cursor, one main and six post-cursors) with weights w_i programmable in M=16 steps. It thus generates 16 levels going from $-V_{LDO}/2$ to $V_{LDO}/2$.

A full-rate architecture has been chosen instead of half-rate [50]. The latter one requires either two replicas of the driver [29], [51] or interleave two sequences at the input of a single driver [52].

Fig. 2.29a shows an octave-transmitter without FFE where the driver has been replicated eight times and each data stream at frequency $\frac{f_{DATA}}{8}$ drives one of the replicas [51]. This is clearly an extremization of the first option, focused on power saving rather then on having an architecture feasible for a complex FFE scheme. Fig. 2.29b presents the architecture of a half-rate transmitter with one pre-cursor tap FFE where a 2:1 mux selects the data stream at frequency $\frac{f_{DATA}}{2}$ which drives the final predriver and driver stages [52]. This is the typical case for the second option where two half-rate data sequences are interleaved at the input of a single driver.

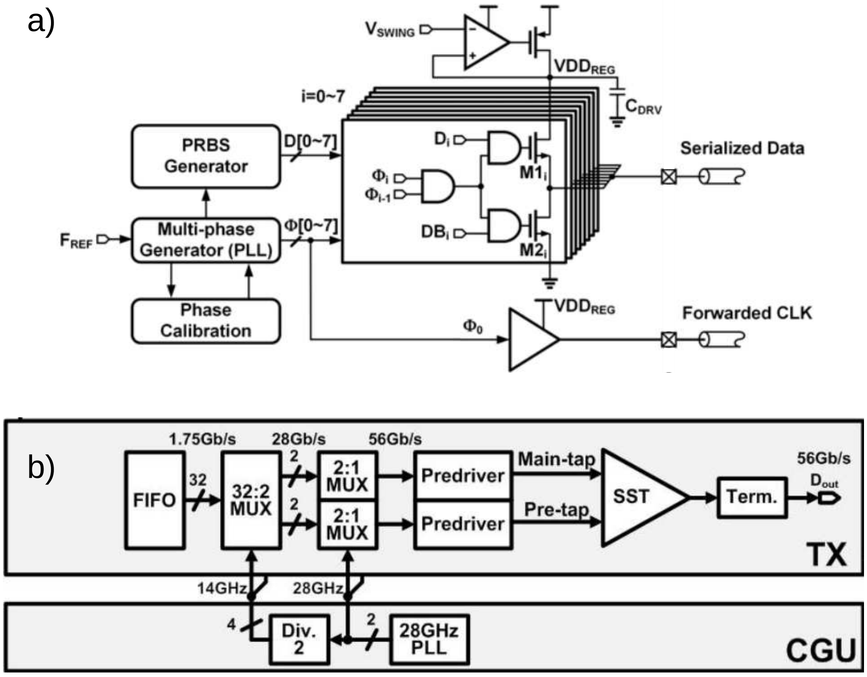

Figure 2.29: a) Architecture of an octave rate transmitter without FFE where the driver has been replicated eight times and each data stream at frequency $\frac{f_{DATA}}{8}$ drives one of the replicas [51] and b) architecture of a half-rate transmitter with one pre-cursor tap FFE where a 2:1 mux selects the data stream at frequency $\frac{f_{DATA}}{2}$ which drives the final predriver and driver stages [52].

In both cases, the flexibility in our choice of the weights w_i would make exceedingly complicate and inefficient the distributions of the delayed versions of the input data. On the other hand, in a full-rate implementation, the flip-flops must operate at 10 Gbps, which makes them more power hungry.

Chapter 3

Transistor Level Design of the Transmitter

3.1 Introduction

In the previous chapter, we investigate how to implement FFE over channels typical of automotive environment pointing out the effort of the tap number and of tap quantization. Moreover, system equalization will be even more challenging when dealing with automotive standards that require a broad range of temperature and supply voltage variations. Hence, suitable equalizers for automotive standards must be highly tunable systems that usually ask for more area and power. At the end of the previous chapter we have then introduced the architecture of the transmitter.

In this chapter we will describe the transistor level design of the transmitter including its Feed-Forward Equalization, following the analysis carried out in Chapter 2. Here design choices coming from literature study will be analyzed, justified and explained, along with the results deriving from schematic and post-layout simulations. The system has been designed at the transistor level and then layouted using a 28 nm planar technology. After this, a schematic analysis of the effect of parasitic inductance on FFE will be carried out.

3.2 Flip-Flops

As we have seen in Fig. 2.27, the main building blocks of the system are the flip-flops (needed for the shift-register and for the re-sampling stage), the driver slices, the LDO and the switch matrix.

The schematic of flip-flops with a pseudo-differential architecture is reported in Fig.3.1 and is a slight adaptation of [53], where the clock pass transistor has been substituted with a t-gate. The layout of the cell is reported in Fig. 3.2a, whereas the waveforms in Fig.3.2b demonstrate that such circuit is able to correctly operate with a 10 GHz clock with a V_{DD} as low as 750 mV (the nominal supply voltage for the flip-flops is 1 V, but due to IR and $L \cdot \frac{dI}{dt}$ drop and supply parasitics it may drop to values as low as 750 mV). In order to contain the effect of IR and parasitic inductance drop, two different supply domains have been used [29]: one for the LDO and the predriver (the driver is under the LDO), the other for all the rest.

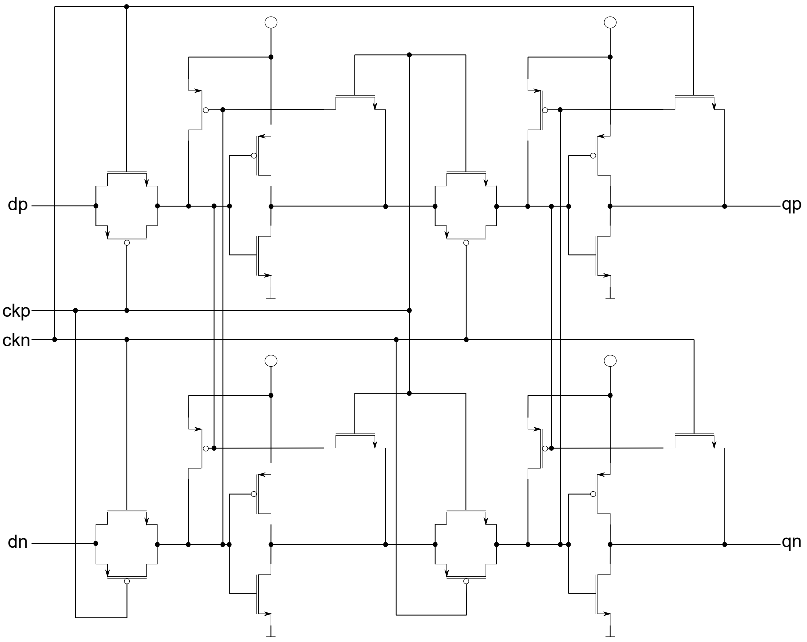

Figure 3.1: Schematic of the pseudo-differential flip-flop used in the transmitter.

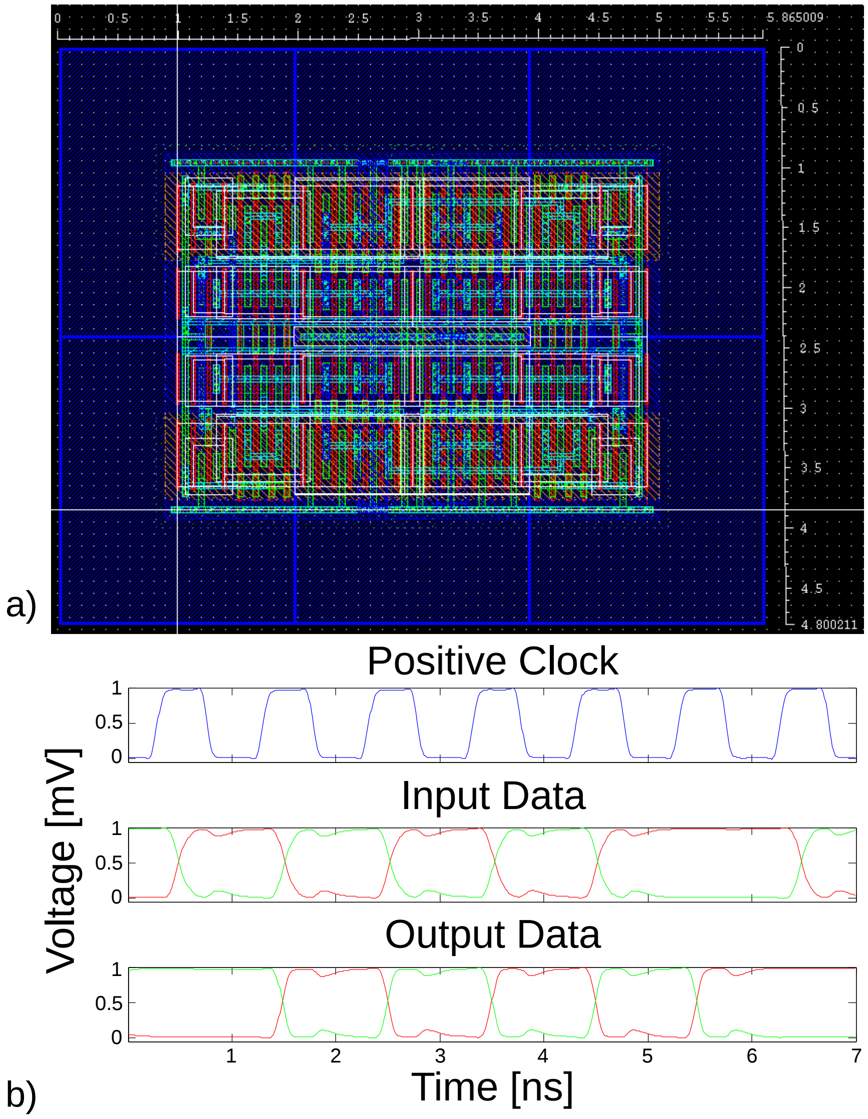

Figure 3.2: a) Layout of the pseudo-differential flip-flop of Fig. 3.1 and b) simulated post-layout transient operation for the same cell with a 10 GHz clock. In this picture, waveforms of positive clock signal and both positive and negative data for a supply of 950 mV are presented.

3.3 Switch Matrix

Another critical part is the switch matrix. The inputs and the outputs of the switch matrix require some buffering, which is implemented by the inverters *invA* and *invB* in Fig. 3.3. *invA* is needed to lower the capacitive load seen by the sign-selection circuit, whereas *invB* to lower the capacitive load seen by the switch matrix. Moreover, by design, the signal delay in the critical path going from the shift-register to the re-sampling stage needs to be lower than the clock period. The breakdown of the delay of the single stages is reported in Fig. 3.3 for the worst case scenario where the supply voltage has the lowest values due to IR drop (750 mV) and considering supply parasitics (0.8 nH series inductance and 0.8 Ω series resistance for both power supply domains and 0.2 nH and 0.2 Ω for ground connection). Fig. 3.4 shows the schematic of the core switches of the switch matrix for one of the eight selectable driver slices, along with inv_B. The core switch has been implemented as a tristate inverter for each possible input signal to each selectable slice, since only one signal among the eight available (pre-cursor, main and six post-cursor taps) will drive a slice, so seven tristate inverters out of eight will be in high-impedance. For this reason, the function of inv_B is particularly important to fasten the transition of the core switch.

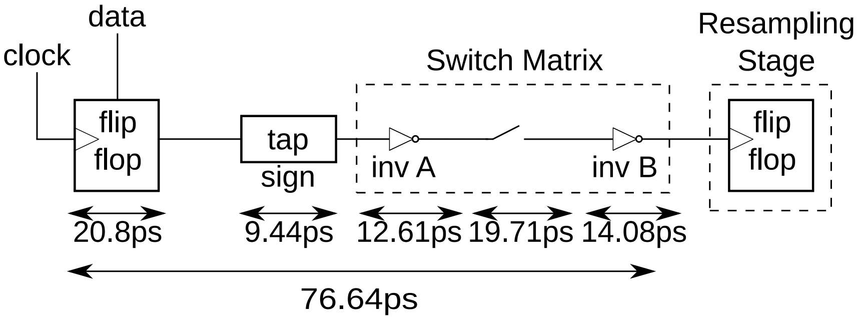

Figure 3.3: Critical data path between shift-register and re-sampling along with the worst-case delay of the principal blocks. With respect to Fig. 2.27, the switch matrix has been exploded into input buffer (*invA*), core switch and output buffer (*invB*).

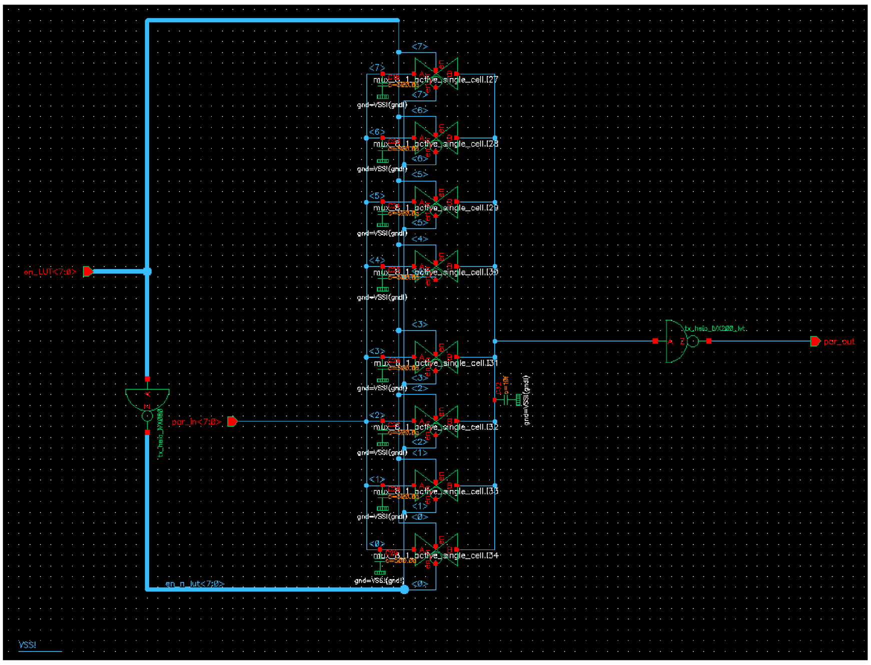
en_LUT<7:0>
par_in<7:0>
en_n_lut<7:0>
tx_hsio_IVX080
mux_8_1_active_single_cell.I27
mux_8_1_active_single_cell.I28
mux_8_1_active_single_cell.I29
mux_8_1_active_single_cell.I30
mux_8_1_active_single_cell.I31
mux_8_1_active_single_cell.I32
mux_8_1_active_single_cell.I33
mux_8_1_active_single_cell.I34
gnd=VSS!(gnd!)
tx_hsio_IVX200_lvt
par_out
VSS!

Figure 3.4: Schematic of the switch matrix. On the left we see the eight enable signals for the eight tristate inverters of which each core switch is composed. The active tristate inverter among the eight, at the center of the picture, has to load a very big output capacitance due to high-impedance connections of the other seven. In order not to add to this node also the input parasitic capacitance of the re-sampling stage, inv_B on the right side, has been inserted.

3.4 LDO

Concerning the LDO that supplies the driver slices, a folded cascode architecture has been selected, see Fig.3.5, since it is suitable for the low voltage supply environment. The pass device (highlighted in red in Fig. 3.5) occupies most of the layout area for LDO, being the total sum of all its fingers width 625μm, along with the eternal load capacitance of 15pF. The layout of the LDO is reported in Fig. 3.6.

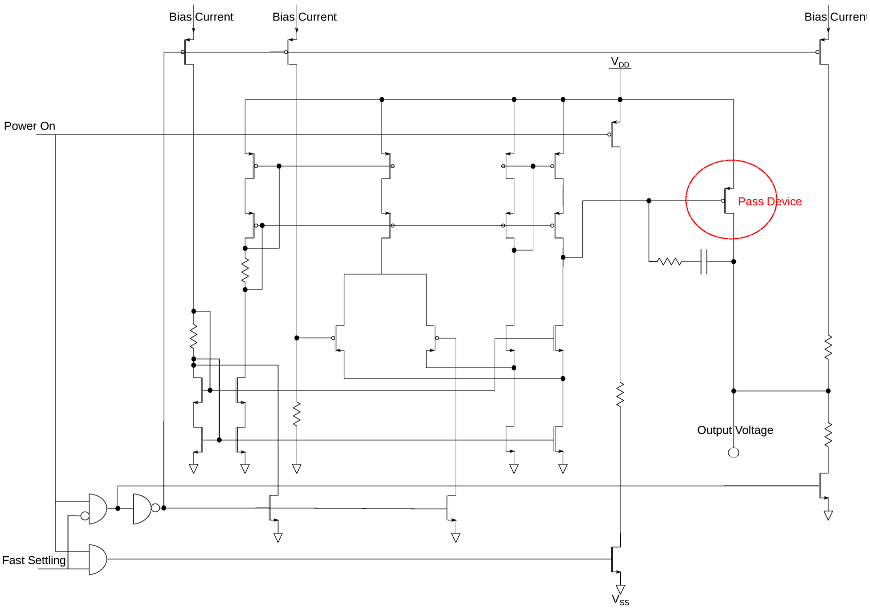

Figure 3.5: Schematic of the LDO.

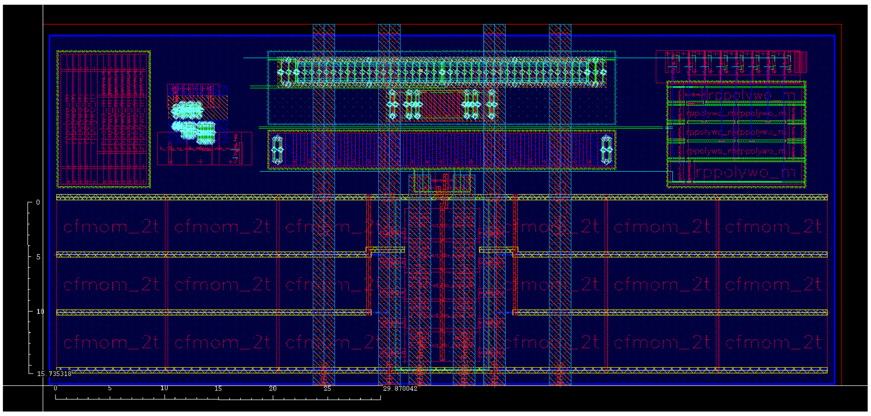

Figure 3.6: Layout of the LDO.

Fig. 3.7 illustrates the parameters we were more focused on when designing the LDO, Power Supply Rejection Ratio (PSRR) and loop gain. PSRR is

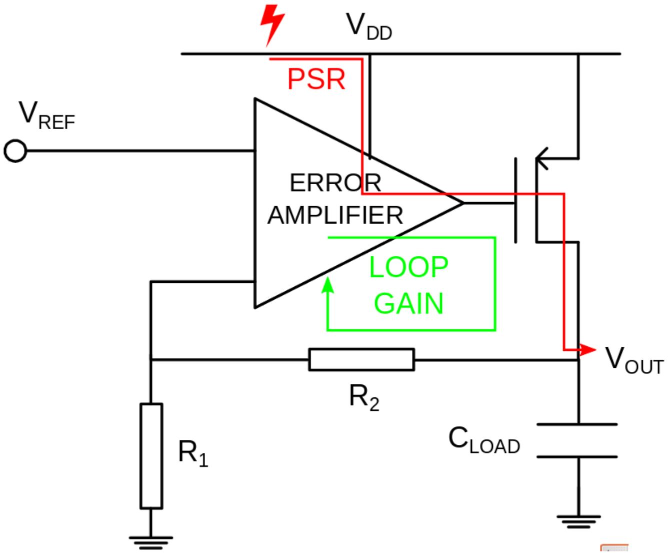

Figure 3.7: Main parameters of an LDO.

a widely known concept in electronics and it is the capability of a system to suppress any variation in the power supply to its output signal, whereas here when we talk about loop gain we refer to the LDO closed loop gain. Finally, a third parameter to take care of is the phase margin of the loop gain, also a widely known concept, in this case defined as the difference between the phase of zero dB loop gain and 180°. In nominal corner (V_{DD}=900mV, T=25° C, top-top technology corner), the LDO features a 45.85 dB loop gain, -38.83 dB of PSR over a bandwidth of 20 MHz and a phase margin of 94 degree. The curves for loop gain, power supply rejection (PSR) and phase margin for PT corners (-40°C, 170°C, slow-slow, fast-fast) for V_{DD} from 1V (best case supply voltage) down to 600mV (even if the worst case supply voltage is 800mV, just to check at which supply voltage the LDO starts to fail) are reported in Figs. 3.8-3.10.

Fig. 3.11 shows that the settling time of the LDO is $\approx 50ns$, which is comparable to the state-of-the-art power-on time for rapid on/off links used for burst-mode communication [54].

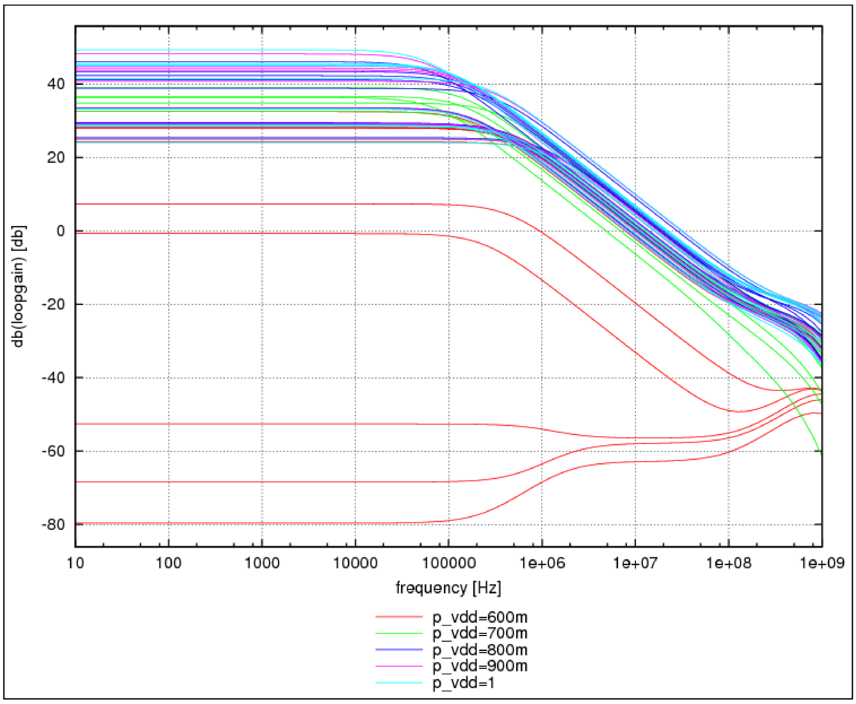

Figure 3.8: Loop-gain of the LDO for all PT corners (-40°C, 170°C, slow-slow, fast-fast) and for V_{DD} from 1V (best case supply voltage) down to 600mV. All failing cases are with V_{DD}=600mV, which is way lower than our worst-case V_{DD} even after having considered IR-drops.

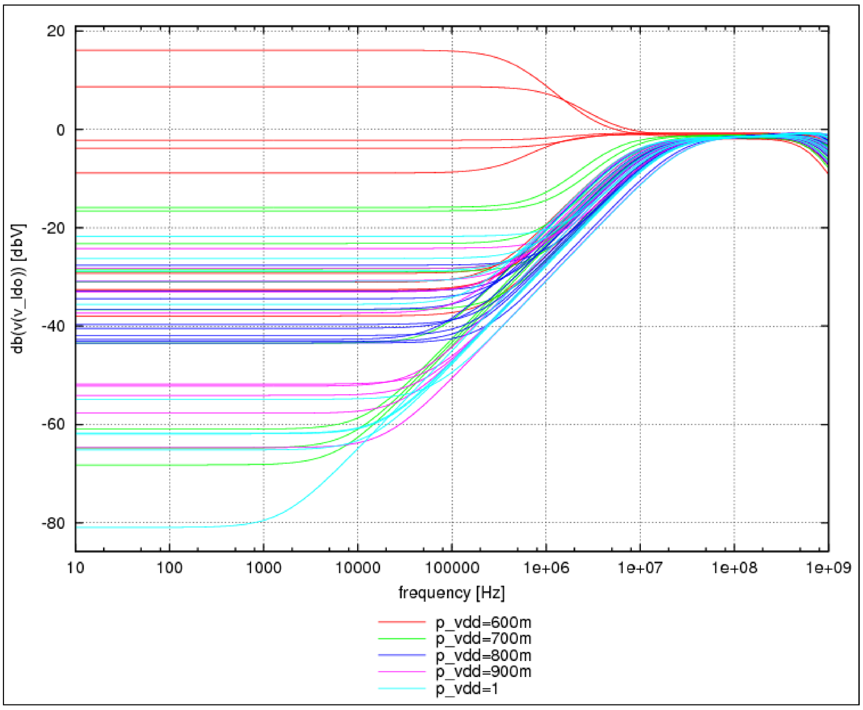

Figure 3.9: Power Supply Rejection (PSR) of the LDO for all PT corners (-40°C, 170°C, slow-slow, fast-fast) and for V_{DD} from 1V (best case supply voltage) down to 600mV. All failing cases are with V_{DD}=600mV, which is way lower than our worst-case V_{DD} even after having considered IR-drops.

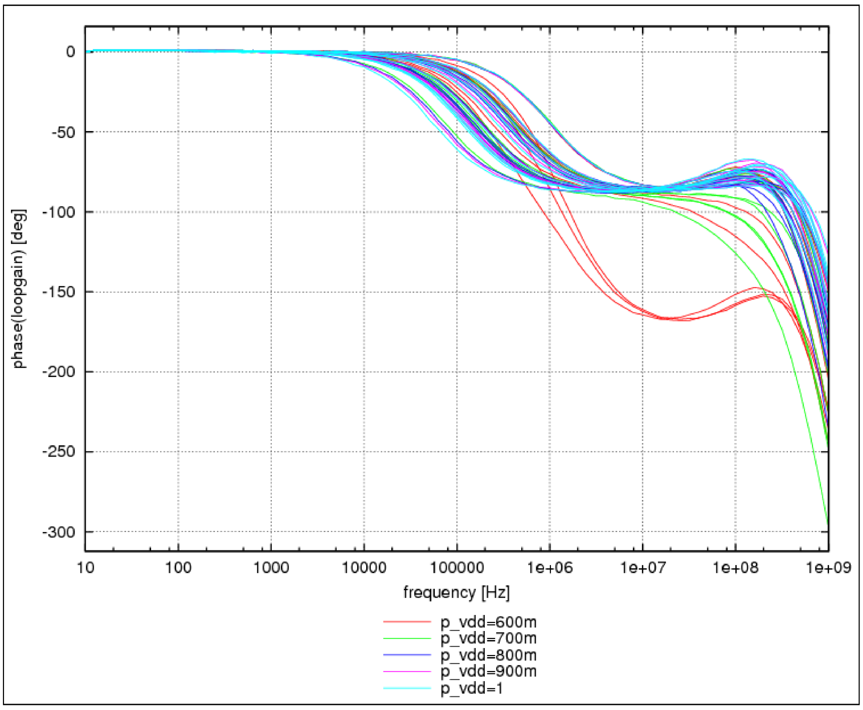

Figure 3.10: Phase Margin of the LDO for all PT corners (-40°C, 170°C, slow-slow, fast-fast) and for V_{DD} from 1V (best case supply voltage) downto 600mV.

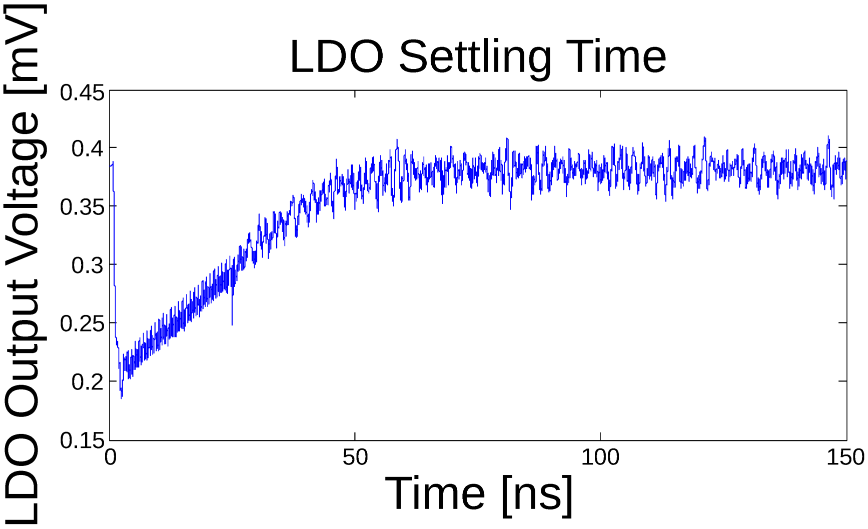

Figure 3.11: Simulated waveform of the output voltage of the LDO of Fig. 3.5 when the *Power On* signal ramps up from 0 (at time $t = 0$) to V_{DD} within 1 ns.

3.5 Simulation Results

The simulated eye diagram for the transistor-level implementation of the transmitter of Fig. 2.27 with FFE activated is reported in Fig. 3.12. The picture highlights the presence of 16 output equalization levels even when operating at 10 Gbps as well as short rise and fall times. No channel has been attached at the output of the transmitter, only a 100 Ω differential load, ESD protection, pad and wiring parasitics. The capacitance at the output node due to these three contributors sums up to 300fF as from post-layout evaluation, due to the large area of the pads that will be shown in the layout view of the test-chip in the next chapter. At this stage, an on-chip decoupling capacitance of 50pF for each supply domain is considered.

The simulated power consumption is 22.5 mW in the least consuming corner (V_{DD}=820mV, temperature of -40°C and slow-slow technology corner), corresponding to a remarkably low 2.25 pJ/bit, aligned with state-of-the-art (Tab. 3.1). We highlight that the reported eye height of 300mV is the one of the same least consuming corner: when checking this figure of merit in other corners, this value rises up to more than 400mV. In fact, if needed, one can also scale down the eye height even more by reducing the LDO output voltage that supplies the driver slices.

We highlight that all other works cited in Tab. 3.1 except for [46] are intended for consumer electronics, not for automotive. At the same time, the high programmability offers the possibility to tune the transmitter impedance and weights over all possible technology corners, extremely simplifying the design of the corresponding receiver and significantly lowering also its power consumption. In fact, the equalization capability is exactly where our transmitter outdistances from the others. The driver and predriver alone consume 0.633 pJ/bit, whereas the shift-register, tap sign selection, switch matrix and re-sampling consume 0.855 pJ/bit. Remaining power consumption (0.762 pJ/bit) derives from LDO and clock tree, the latter being power-hungry due to the presence of the re-sampling stage.

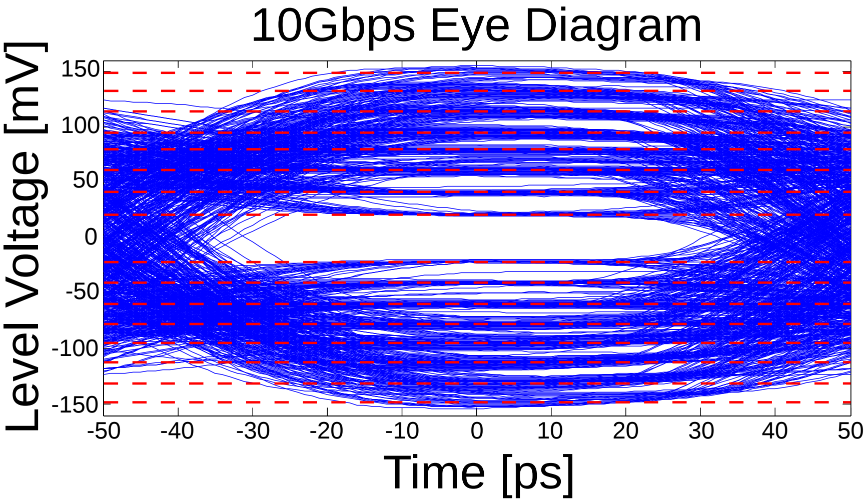

Figure 3.12: Eye diagram at the output of the transmitter resulting from the transistor level, post-layout simulation. $V_{DD} = 1V$. An internal PRBS-12 signal has been used as input data. The tap weights has been set as $w_0 = \frac{9}{16}$, $w_1 = -\frac{3}{16}$, $w_3 = \frac{1}{16}$, $w_4 = -\frac{2}{16}$ and $w_5 = \frac{1}{16}$. The red lines represent the 16 voltage levels at the center of the eye.

Table 3.1: BENCHMARKING OF TRANSMITTER PERFORMANCE WITH LITERATURE

	This work	[66]	[52]	[62]	[51]	[26]	[71]	[56]	[67]	[39]	[97]	[40]
Tx Arch.	SST	SST	SST	SST	SST	SST	SST	VM	CM	CM	CM	CM
Technology [nm]	28	40	90	65	22	32-SOI	28	90	90	28	28	65
Data Rate [Gbps]	10	2.5	6.25	8.5	16|32	28	56.2	5	10	20	40	64.5
TX Eq. [FFE Taps]	8	2	none	2	3	4	2	PWM-based	4	2	2	4
Eye Height [$V_{PKPK-DIFF}$]	0.3	0.9	0.125	1	0.515|0.630	1.05	0.2	0.66	1	0.5	0.32	0.85
Efficiency [pJ/bit]	2.25	3.35	0.8	11.3	4.6|6.4	7.75	1.87	3.1	7	6.5	6.15	3.1
V_{DD} [V]	1	1.25	1	1.5	0.9|1.07	1.1	0.96	1.1	1.2	1.35	1.1	1.2
Area [mm^2]	0.05^{*}	$1^{\S}$	$0.535^{\S}$	0.0682^{*}	$0.08\|0.08^{\S}$	$0.81^{\dagger}$	$1.4^{\ddagger}$	$0.627^{\pounds}$	0.16^{*}	$0.12^{\S}$	$0.81^{\dagger}$	1.2^{*}

Layout area includes the following parts: * Only TX, § TX+RX+PLL, † TX+RX+PLL/4, ‡ TX+RX, $^{\pounds}$ TX+PLL

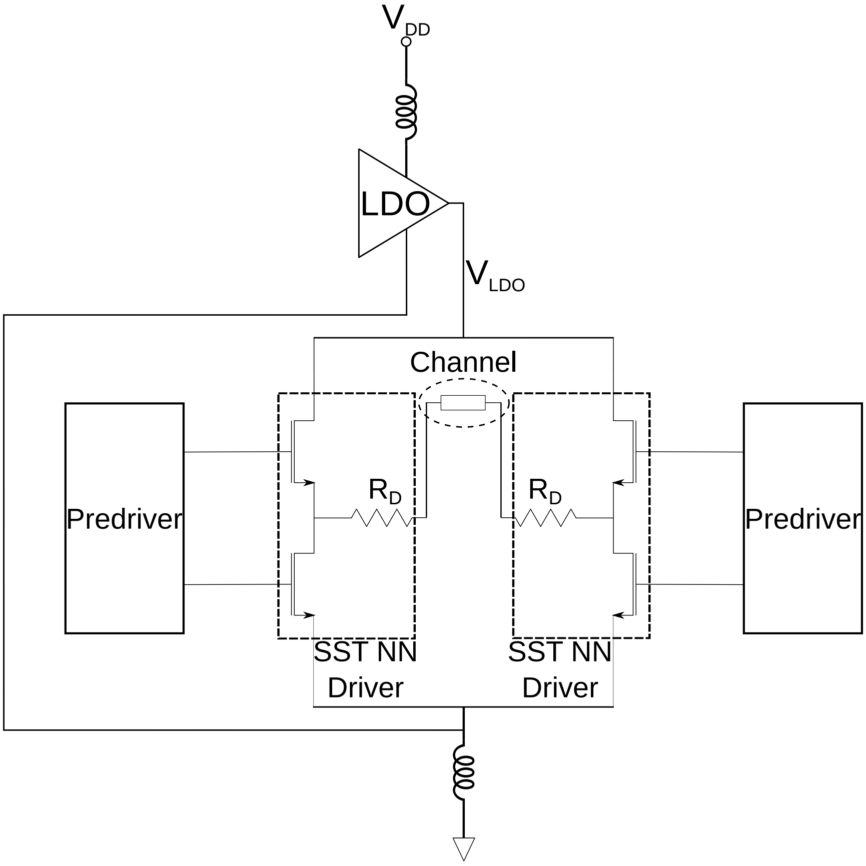

Figure 3.13: Scheme of the predriver and driver slices on the right of Fig. 2.27. Each driver slice is supplied by the LDO output (nominally, 400 mV). The channel has 100 Ω differential resistance.

3.6 Analysis of the effect of parasitic inductances

In the following sections, we use transistor level and system level simulations of such transmitter coupled with s-parameters description of realistic channels obtained by accurate electromagnetic simulations to assess the performance of the link and the effectiveness of FFE in presence of realistic parasitic inductance on the supply connections [55] in a scenario where the chip is bonded into its package, showing that they have a non-negligible impact on the link performance and tend to cancel the benefits of FFE. To our best knowledge, a systematic study of the impact of supply inductance on serial link transmitters has not been presented so far and complements similar analysis [56] devoted to the effect of parasitic inductance on interconnects. Thanks to the presence of the re-sampling stage that eliminates all jitter up to that point, we can just focus on the supply inductances present on the driver slices, as shown in Fig. 3.13

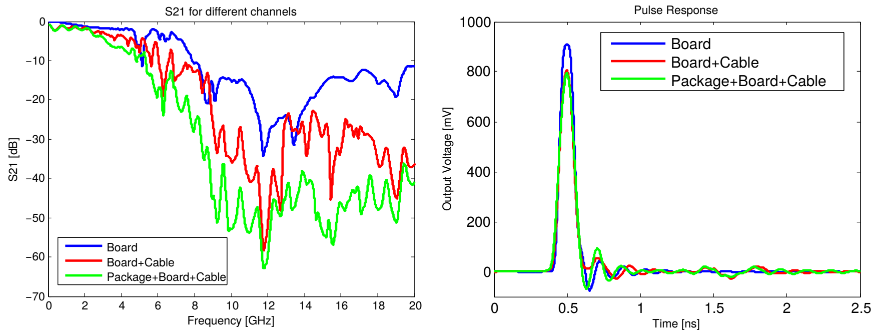

a) b)

Figure 3.14: S_{21} (a) and pulse response (b) for three different channels: board and 50 Ω termination (blue curve), board, cable and 50 Ω termination (red curve) and package, board, cable and 50 Ω termination (green curve). The cable is a 10 cm long Samtec cable and the package is a BGA.

Table 3.2: Tap weights for FFE for the channels considered in Fig. 3.14.

	Board	Board+Cable	Package+Board+Cable
w(-1)	-0.003	-0.0243	-0.0464
w(0)	0.8638	0.8298	0.7865
w(1)	-0.08	-0.0627	-0.0203
w(2)	-0.002	-0.0458	-0.0894
w(3)	0.0215	0.0131	0.0384
w(4)	-0.0214	-0.0076	-0.0069
w(5)	0.0032	0.0074	-0.0101
w(6)	0.0052	-0.0092	0.0020

3.6.1 Selection of the FFE taps

In order to evaluate the improvement obtained with FFE and the detrimental effect of supply parasitic inductance, we have to adapt the method introduced in Chapter 2 to our transistor design specifications. Once again, we have considered three realistic channels, whose $|S_{21}|$, as obtained from EM simulations with Ansys SIWave [45] is shown in Fig. 3.14a. We consider a system composed by package+board+cable and some of its subsystems (board alone and board+cable). The overview of the package, a BGA used for automotive applications, as it appears in SIWave is reported in Fig. 3.15.

After simulating the response to a 100ps pulse (using the software ANSYS Electronic Desktop [44]), see Fig. 3.14b, we use Eq. 2.13 to obtain the weights of the FFE taps. The resulting w_i from Eq. 2.13 are reported in Tab. 3.2 and have to be quantized in steps of $\frac{1}{16}$.

In the following, we will compare system level simulations where PRBS transmission is implemented using a parametric high-level transmitter model built-in into ANSYS Electronic Desktop and s-param representation of the channel with a more involved (and time-consuming) procedure, where the s-param description of the channel and the transistor level description of the circuit (including layout parasitics) are coupled together [46]. The rise and fall times of the high-level transmitter model have been adjusted in order to match the

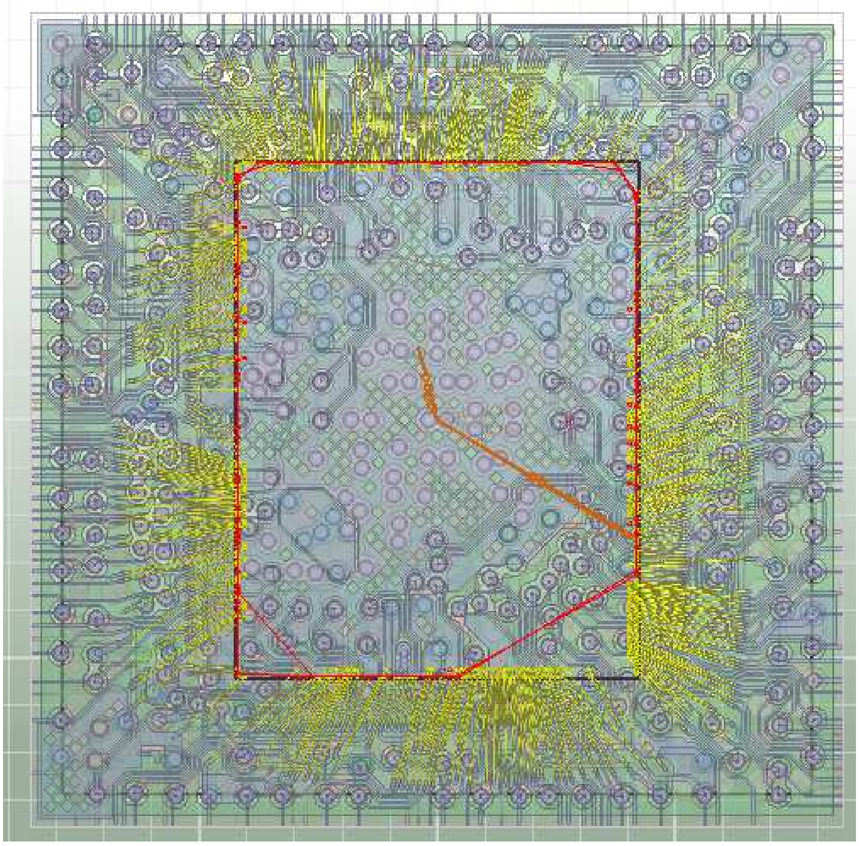

Figure 3.15: Overview of a BGA292 package as it appears in Ansys SIWave.

pulse at the output of the transmitter as obtained from the transistor level simulations without V_{DD} and V_{SS} inductance, as can be seen in Fig. 3.16. We include in all simulations (transistor level and system level, for all channels) the bonding inductance (1nH) on the output pad and the corresponding pad capacitance (300fF).

3.6.2 Results without V_{DD} and V_{SS} inductance

Here we compare the performance extracted from the eye diagram considering the full transistor level description of the transmitter or its replacement with its system level description. We set to zero the V_{DD} and V_{SS} inductance. The value k of the active slices among the 12 available is set to 5. Sample eye diagrams with and without FFE are reported in Fig. 3.17.

Fig. 3.18 shows the jitter and signal-to-noise ratio (SNR) for the three channels of Fig. 3.14 without activating FFE. As expected, channels with larger attenuation show larger jitter and reduced SNR. We also see that system level and transistor level simulations provide similar results, consistently with the matched pulses shown in Fig.3.16. We observe jitter also without channel since

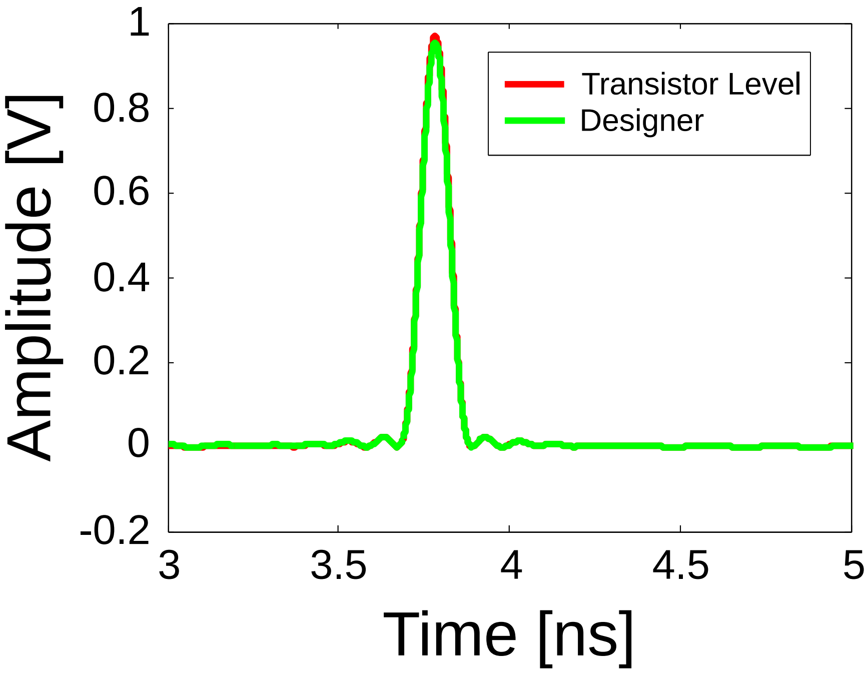

Figure 3.16: Comparison between pulse response at the output of the transmitter as obtained from the transistor level simulations without V_{DD} and V_{SS} inductance (red curve) and pulse response as obtained from high-level transmitter model (green curve).

we have included the bonding inductance and the capacitance on the signal pads.

When we activate FFE, Fig. 3.19, SNR improves, whereas the jitter does not. This is expected, since the ZFM (Eq. 2.10) targets opening in the center of the eye (see Fig. 3.17). Furthermore, the effectiveness of the FFE is affected by the need to quantize the value of the tap weights (blue vs green bars in Fig. 3.19) as required by the use of a finite number of driver slices.

The agreement between transistor level and system level simulations is good in Fig. 3.19 as was in Fig. 3.18, provided the tap weights are quantized in the system level simulations (green vs red bars). A designer can thus exploit system level simulations in an early phase of study, since these are on average thousand times faster than transistor level ones.

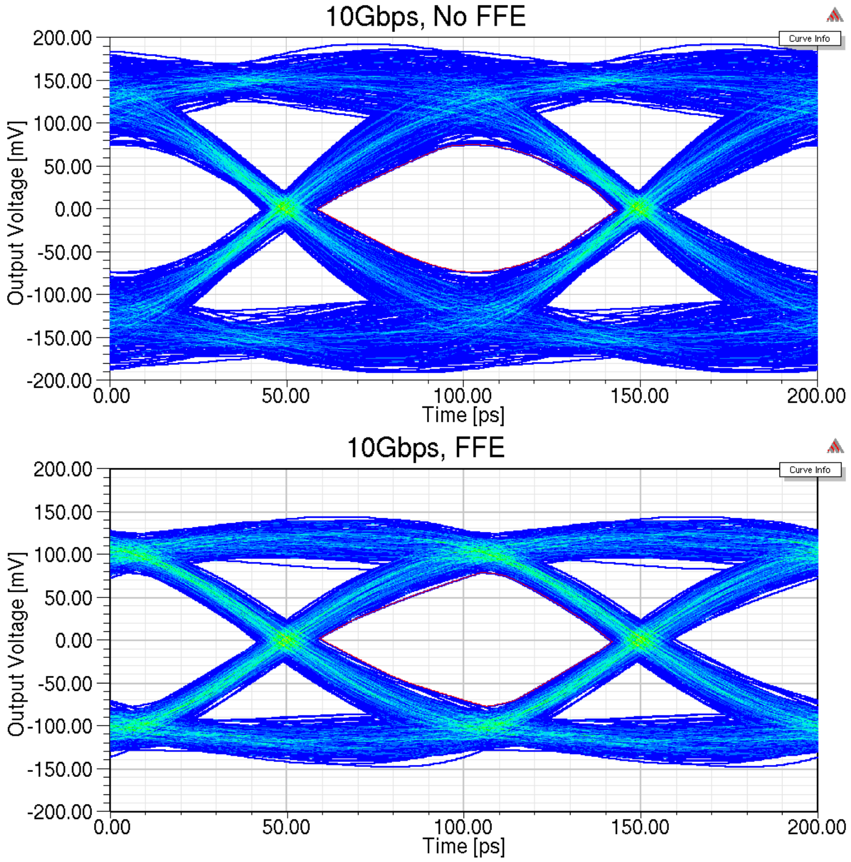

Figure 3.17: Comparison between eye diagrams obtained with Transistor Level simulations without FFE (up) and with FFE activated (bottom). Both eye diagrams refer to the channel composed by package, board and cable and without considering the inductance on V_{DD} and V_{SS}.

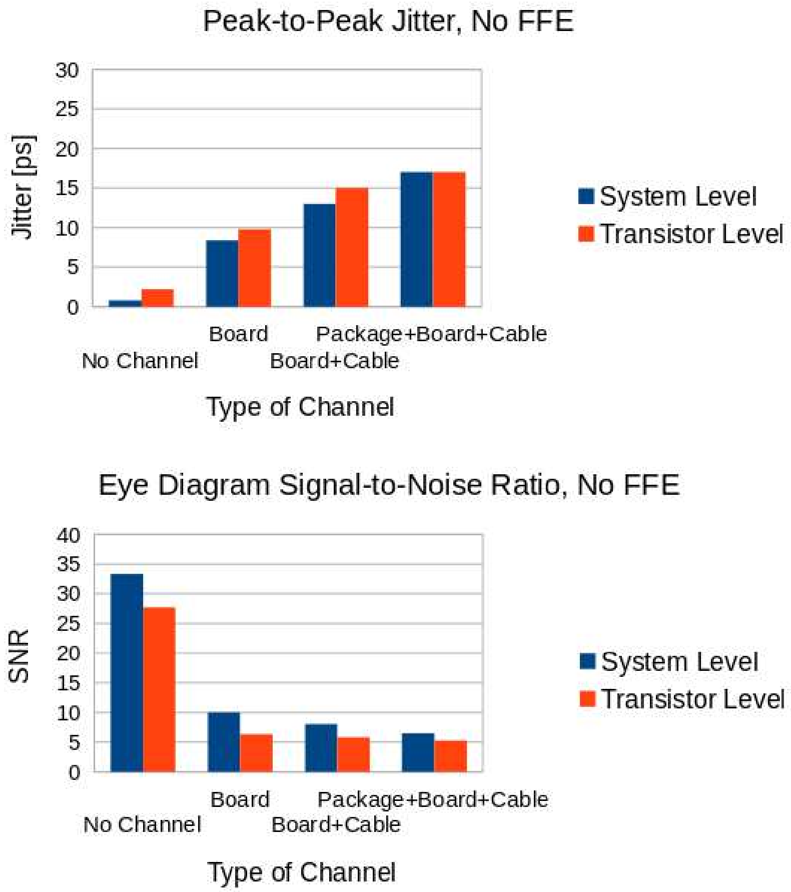

Figure 3.18: Peak-to-Peak Jitter (width of the horizontal histogram across the eye-crossing point) and Signal-to-Noise Ratio (difference between the average ’1’ and ’0’ levels divided by the sum of the standard deviations of the two levels at the center of the eye) for all channels when FFE is not applied. System (TX modeled as a PRBS generator as available in Ansys Electronic Desktop) and transistor level simulations.

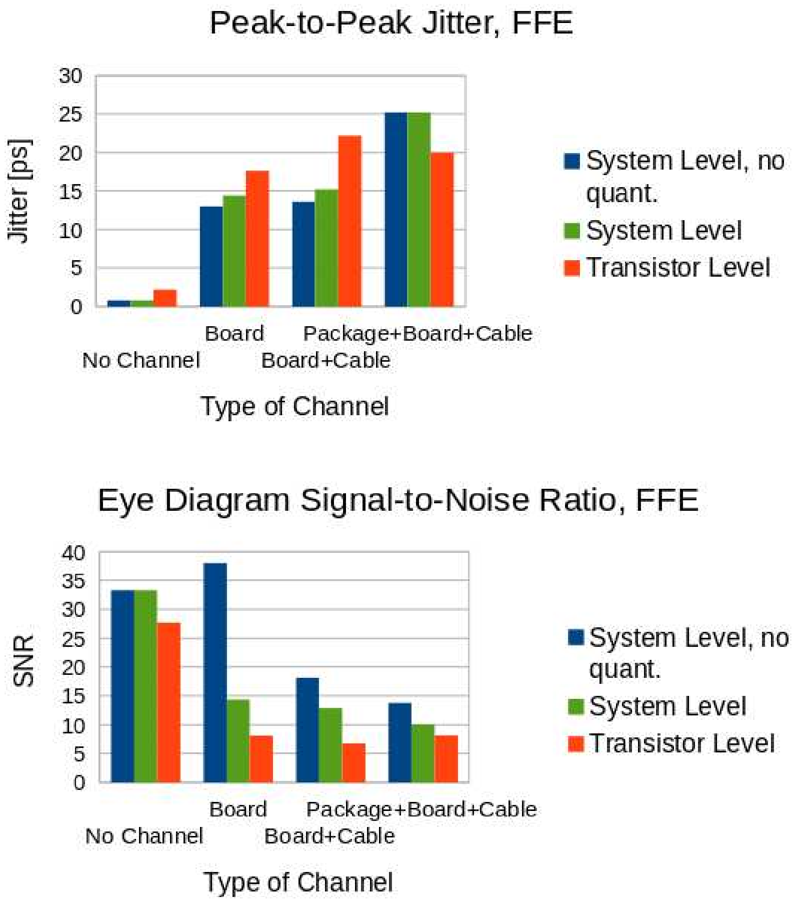

Figure 3.19: Peak-to-Peak Jitter and Signal-to-Noise Ratio as Fig. 3.18 but with FFE. The blue curves report results of system level simulations in which the equalization weights w_i have not been quantized, the green curves of system level simulations in which equalization weights w_i have been quantized in steps of $\frac{1}{16}$ and the orange ones of transistor level simulations having the same w_i quantized in steps of $\frac{1}{16}$.

3.6.3 Results including VDD and VSS inductance

Current peaks as large as 8mA (referred to the DC current) and with duration of less than 1ns are observed in the transistor-level simulations (Fig. 3.20), pointing out the importance of considering the parasitic supply inductance on V_{DD} and V_{SS}. We have also included typical parasitic series resistance (0.4-0.8 Ω per supply domain), but they do not play any role (not shown). The histogram in Fig. 3.21 refer to a situation when FFE is activated and shows that the presence of the parasitic inductance significantly increases the jitter and degrades the SNR. This can be observed (for the specific channel with package+board+cable) in the eye diagrams in Fig. 3.22. Comparison between Fig. 3.21 and Fig. 3.19 shows that the supply inductance partly cancels the performance enhancement associated with FFE: the time- and patter-dependent voltage drop results in additional ISI.

To further investigate the effect of the parasitic inductance, we plot in Fig. 3.23 the jitter and SNR for different values of k (number of activated replica of the vector of 16 slices). The figure also summarizes some of the findings of the previous analysis. We see that:

i) FFE has a minor impact on jitter;

ii) the effect of parasitic inductance on jitter is large;

iii) even when FFE is not activated, the effect of the parasitic inductance on SNR is large;

iv) the parasitic inductance tends to hamper the SNR improvement associated to FFE.

About the effect of increasing k, SNR is enhanced and jitter is reduced (except than the case of FFE with inductance where the jitter is almost constant), mainly because the lower output resistance of the driver results in larger signals at the output. The corresponding impedance mismatch, however is expected to have a detrimental effect in the presence of impedance discontinuities along the line or at the receiver side (that here is assumed to be an ideal 100Ω differential load), since this would result in multiple reflections (i.e. additional ISI).

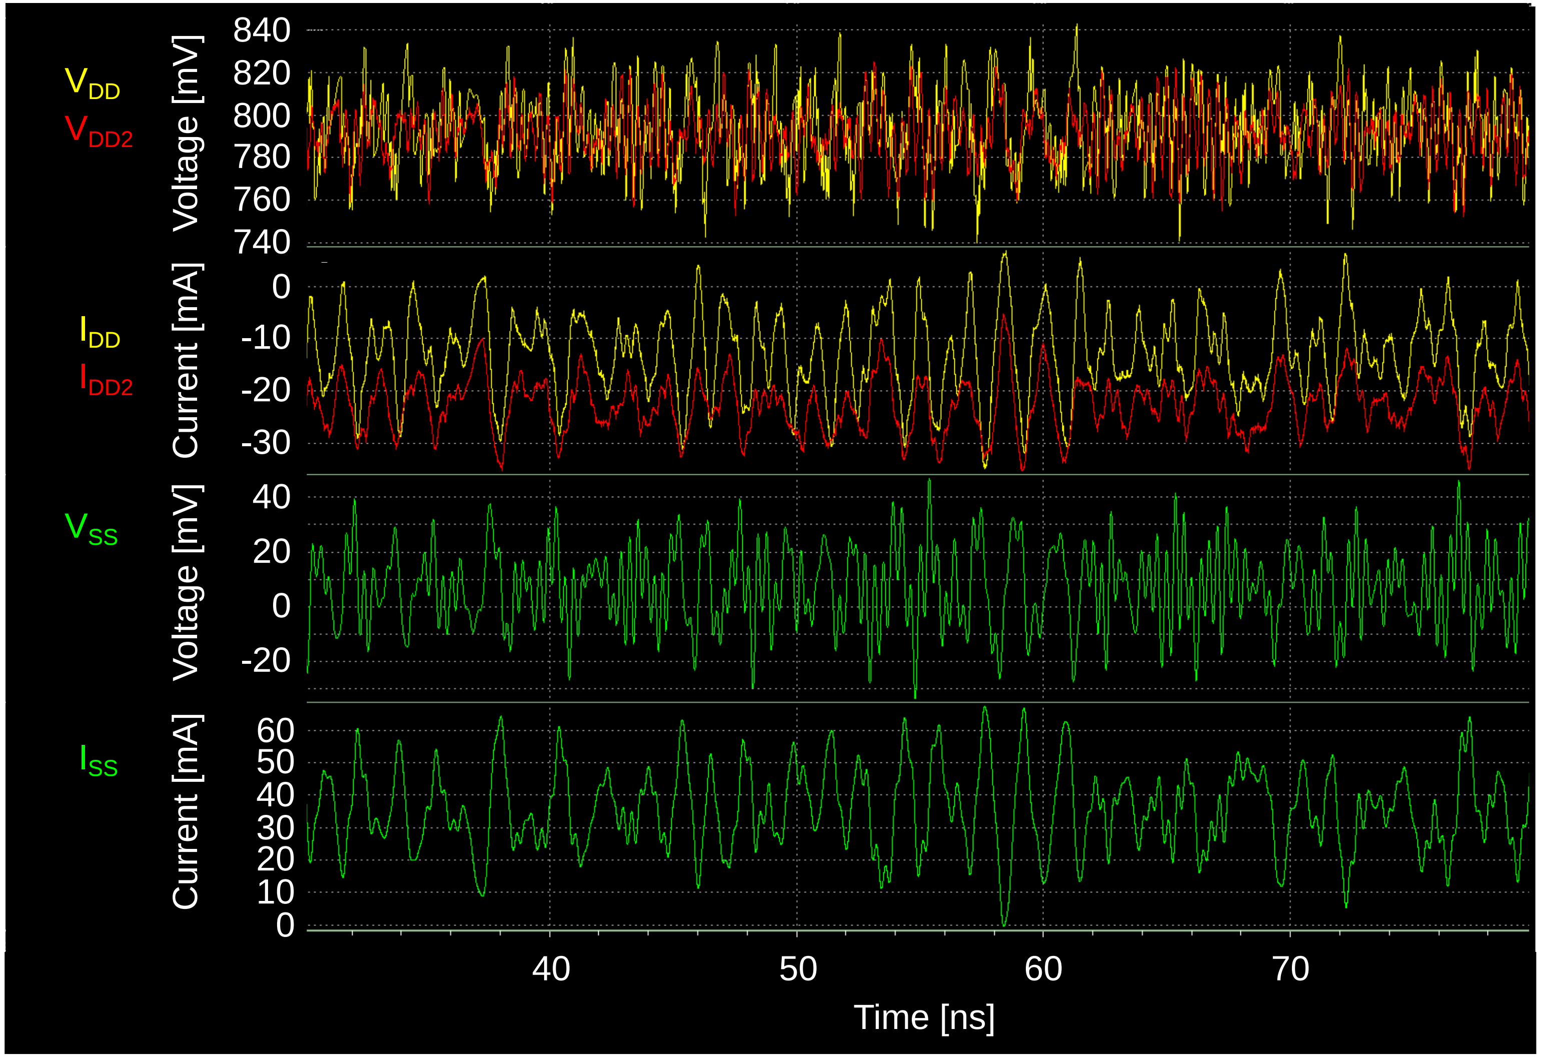

Figure 3.20: Current spikes on both supply domains (V_{DD} and V_{DD_2}) and V_{SS}. For each supply domain, along with the current consumption, also the voltage curves are reported, so to highlight the correlation between current and $L\frac{dI}{dt}$ voltage drops.

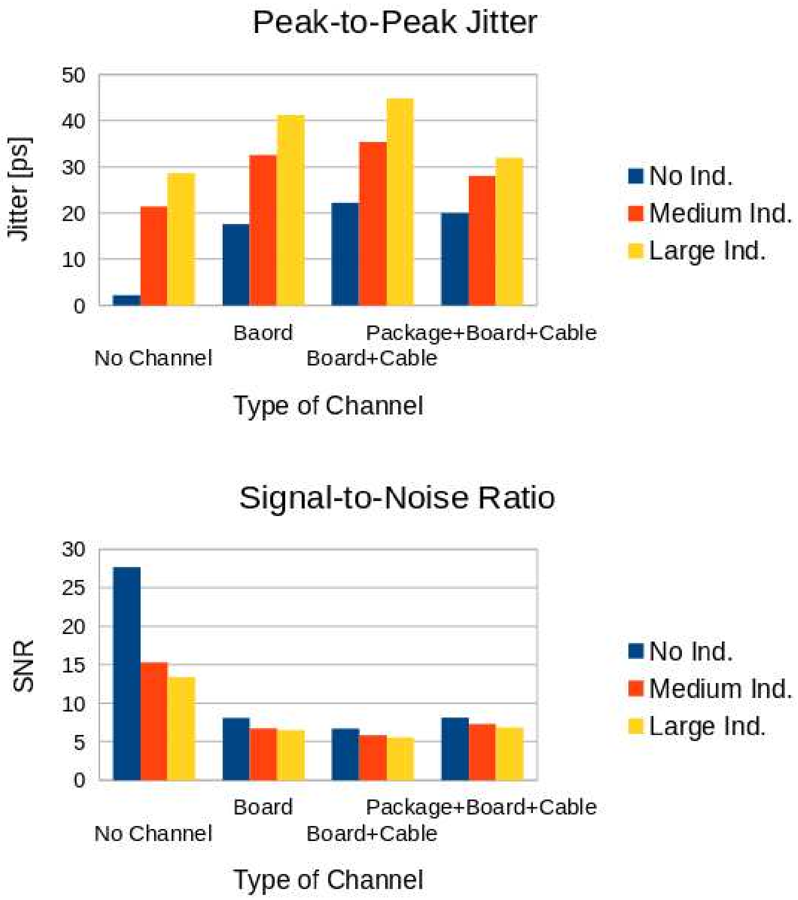

Figure 3.21: Effect of V_{DD} and V_{SS} parasitic inductance on both Peak-to-Peak Jitter and Signal-to-Noise Ratio when FFE is applied. All results refer to transistor level simulations. *Medium Ind.* means L=0.4nH on V_{DD} and L=0.2nH on V_{SS}, whereas *Large Ind.* means 0.8nH and 0.2nH on V_{DD} and V_{SS}, respectively.

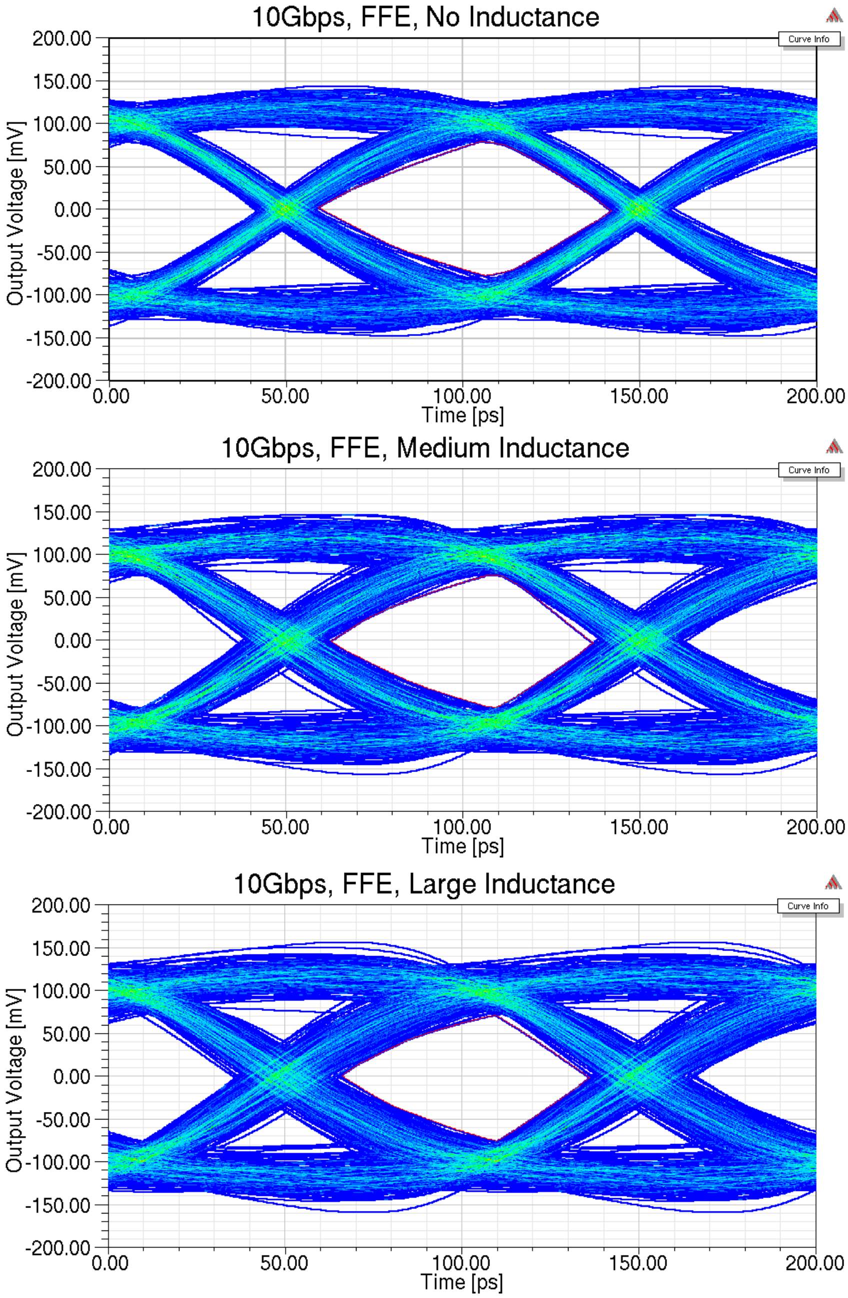

Figure 3.22: Comparison between eye diagrams obtained with transistor level simulations for different V_{DD} and V_{SS} parasitic inductance: no inductance (top), *Medium Ind.* (middle) and *Large Ind.* (bottom). All eye diagrams refer to the channel composed by package, board and cable of Fig. 3.14 and are obtained with FFE activated. The weights w_i used are the ones reported in the rightmost column of Tab. 3.2.

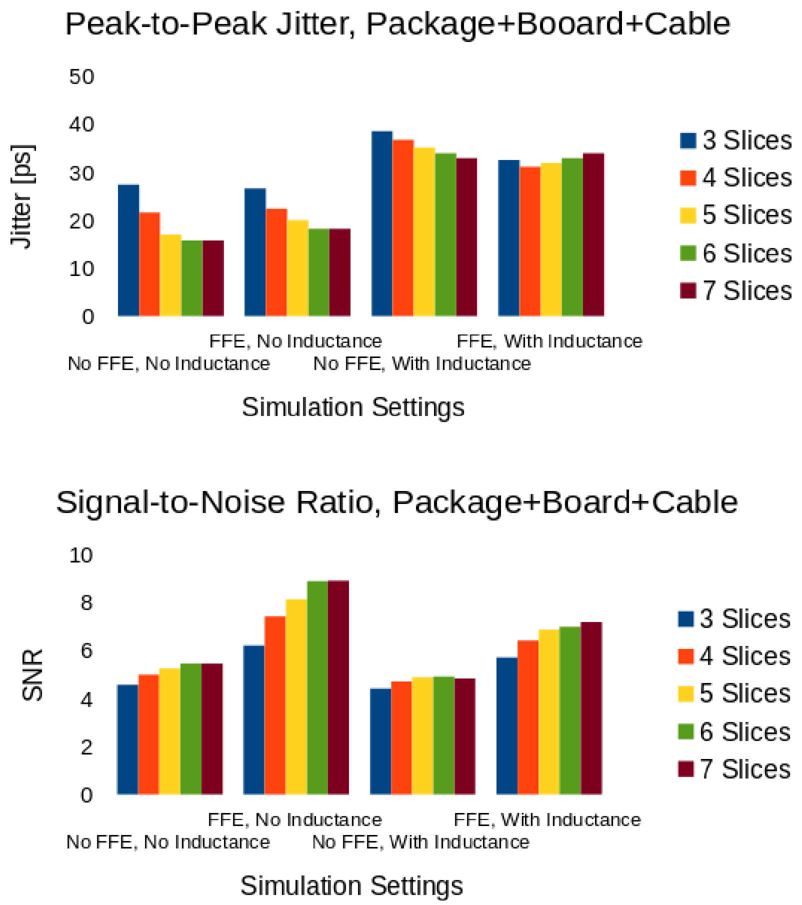

Figure 3.23: Effect of the number *k* of activated slices on Peak-to-Peak Jitter and SNR for the channel composed by package, board and cable.

3.7 Conclusions

In this chapter, we have described the transistor level design of a 8-tap FFE 10Gbps transmitter [57] and assessed its capability to operate with realistic channels by means of back-checking with the methodology proposed in Chapter 2. First, we analyzed the performances of the single blocks as from post-layout simulations and checked the correctness of the resulting eye diagram at the output of the transmitter. Then, we have compared the performances of our transmitter with the literature, highlighting how the power consumption and the equalization capability align this work with the state-of-the-art. Finally, we have focused on the effect of parasitic inductances on the performances of the transmitter, especially on the eye diagram parameters. Keeping under control the supply parasitic inductance is mandatory in order to fully exploit the advantages of feed-forward equalization. For the particular case under study, keeping the V_{DD} inductance below 0.4nH and the V_{SS} inductance around 0.2nH allows recovering an almost ideal situation. When these inductances are low, efficient system level simulations provide accurate performance estimation (in terms of eye parameters such as jitter and SNR) comparable to more accurate (but time consuming) transistor level analysis.

Chapter 4

Experimental Characterization of the Transmitter

In this chapter we analyze the results of the measurements of two different test-chips of the 10Gbps transmitter, the first one (in the following called RMTX) taped-out after the design described in the previous chapter, the second one(in the following called InnoTC) already containing the full link. For each one of the two test-chips, we will start with a quick overview of the test-chip, then pass to the eye diagrams obtained with different timing settings, with and without FFE.

4.1 RMTX Test-Chip Overview

Fig. 4.1 shows the structure of RMTX test-chip, taped out in March 2017 and started to be measured in August 2017. The 10Gbps transmitter can be seen on the left of the picture, at the left of to the two high speed output pads. The area of the test-chip is 0.98mm^2, but it can be seen that the pads are occupying almost half of the test-chip area, whereas a large portion is occupied by the 20GHz free-running VCOs which supplies the clock to the transmitter, after this being divided by 2. At the last stage of development, the chip has been integrated with a hybrid SPI/JTAG interface in order to program the settings for both VCO and transmitter. The remaining area has been filled with decoupling capacitance, both high and low-Q in ratio 1:1, for a total of roughly 40pF for each one of the two power domains.

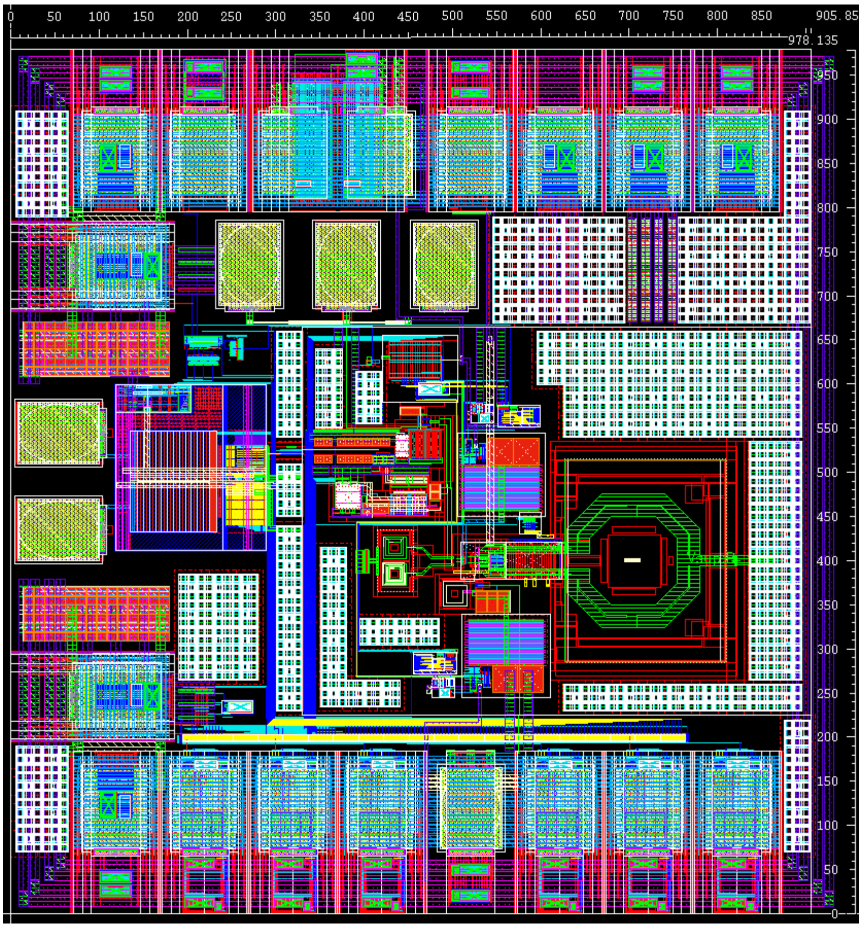

Figure 4.1: Layout view of the RMTX test-chip. On the top and the bottom of the test-chip analog and digital pads can be seen, whereas on the left-hand side two high speed output pads are present. On the low side of the picture, the hybrid SPI/JTAG interface can be seen (yellow). On the right side, the VCO and its coil are occupying a big chunk of the area.

4.2 RMTX Measurements Results

Fig. 4.2 shows an eye diagram for the transmitter without FFE with a voltage supply of 1.04V for both power domains. The VCO has been set to obtain an oscillation frequency of roughly 20GHz, so that the transmitter is actually running at 9.97Gbps. In fact, the settings for the VCO are far from being optimized, even if the Random Jitter shows up to be as low as 651fs. The eye diagram looks clearly to be RC limited and shows the typical behavior of a not sufficiently decoupled eye diagram (having 50pF of high-Q decoupling capacitance for each supply domain would have been better). On top of that, in case of this test-chip it was not possible to double bond the chip to the supply pads in order to halve the parasitic supply inductances and the bonds look quite long, around 1mm, which in turn translates in circa 1nH, way over the guidelines derived from the previous chapter. The '0' and '1' levels are smaller than expected: in fact, the LDO has been designed to supply roughly 400mV at the driver replicas when the supply is 900mV, so we would expect this voltage to scale to circa 460mV. In reality we measure an eye amplitude of 316mV.

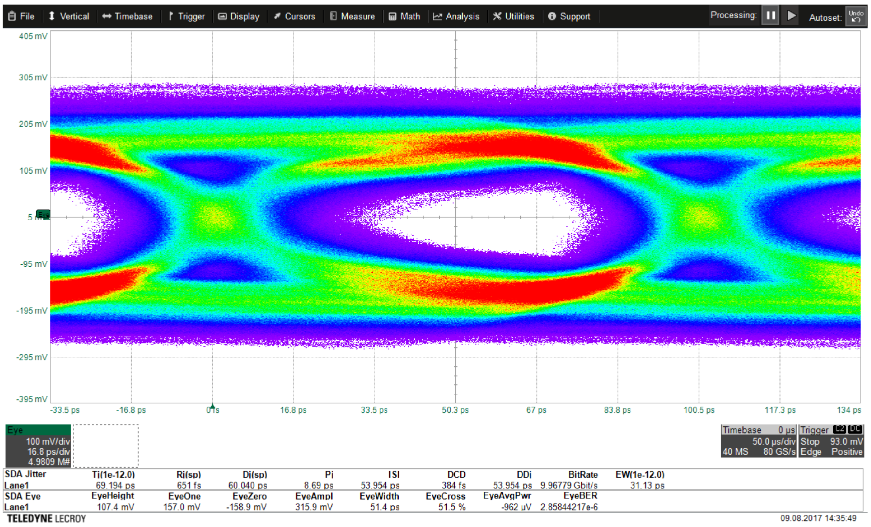

SDA Jitter	Tj(1e-12.0)	Rj(sp)	Dj(sp)	Pj	ISI	DCD	DDj	BitRate	EW(1e-12.0)
Lane1	69.194 ps	651 fs	60.040 ps	8.69 ps	53.954 ps	384 fs	53.954 ps	9.96779 Gbit/s	31.13 ps
SDA Eye	EyeHeight	EyeOne	EyeZero	EyeAmpl	EyeWidth	EyeCross	EyeAvgPwr	EyeBER	
Lane1	107.4 mV	157.0 mV	-158.9 mV	315.9 mV	51.4 ps	51.5 %	-962 µV	2.85844217e-6	

Figure 4.2: Measured eye diagram at the output of the transmitter for a supply voltage of 1.04V on both power domains. FFE has not been activated. Data Dependent Jitter is the main source of noise and sums up to 53.9ps, to give an eye width of 51.4ps. The eye height is 107.4mV.

After a 10Gbps measurement, we set the VCO to the best possible settings, corresponding to a 18.8GHz oscillation frequency, thus obtaining a 8.92Gbps transmission rate. As done before, we evaluate the results when setting both power domains at 1.04V, obtaining the eye diagram reported in Fig. 4.3. In fact, now we obtain a eye amplitude of 516mV and a data dependent jitter of 44.9ps that, compared to the 53.9ps of a 10Gbps, demonstrates that this main contributor stems mainly from parasitic supply inductances, as by increasing the data transmission rate from 9 to 10Gbps, what changes drastically is $L \cdot \frac{dI}{dt}$, increasing the supply bouncing and, consequently, data dependent jitter.

Since from the eye diagram a large dispersion in '0' and '1' levels is visible, one driver slice out of the eight switchable ones have been connected to the first post-cursor tap, resulting in $w_0 = 0.9375$, $w_1 = -0.0625$. A clear improvement at the center of the eye, which directly reflects in an increase of eye height (from 255.7mV to 295.8mV), is clearly visible in Fig. 4.4 as direct consequence of FFE, and is obtained at the cost of 3.6ps less in eye width (from 64.5ps to 60.9ps). The eye amplitude reduces by circa 12.5% as expected, and is now 445mV. An extended characterization of the transmitter placed on RMTX has not been carried out since the results obtained with the transmitter placed on InnoTC have resulted to be way better than what obtained with RMTX.

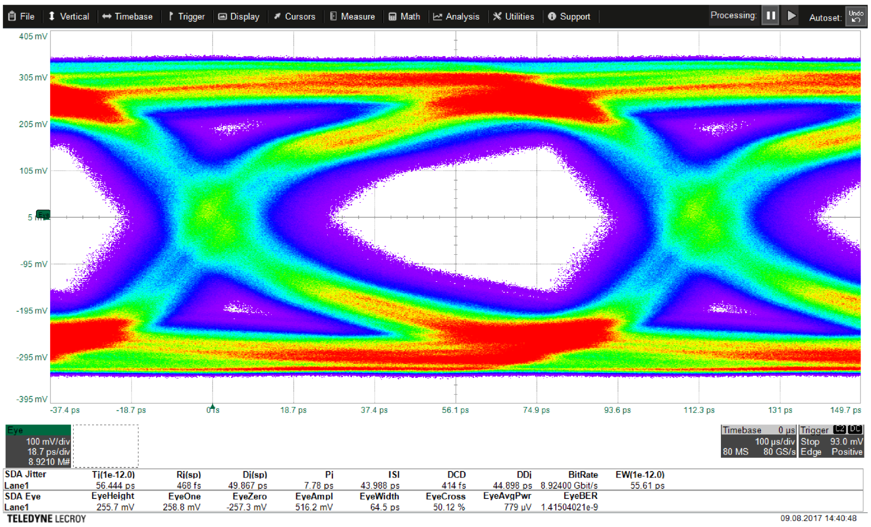

Eye
100 mV/div
18.7 ps/div
8.9210 M#

Timebase 0 µs
100 µs/div
80 MS 80 GS/s

Trigger C2 DC
Stop 93.0 mV
Edge Positive

SDA Jitter	Tj(1e-12.0)	Rj(sp)	Dj(sp)	Pi	ISI	DCD	DDj	BitRate	EW(1e-12.0)
Lane1	56.444 ps	468 fs	49.867 ps	7.78 ps	43.988 ps	414 fs	44.898 ps	8.92400 Gbit/s	55.61 ps

SDA Eye	EyeHeight	EyeOne	EyeZero	EyeAmpl	EyeWidth	EyeCross	EyeAvgPwr	EyeBER
Lane1	255.7 mV	258.8 mV	-257.3 mV	516.2 mV	64.5 ps	50.12 %	779 µV	1.41504021e-9

Figure 4.3: Measured eye diagram at the output of the transmitter for a supply voltage of 1.04V on both power domains. FFE has not been activated. Data Dependent Jitter is the main source of noise and sums up to 44.9ps, to give an eye width of 64.5ps. The eye height is 255.7mV.

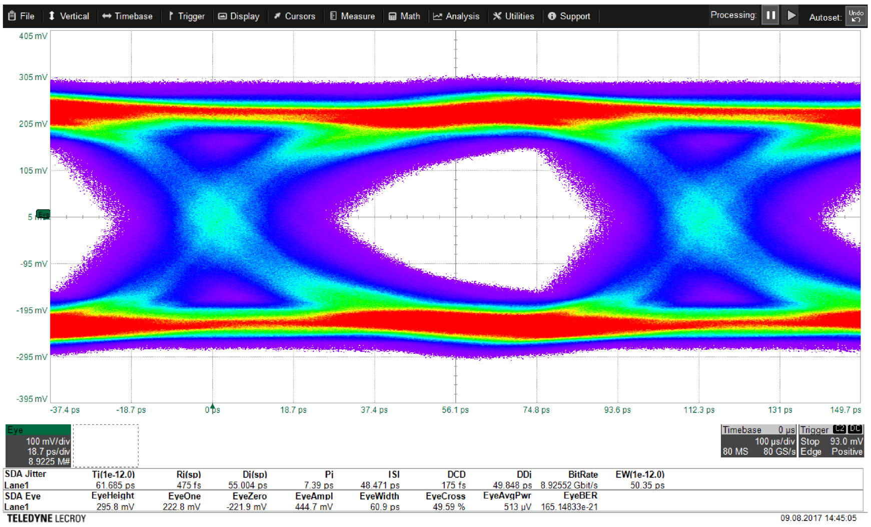

SDA Jitter	Tj(1e-12.0)	Rj(sp)	Dj(sp)	Pj	ISI	DCD	DDj	BitRate	EW(1e-12.0)
Lane1	61.685 ps	475 fs	55.004 ps	7.39 ps	48.471 ps	175 fs	49.848 ps	8.92552 Gbit/s	50.35 ps

SDA Eye	EyeHeight	EyeOne	EyeZero	EyeAmpl	EyeWidth	EyeCross	EyeAvgPwr	EyeBER
Lane1	295.8 mV	222.8 mV	-221.9 mV	444.7 mV	60.9 ps	49.59 %	513 µV	165.14833e-21

Figure 4.4: Measured eye diagram at the output of the transmitter for a supply voltage of 1.04V on both power domains. FFE weights are $w_0 = 0.9375$, $w_1 = -0.0625$. Data Dependent Jitter is the main source of noise and sums up to 49.8ps, to give an eye width of 60.9ps. The eye height is 295.8mV.

4.3 InnoTC Overview

The layout of this test-chip and its description will be discussed in Chapter 6, however there are four major differences between the two test-chips for what concerns the transmitter. The first one is that the only available PRBS sequence is a PRBS32, which is not recognizable as pattern by our scope: this means that all jitter analysis have not been available for these measurements, thus we cannot separate the various jitter components. The second one is that a different decoupling concept has been used: for each one of the two transmitter supply domain, 105.6pF of high-Q and 105.6pF of low-Q decoupling capacitance have been used, so to be sure not to see any jitter effect related to self-induced noise on the power supply due to parasitic inductance.

The third one is that for each power supply two pads have been used, so to reduce the parasitic supply inductance by a factor of 2. Anyway, the bondwires for this test-chip are longer than for RMTX (circa 2-3mm), so that in the end for each power supply domain we end up having comparable parasitic supply inductances compared to what we had in RMTX (roughly 1-1.5nH). Finally, the clock is not directly supplied by a free-running VCO, but from a digital PLL on chip, which takes itself the clock from a DCO. Anyway, we don't expect to see a big difference between the RMS jitter for the two cases, given the fact that the PLL used is designed for radar applications, so the phase noise of such a PLL is not much higher than the one of a free-running VCO and anyway much smaller than the one of PLLs normally used for such transmission speeds.

4.4 InnoTC Measurements Results

Fig. 4.5 shows an eye diagram for the transmitter without FFE with a voltage supply of 0.95V for both power domains. The VCO has been set to obtain the highest possible oscillation frequency, roughly 18.4GHz, so that the transmitter is actually running at 9.23Gbps. Compared to Fig. 4.3, there are a differences: first of all, the transmitter placed on InnoTC is capable of operating at much lower supply voltage. In fact, 0.95V is the voltage we set at the external supply and, even if the detrimental effect of supply inductance has been almost completely removed, the measurement setup for this test-chip consist of many more components that contribute to IR supply drop and IR drop due to the power distribution is much bigger due to the fact that the circuit placed on this second test-chip is much bigger, therefore also the power distribution lines are way longer. In fact, post-layout RC-extracted simulations hint the supply that reaches the chip pin is on average roughly 10% lower than what placed on the external power supply: so, for 0.95V external supply, we can expect that the supply the chip effectively sees is roughly 0.85V. The eye amplitude is therefore much lower, both because the power supply is lower and only six replicas have been turned on, compared to the 12 we had turned on in Fig. 4.3. We recall that, in nominal conditions, the LDO output voltage, therefore also the eye amplitude, should be roughly 380mV.

Anyway, the eye looks greatly improved: in fact, the horizontal opening is much better in InnoTC than in RMTX (80.1ps vs 64.5ps), and gets even better when factoring the different speed between Fig. 4.5 and Fig. 4.3. The eye of InnoTC shows an eye opening of 0.74UI, compared to the 0.57UI we had in

RMTX. Again, we have to stress that this has been achieved with a much lower power supply. We believe this improvement is mainly given by two facts: the first one, the decoupling concept now suits perfectly our needs; second, a re-layout of the final stage of the driver has reduced the parasitic supply we see at the output of the transmitter by roughly 20%.

Since it looks like the measurement setup introduces a loss of a couple of dB, we tried to activate the FFE: the chosen settings are the same as in Fig. 4.4, so one out of the eight programmable slices has been assigned to the 1-post tap. The resulting eye diagram is shown in Fig. 4.6. Once again, the improvement in eye height is remarkable: from 131.6mV we had in Fig. 4.5 to 173mV of Fig. 4.6, and this improvement comes at a cost of just 2.4ps of horizontal opening.

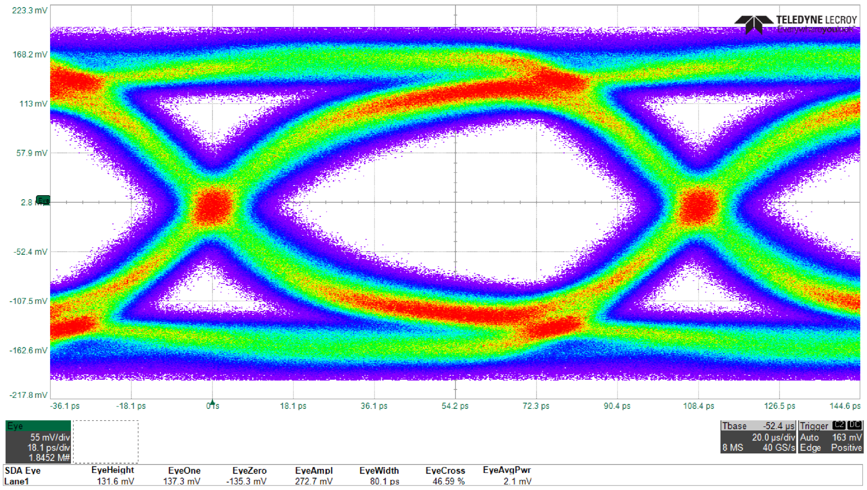

Eye
55 mV/div
18.1 ps/div
1.8452 M#

Tbase -52.4 µs
20.0 µs/div
8 MS 40 GS/s

Trigger C2 DC
Auto 163 mV
Edge Positive

SDA Eye	EyeHeight	EyeOne	EyeZero	EyeAmpl	EyeWidth	EyeCross	EyeAvgPwr
Lane1	131.6 mV	137.3 mV	-135.3 mV	272.7 mV	80.1 ps	46.59 %	2.1 mV

Figure 4.5: Measured eye diagram at the output of the transmitter with six active replicas for a supply voltage of 0.95V on both power domains. FFE has not been activated. The transmission speed is 9.23Gbps. The eye height and width are 131.6mV and 80.1ps, respectively.

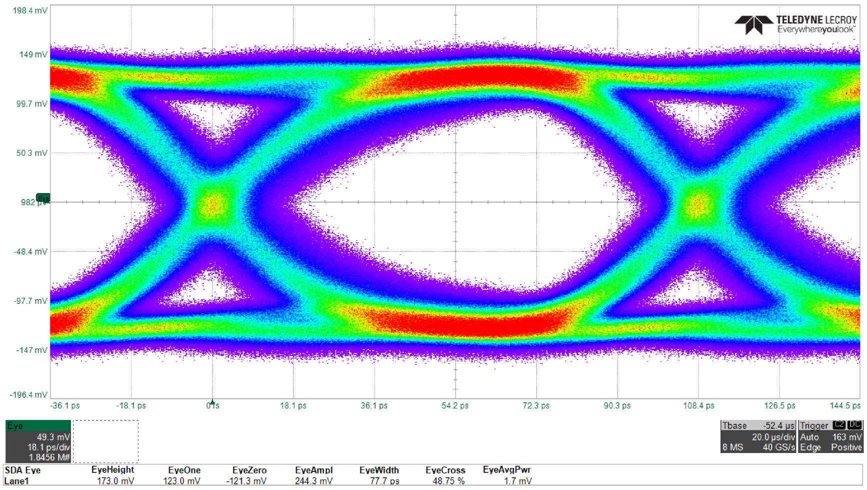

Eye
49.3 mV
18.1 ps/div
1.8456 M#

Tbase -52.4 µs
20.0 µs/div
8 MS 40 GS/s

Trigger C2 DC
Auto 163 mV
Edge Positive

SDA Eye	EyeHeight	EyeOne	EyeZero	EyeAmpl	EyeWidth	EyeCross	EyeAvgPwr
Lane1	173.0 mV	123.0 mV	-121.3 mV	244.3 mV	77.7 ps	48.75 %	1.7 mV

Figure 4.6: Measured eye diagram at the output of the transmitter with six active replicas for a supply voltage of 0.95V on both power domains. FFE weights are $w_0 = 0.9375$, $w_1 = -0.0625$. The transmission speed is 9.23Gbps. The eye height and width are 173mV and 77.7ps, respectively.

A measurement that resulted to be almost impossible to do with RMTX is to check the linearity of the output voltage levels when all 16 possible FFE levels are activated. This measurement can now be done with InnoTC, even tough the sixteen levels are much clearly visible when running at speeds lower than 10Gbps. In fact, this measurement has been run at quarter speed (2.3Gbps). Out of the eight programmable slices, one is connected to the main tap, two to 1-post and 2-post taps and one to 3-, 4- and 5-post taps. The so obtained eye diagram and the measured voltage levels measured at the transmitter output resulting from this measurement are shown in Fig. 4.7: the figure shows the good linearity of the transmitter output voltage, as was previously reported in simulation for RMTX and shown in Fig. 3.12. It must be underlined that, looking at the eye diagram, it results clear that the levels are not equally probable: the probability of hitting one level is in fact just dictated by the chosen FFE weights. For the sake of completeness, we report the eye diagram at the transmitter output for the same settings also for 4.6Gbps (fig. 4.8) and for 9.2Gbps (fig. 4.9): as it can be seen, when increasing the speed the sixteen levels become less and less clearly visible.

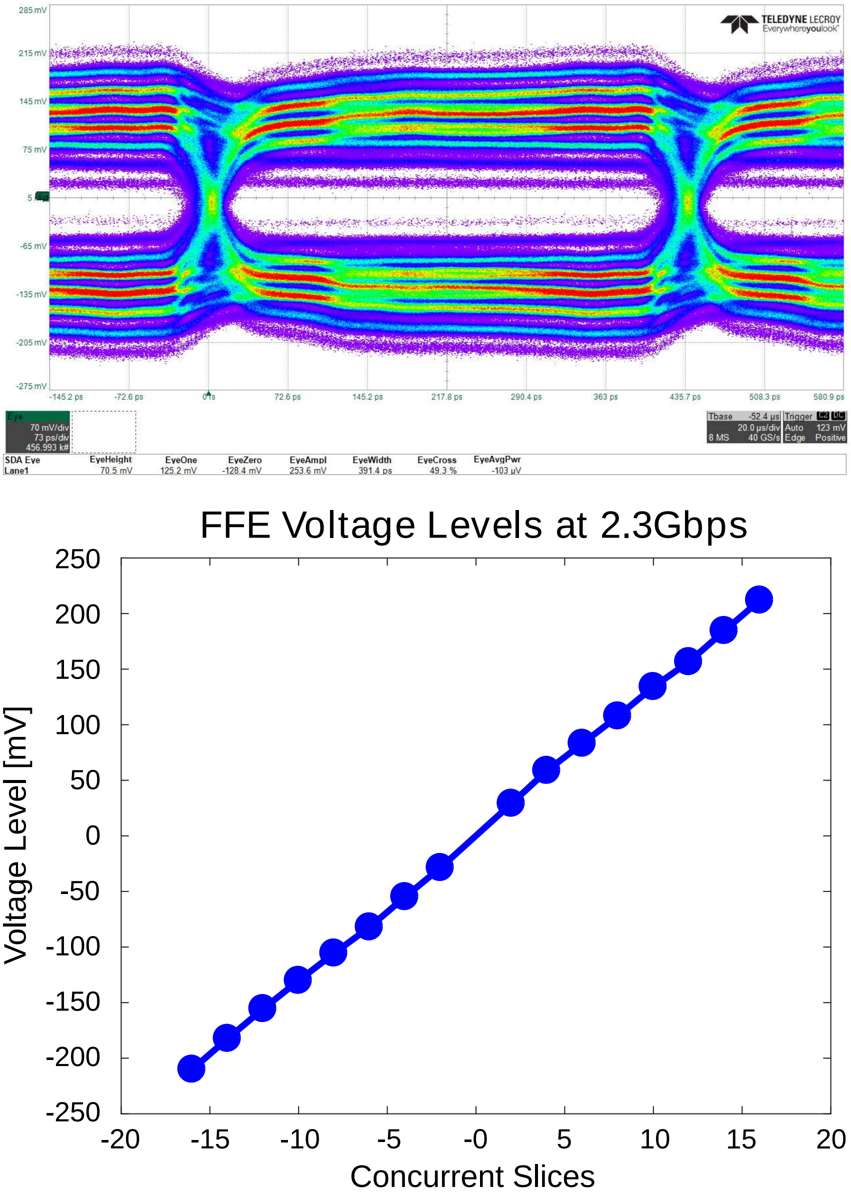

SDA Eye Lane1	EyeHeight	EyeOne	EyeZero	EyeAmpl	EyeWidth	EyeCross	EyeAvgPwr
	70.5 mV	125.2 mV	-128.4 mV	253.6 mV	391.4 ps	49.3 %	-103 µV

Figure 4.7: On top, measured eye diagram at the output of the transmitter with nine active replicas for a supply voltage of 1.05V on both power domains when all 16 FFE levels are activated. On the bottom, the corresponding measured voltage levels FFE weights are $w_0 = 0.56255$, $w_1 = -0.125$, $w_2 = -0.125$, $w_3 = -0.0625$, $w_4 = -0.0625$ and $w_5 = -0.0625$. The transmission speed is 2.3Gbps. The levels show good linearity of the transmitter output voltage.

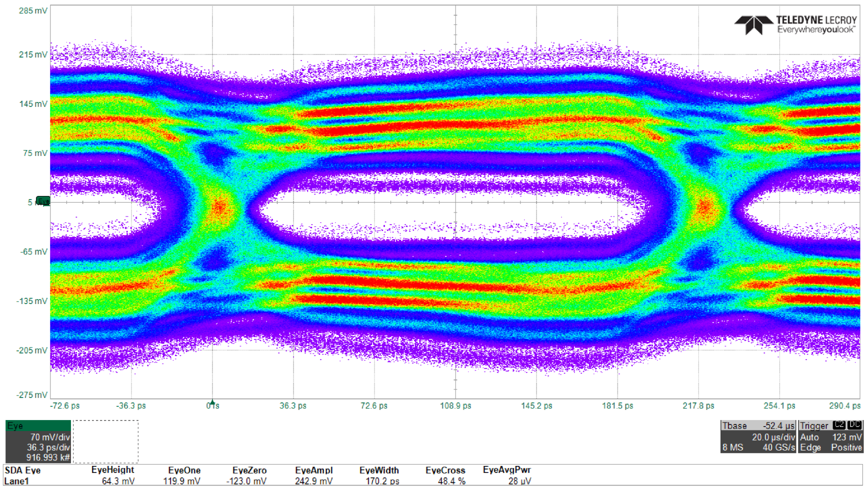

SDA Eye	EyeHeight	EyeOne	EyeZero	EyeAmpl	EyeWidth	EyeCross	EyeAvgPwr
Lane1	64.3 mV	119.9 mV	-123.0 mV	242.9 mV	170.2 ps	48.4 %	28 µV

Figure 4.8: Measured eye diagram at the output of the transmitter with nine active replicas for a supply voltage of 1.05V on both power domains when all 16 FFE levels are activated. FFE weights are $w_0 = 0.56255$, $w_1 = -0.125$, $w_2 = -0.125$, $w_3 = -0.0625$, $w_4 = -0.0625$ and $w_5 = -0.0625$. The transmission speed is 4.6Gbps. The levels show good linearity of the transmitter output voltage.

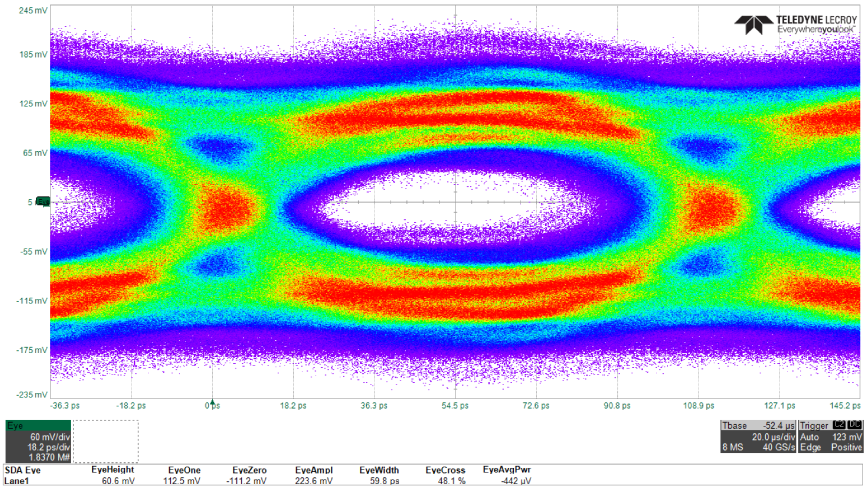

SDA Eye	EyeHeight	EyeOne	EyeZero	EyeAmpl	EyeWidth	EyeCross	EyeAvgPwr
Lane1	60.6 mV	112.5 mV	-111.2 mV	223.6 mV	59.8 ps	48.1 %	-442 µV

Figure 4.9: Measured eye diagram at the output of the transmitter with nine active replicas for a supply voltage of 1.05V on both power domains when all 16 FFE levels are activated. FFE weights are $w_0 = 0.56255$, $w_1 = -0.125$, $w_2 = -0.125$, $w_3 = -0.0625$, $w_4 = -0.0625$ and $w_5 = -0.0625$. The transmission speed is 9.2Gbps. The levels show good linearity of the transmitter output voltage.

After that, our analysis focused on the dependence of the eye diagram parameters on the supply voltage. Fig. 4.10 and fig. 4.11 show the horizontal and vertical opening of the eye diagram at a transmission speed of 9.23Gbps, respectively. Both cases where the driver is under the LDO supply (blue curve) and directly connected to V_{DD} (red curve) are shown. Fig. 4.10 shows that the width of the eye diagram with and without LDO are comparable, thus the LDO is not giving us a big help eliminating jitter: this is a direct consequence of the fact that the supply is now well decoupled from the effect of parasitic inductances, therefore the PSR action of the LDO at high frequency is not so relevant. As can be seen in Fig. 4.11, as expected the vertical opening of the eye diagram greatly improve when bypassing the LDO. We underline that the driver doesn't reach full swing when connected directly to V_{DD} because of the NN slice: in fact, the pull-up network in this case should be done with a P-MOS. Using a NN slice, the N-MOS forming the pull-up network is working in saturation and this greatly reduces the swing of the driver, limiting it from reaching rail-to-rail operation. For the same reason of having a NN slice, increasing V_{DD} the vertical opening increases drastically, confirming that the pull-up N-MOS operates in saturation region.

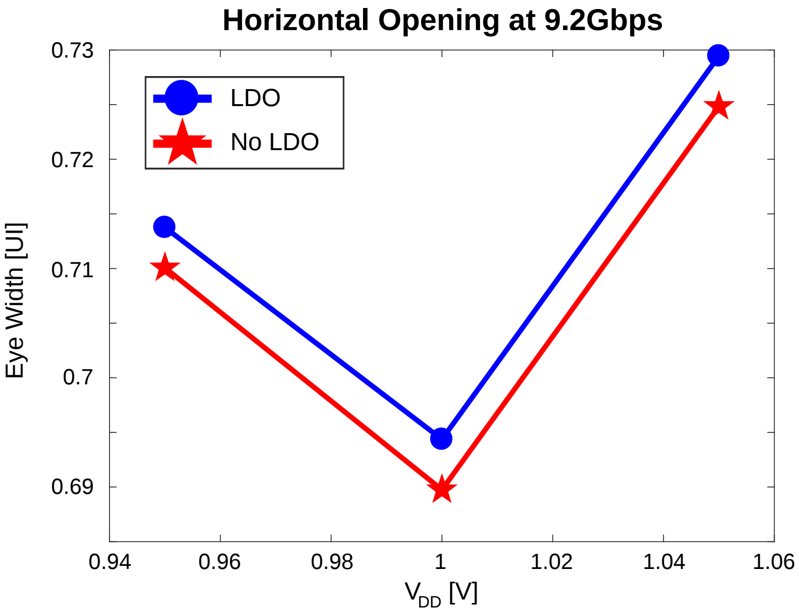

Figure 4.10: Measured normalized horizontal opening at the transmitter output with nine active replicas for different values of V_{DD}. The transmission speed is 9.23Gbps. FFE has not been activated. Both cases where the driver is under the LDO or directly connected to V_{DD} are shown.

Fig. 4.12 shows the height-to-amplitude ratio for the same 9.23Gbps transmission speed versus V_{DD}. The amplitude is defined as the difference between the mean of the '1' level and the mean of '0' level, therefore this figure of merit is indicating how clean the levels '1' and '0' are visible: a high height-to-amplitude ratio indicates a better SNR. As expected, this doesn't change with the power supply in the case with LDO, where it improves with increasing V_{DD} when the driver is directly connected to the supply due to the increase in vertical opening. Finally, the power consumption at 9.23Gbps for different V_{DD} is shown in Fig. 4.13. The higher power consumption for the case without the LDO is due to the increase in the driver power consumption when it is directly connected to V_{DD}.

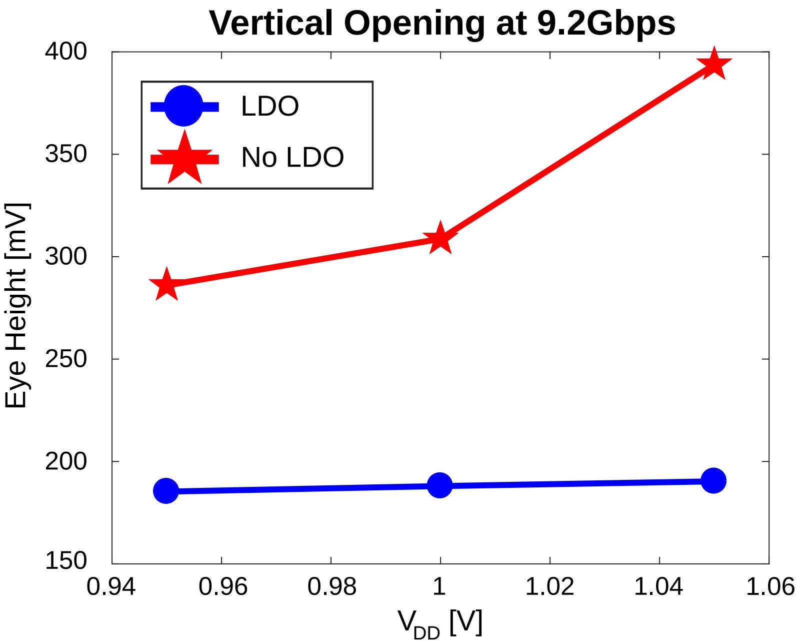

Figure 4.11: Measured vertical opening at the transmitter output with nine active replicas for different values of V_{DD}. The transmission speed is 9.23Gbps. FFE has not been activated. Both cases where the driver is under the LDO or directly connected to V_{DD} are shown.

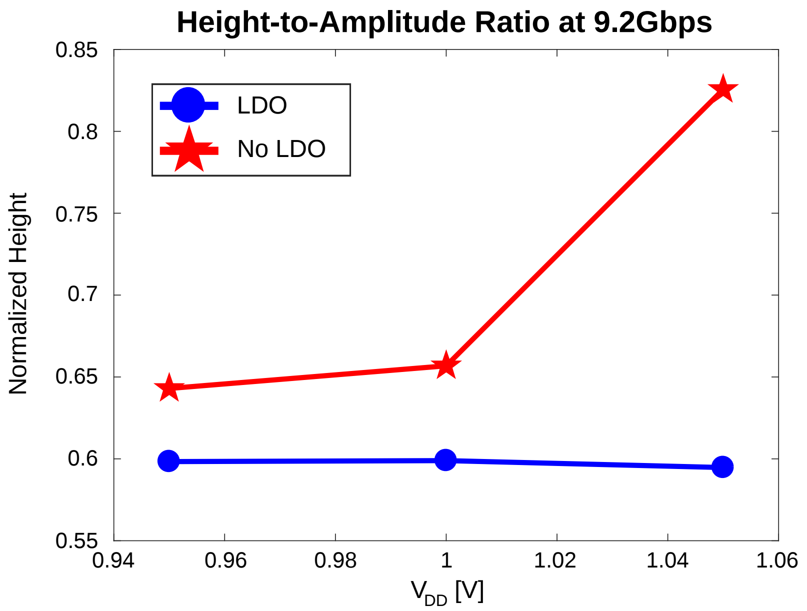

Figure 4.12: Measured height-to-amplitude ratio at the transmitter output with nine active replicas for different values of V_{DD}. The amplitude is defined as the difference between the mean of the ′1′ level and the mean of ′0′ level. The transmission speed is 9.23Gbps. FFE has not been activated. Both cases where the driver is under the LDO or directly connected to V_{DD} are shown.

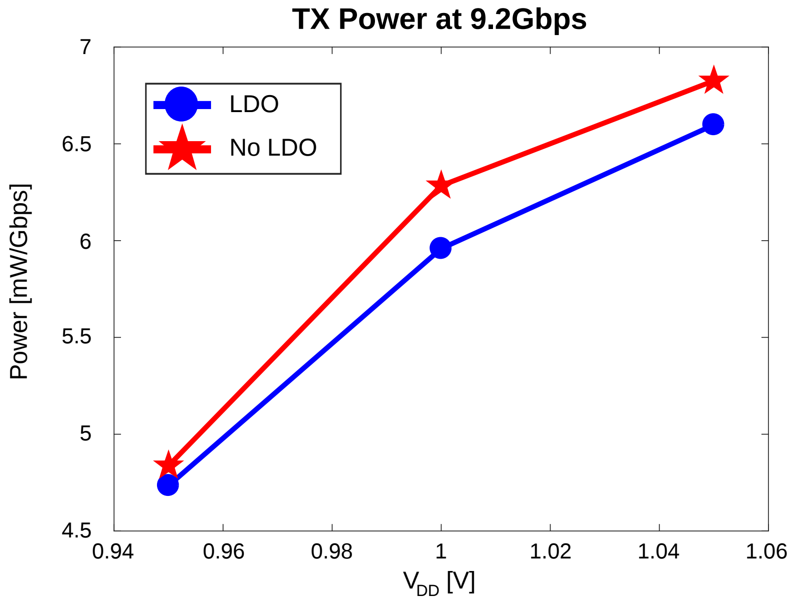

Figure 4.13: Measured normalized power consumed by the transmitter with nine active replicas for different values of V_{DD}. The transmission speed is 9.23Gbps. Both cases where the driver is under the LDO or directly connected to V_{DD} are shown.

The next step of the transmitter characterization has been the analysis of the effect on the eye diagram parameters of activating a different number of replicas. The results of such analysis are shown in Figg. 4.14-4.16. Fig. 4.14 indicates that, as expected, there is no strong dependence of the eye width on the number of active replicas. The same thing doesn't hold for the vertical opening: in fact, fig. 4.15 shows that the eye height increases when increasing the number of active replicas. Also this effect is expected, since by activating more replicas we are lowering the equivalent output resistance of the driver, so the current on the driver side of the resistive divider increases, therefore the transmitter output voltage increases. By increasing the current drawn by the driver, we are also increasing the power consumed by the transmitter, which is what we see in fig. 4.16.

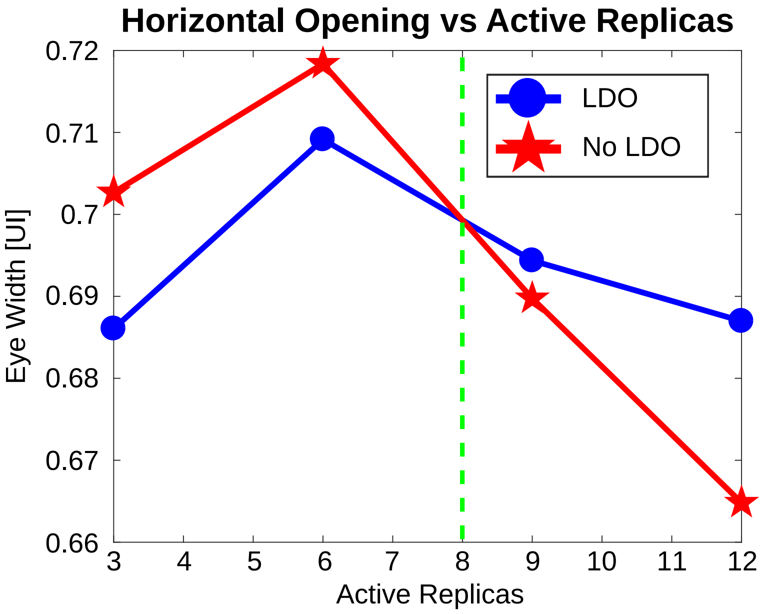

Figure 4.14: Measured normalized horizontal opening at the transmitter output for different number of active replicas for V_{DD}=1V. The transmission speed is 9.23Gbps. FFE has not been activated. Both cases where the driver is under the LDO or directly connected to V_{DD} are shown. The green dashed line indicates the number of replicas that should be activated in order to match 50Ω under nominal conditions (typical-typical technology corner, 900mV chip supply, 25°C).

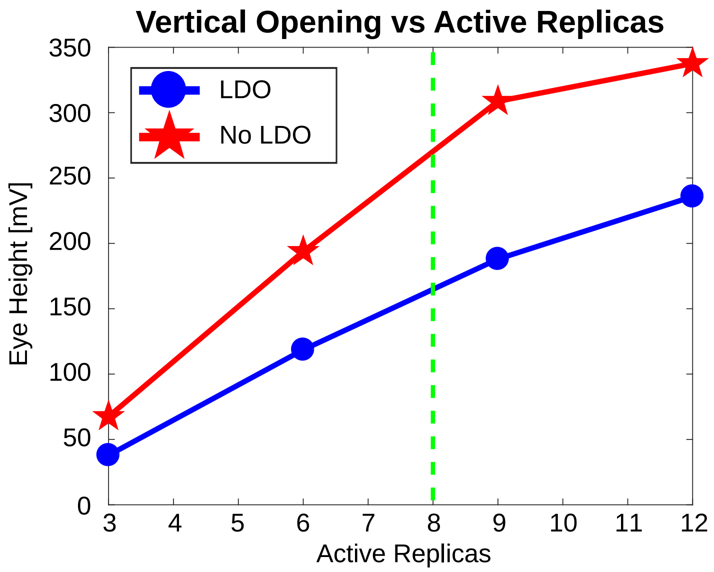

Figure 4.15: Measured vertical opening at the transmitter output for different number of active replicas for V_{DD}=1V. The transmission speed is 9.23Gbps. FFE has not been activated. Both cases where the driver is under the LDO or directly connected to V_{DD} are shown. The green dashed line indicates the number of replicas that should be activated in order to match 50Ω under nominal conditions.

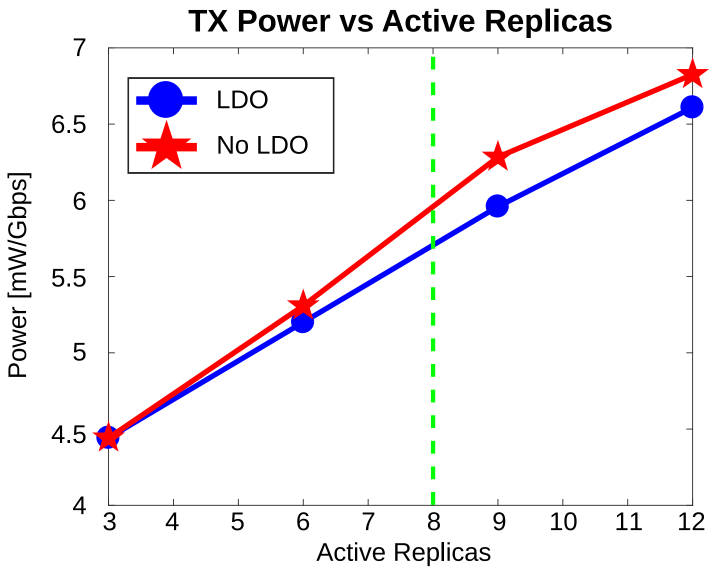

Figure 4.16: Measured normalized power consumed by the transmitter for different number of active replicas for V_{DD}=1V. The transmission speed is 9.23Gbps. Both cases where the driver is under the LDO or directly connected to V_{DD} are shown. The green dashed line indicates the number of replicas that should be activated in order to match 50Ω under nominal conditions.

Finally, figg. 4.17-4.19 show the degradation of the eye diagram parameters when increasing the transmission speed. In order to perform these measurements, the clock provided to the transmitter has been divided by two and by four, so to obtain transmission speeds of 4.6Gbps and 2.3Gbps, respectively. Fig. 4.17 and fig. 4.18 show how both horizontal and vertical opening improve when reducing the transmission speed. The fact that the horizontal width improves is quite trivial, whereas the fact that the vertical opening improves is not so obvious. We believe that this decrease in vertical opening when the transmission speed is increased is given by the IR drop on the power distribution lines: in fact, when increasing the transmission speed, also the current drawn from the supply increases, bringing to a reduced effective supply as seen by the chip. This hypothesis is also supported by the fact that this decrease in vertical opening is more pronounced when the driver is directly connected to V_{DD}, which increases the amplitude of the spikes of the current drawn by the driver.

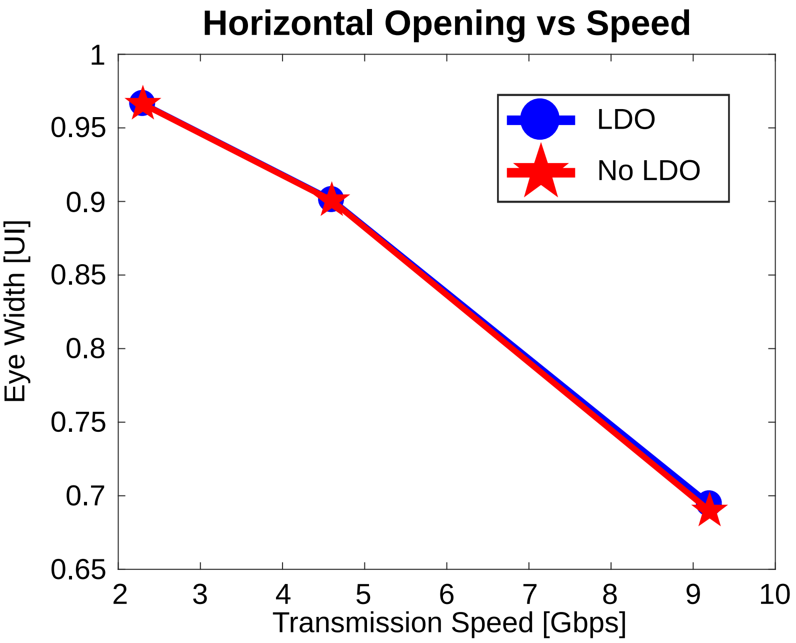

Figure 4.17: Measured normalized horizontal opening at the transmitter output with nine active replicas for different transmission speed for V_{DD}=1V. FFE has not been activated. Both cases where the driver is under the LDO or directly connected to V_{DD} are shown.

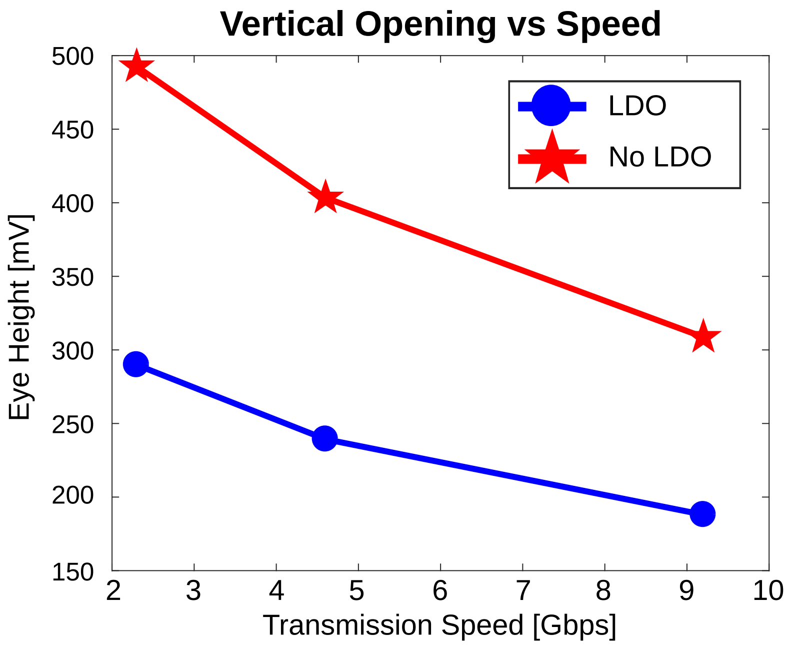

Figure 4.18: Measured vertical opening at the transmitter output with nine active replicas for different transmission speed for V_{DD}=1V. FFE has not been activated. Both cases where the driver is under the LDO or directly connected to V_{DD} are shown.

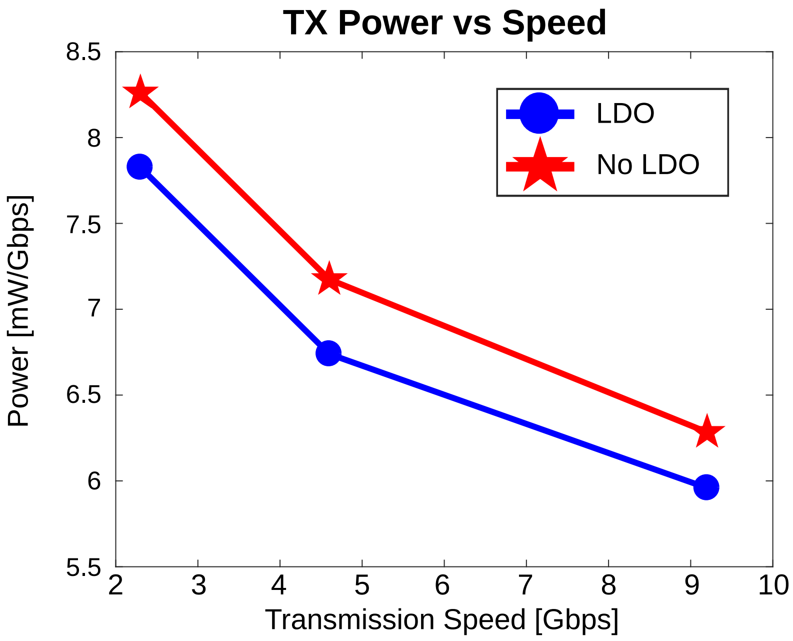

Figure 4.19: Measured normalized power consumed by the transmitter with nine active replicas for different transmission speed for V_{DD}=1V. Both cases where the driver is under the LDO or directly connected to V_{DD} are shown.

Chapter 5

System and Transistor Level Design of the Receiver

As we have already said, channels with different characteristics, including ones with many connectors introducing notches in the S_{21}, may be employed in automotive environment, which demand for advanced equalization strategies with high reconfigurability [57]-[27]-[58]. As a consequence, linear equalization as IIR in the receiver must be combined with DFE [35], especially at high data rates.

The typical HSSI receiver architecture is shown in Fig. 5.1. DFE is employed as non-linear equalization technique, where the information about the previous reconstructed bits is used to mitigate the effects of ISI on the currently received bit. The feedback loop relative to the first post-cursor (h_1 in Fig.5.1) has very stringent timing constraints, so that speculative DFE (loop-unrolling) is often employed at high data rates [33]-[29]-[31]. The architecture of a full-rate receiver with loop-unrolled DFE is shown in Fig. 5.2.

The receiver in Fig.5.1 works at full-rate (i.e. clock frequency equal to the data-rate). Half-rate clocks are required to relax timing [59], save power and simplify the latches. The architecture of a very basic half-rate receiver without DFE is shown in Fig. 5.3. Combining loop-unrolling DFE, CDR and half-rate architecture at high-data rates is challenging. In particular, DFE opens the received eye diagram at the sampling points. However, it has detrimental effects on the edges distribution of the partial response signals, making difficult the use of bang-bang digital CDR [60]. Edges can be sampled before the DFE summer, however requiring enhanced samplers' sensitivity and transition filtering, that reduces the bandwidth of the CDR [47].

In this chapter, we propose a new design for a half-rate receiver with 4-tap DFE (one speculative) for HSSI at 10Gb/s for automotive applications which is able to employ the edge samples from the speculative DFE paths to adapt the CDR while still keeping high CDR bandwidth.

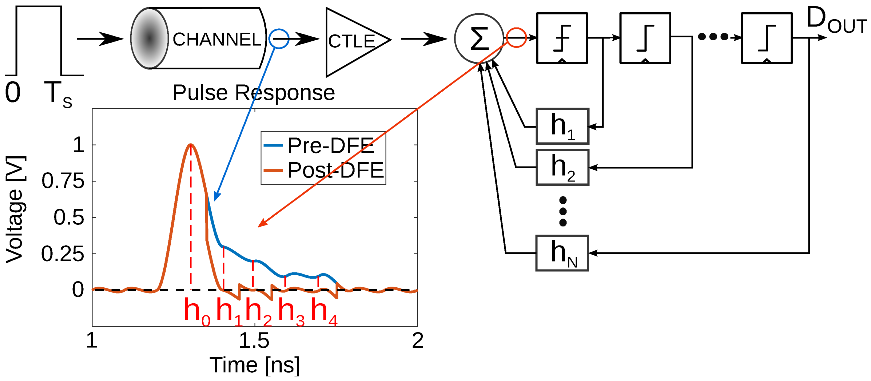

Figure 5.1: Sketch of a full-rate receiver with DFE. A pictorial example of pulse response before and after DFE is reported at the bottom. We subtract to the received (analog) signal the value of the previously received bits multiplied by suitable coefficients (tap weights) h_i.

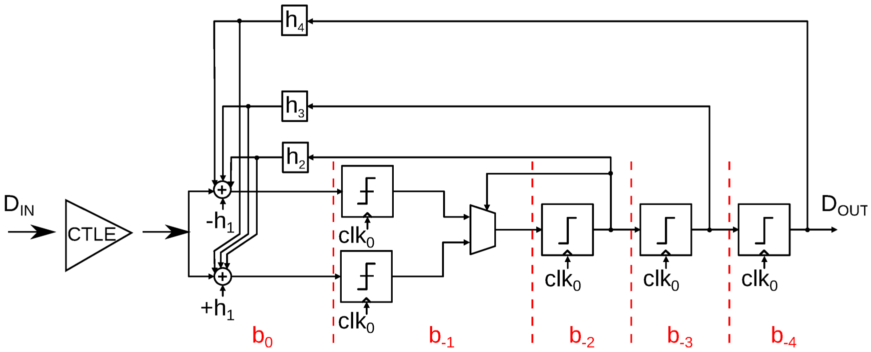

Figure 5.2: Architecture of a full-rate receiver with 4-taps, 1-tap loop-unrolled DFE.

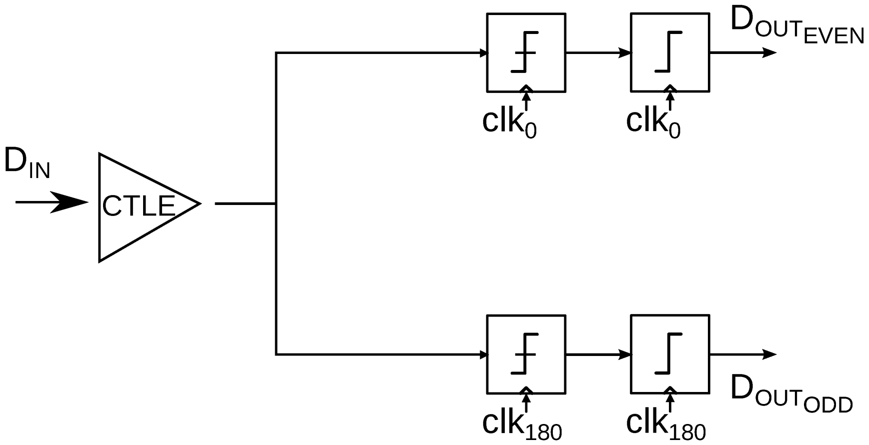

Figure 5.3: Architecture of a half-rate receiver without DFE.

5.1 Structure of the receiver

In this section we will motivate the general choices about the structure of the receiver. The transistor level design will be then analyzed in detail in a later section of this chapter.

5.1.1 Input Amplifying Stage

The first block that we meet in the receiver is the input amplifying stage, which can be just as simple as a single stage CTLE or a much more complicated circuit. CTLE is present in almost all high-speed transceivers in literature: a good CTLE should feature the possibility of having a good gain at Nyquist frequencies with respect to low frequencies, low input referred noise and low DC offset [61]. However, some works [62]- [59] claim that CTLE amplifies also HF noise and crosstalk in addition of extra-power needed, resulting in both being amplified with respect to the eye diagram opening. For this reason, they normally have a DFE-IIR doing the same job (or even having more than one IIR [63]): however, we have decided that the amount of required CTLE in our case is not such to justify such a non-conventional solution.

In front of the first amplifying block, an AC coupling is needed in order to set the bias point independently from the one of the received data [64]- [58]. However, long sequences of 1s or 0s at the input of the receiver may induce a transient shift of the low-frequency part of the received signal after the AC coupling itself, which in turn becomes reduced sensitivity for the receiver. In order to avoid this problem, many circuital solutions have been proposed [64]- [58] [65], but we decided to solve it by encoding the transmitted data with the already introduced "8b10b encoding" [17], so that the maximum length of a sequence of identical bits in a row is five. Along with this, an on-board AC coupling capacitance of 25nF has been used.

A conventional solution for an amplifying stage is proposed by [29] and is sketched in Fig. 5.4. After a T-coil peaking, a parallel amplifier for a wide range of input signal levels and a two-stage peaking amplifier, described more in detail in Fig. 5.5, are employed. A similar solution can be found also in [66].

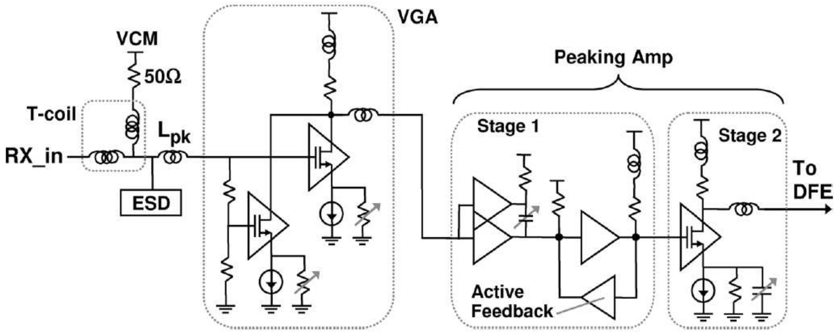

Figure 5.4: Input Amplifying Stage as presented in [29].

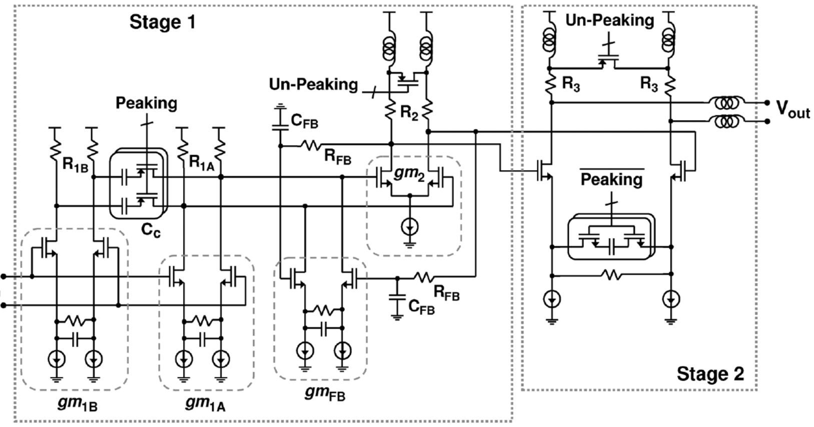

Figure 5.5: Two-stages peaking amplifier used in [29]. The first stage presents capacitively coupled parallel stages, the second has zero-peaked topology with switched cap degeneration. In the first stage, RC source degeneration for improved linearity and bandwidth extension is present.

In literature, many more complicated solutions are presented: [67] and [68] propose a CTLE that demuxes the full-rate data stream to following stages working at half-rate; [69] presents a source-degenerated CTLE followed by a 5-stages limiting amplifier; [52] shows a CTLE formed by a high-frequency equalizer, formed by a source-degenerated differential pair, and a low-frequency equalizer with a shunt inductor, which is a not suitable solution for an automotive product; finally, some other works [70]-[71] use active peaking as a mean of adding a programmable boosting high-frequency gain after a degenerated-source differential pair.

Anyway, all these proposed solutions are meant to be used for much higher speed than what we require, and are also more power hungry and less robust. For this reason, we decided to implement a single-stage RC source-degenerated differential pair CTLE, similar to the first peaking stage of [29] and [66].

5.1.2 Half-Rate, CDR and DFE

As said both in Chapter 1 and in the introduction to this chapter, at high-speed it's almost impossible to analyze separately the following three aspects of a receiver: whether it should be full-rate or half-rate, its CDR scheme and the DFE.

A very basic idea of how a full-rate and a half-rate receiver work is sketched in Fig. 5.6. The main difference between the two is that in the full-rate one, a full-rate clock (in our case, 10GHz) is used to sample both data and edges, whereas in a half-rate architecture two different half-rate clocks (in our case, 5GHz) shifted in phase of 90°, are used, one for sampling the data, the other one for the edges. These two clocks are usually called clock I and clock Q.

Even with a deep investigation of the literature, it's impossible to find a 10Gbps full-rate receiver architecture, because the analog path of the circuit would be too power hungry and it would be extremely difficult to cope with

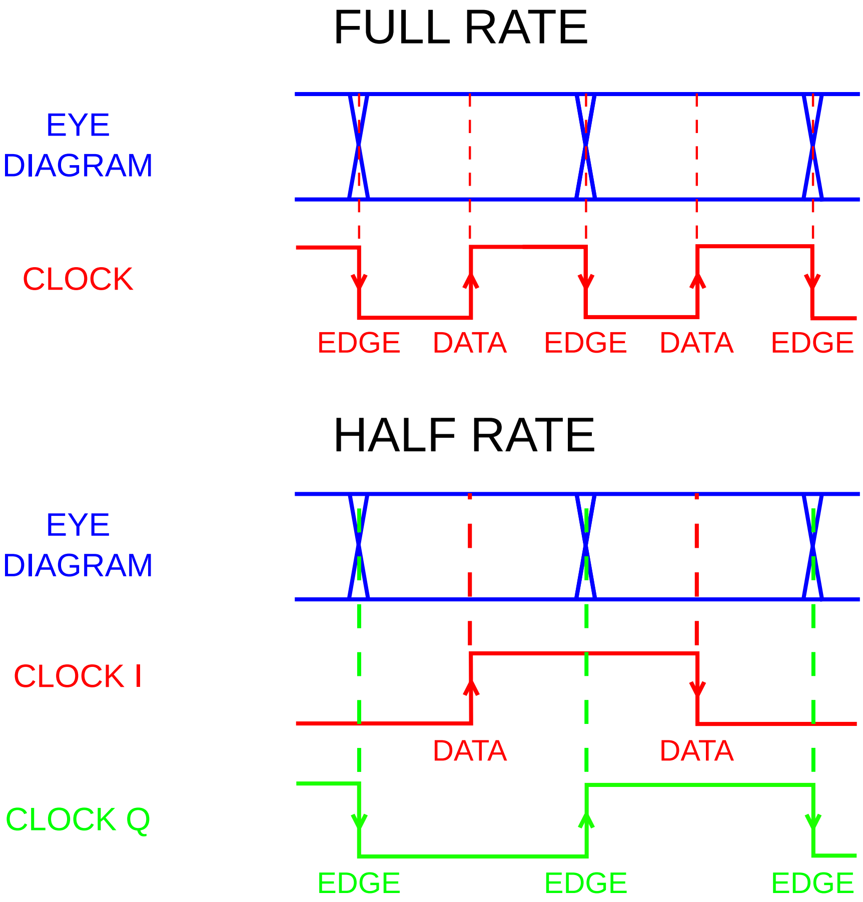

Figure 5.6: Sketch highlighting the main difference in terms of clocking scheme between full-rate and half-rate receivers. The dotted lines show the data and edge sampling instants.

the timing of the first DFE tap. For these reasons, people started investigating on half-rate [72], [73], [62], [59] and quarter-rate [29], [30], [74] architectures. The first ones seem anyway more suited to our target speed, avoiding also an unnecessary complexity overhead.

Regarding the CDR, the literature showing the circuital implementation is in fact very old and the focus nowadays is completely on the algorithm side [52]. There are very few paper describing the transistor level design of the CDR [75], [64] and they don't vary that much one from another: a good overview has already been given in chapter one. An in-depth analysis of the state-of-the-art CDR algorithms and our proposed one will follow in the next section.

Finally, concerning DFE, we decide to avoid using IIR [76] because it's more suited for long silicon channels as said in the introduction. Our idea would be to share the equalization duties between transmitter and receiver, using for certain taps FFE, for some others DFE, as already done by [31], [33]. At 10Gbps the speed forces us to use the loop-unrolling technique, which we already introduced before. There are works that unroll up to three taps [61], [29], [66], [31], but this introduces a significant overhead in both area and power consumption, so to make this option not too attractive for our speed target. We decided to go with a more-traditional 1-tap loop-unrolled DFE [58], [74], [77].

Finally, the summers have been realized with a traditional current-mode architecture [74], which we will analyze later on when describing the transistor level design.

5.1.3 Architecture of the Receiver

The block diagram of the DFE and sampling portion of the proposed receiver is sketched in Fig.5.7. For simplicity, only the even path is shown. Loop-unrolling requires two summers per path. The tap h_1 is speculative and the selection between $+h_1$ and $-h_1$ is decided by the output of the other path. Feedback loops and shift registers within the same path are used for h_2 and h_4, while h_3 requires crossing between even and odd path. We limit the DFE to 4 taps based on the system level analysis carried out in [78] showing that, for channels typical for board-to-board communication in automotive, the pulse response does not extend significantly outside the 4th post-cursor, as shown also in Figs. 2.13-3.14b).

At the bottom of Fig.5.7 we can see the additional latches needed for edge sampling, to be input to CDR algorithm together with the data samples. Notice that, as we will detail in the next section, we consider the edges at the output of all summers (4 in total). A four phases, half-rate (i.e. 5GHz) clock samples data and edges. Fig. 5.8 shows how these 5GHz four-phases are generated. A digital PLL outputs a 10GHz into a clock divider, which outputs a four-phases, 5GHz clock. These four phases are fed into two phase interpolators, one for clock I phases, the other one for clock Q phases. Both are controlled by the CDR and align their output clock edges with the data stream.

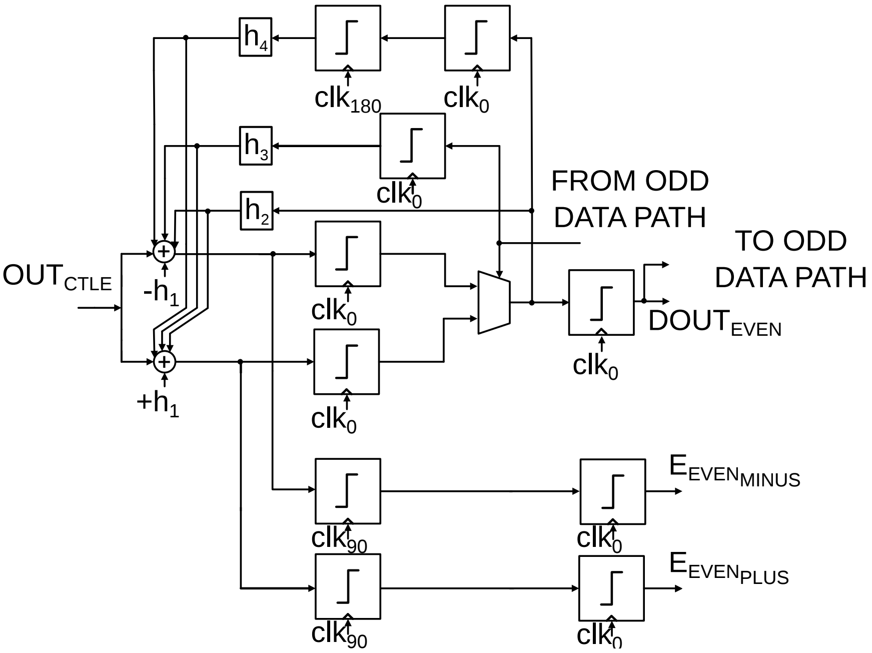

Figure 5.7: Sketch of one of the two paths (in this case the even) of the half-rate receiver using loop-unrolled DFE. To the far left, the data coming out from a CTLE are the input to this part of the circuit. The sampled data and edges are then sent to the deserializer.

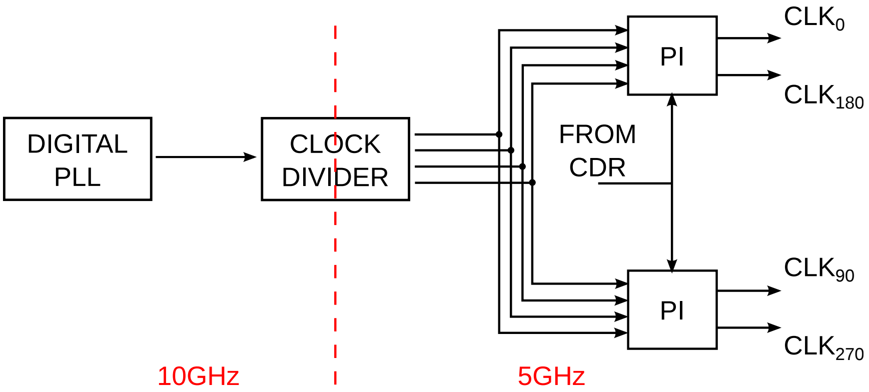

Figure 5.8: Sketch of the clocking scheme of the receiver. A digital PLL outputs a 10GHz clock into a clock divider: from there on, only 5GHz clocks are distributed.

5.2 CDR Algorithm

We consider digital CDR with transition-based phase detection. In the presence of ISI this requires transition filtering [60]. If the channel introduces significant ISI only on four post-cursors or less, as in our case, this would imply discarding a lot of bit sequences (only the pattern 101010 is useful), resulting in very low CDR bandwidth. This is due to the fact that a transition-based phase detection needs transitions that are symmetric around 0V in order to work correctly, and this is not the case if the channel introduces ISI.

To avoid interactions between DFE and CDR and leverage exact transition filtering, one should sample the edge transitions before the DFE summer [79], however increasing the load capacitance of the first stage of the receiver (the continuous-time-linear-equalizer) and thus in turn reducing its bandwidth. Not surprisingly, many algorithms have been proposed for CDR in the presence of DFE to reduce the need for transition filtering [60]-[52]-[80]-[79], but they rely on adaptive techniques, whereas the focus of this work is to propose a phase detection algorithm for CDR which is stable also when (simpler) non-adaptive DFE is used.

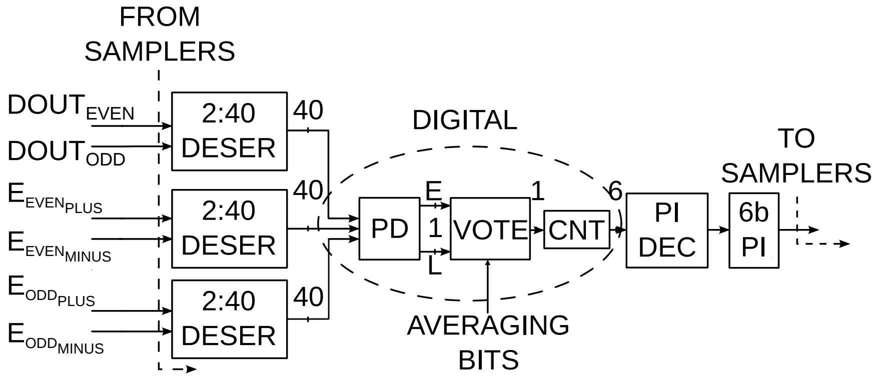

Figure 5.9: Scheme of the CDR circuit. Three packets of 40 bits enter the digital state machine: even and odd data, as well as even and edges for all speculative paths. The CDR circuit controls a phase-interpolator through a PI counter and a PI decoder.

The block diagram of the CDR is sketched in Fig.5.9, whereas the state diagram of the phase detection (PD) algorithm is in Fig.5.10. The first step of the PD is to find a 3-bit clock pattern (101 or 010) in a packet of 40 bit (four bytes of information encoded with a 8b10b encoding). In the selected pattern we analyze the second transition (e.g. 1 $\rightarrow$ 0 in the 010 pattern, see Fig.5.11), where there is a symmetry between the speculative paths (+h_1 and -h_1) at the edge sampling instant compared to the first transition (e.g. the 0 $\rightarrow$ 1 in the 010 pattern, compare the horizontal arrows in Fig.5.11). If the sampling point is optimally set, the +h_1 edge sampler will sample a 1, the -h_1 edge sampler will sample a 0 (Fig. 5.12 a)). In this case, the phase detector outputs (*early* or *late*) will both be 0. If both edge samples are equal to the second bit of the 3-bits clock pattern, then the early output of the PD will go high (Fig. 5.12 b)). If both edge samples are equal to the third bit of the 3-bits clock pattern, then the late output of the PD will go high (Fig. 5.12 c)).

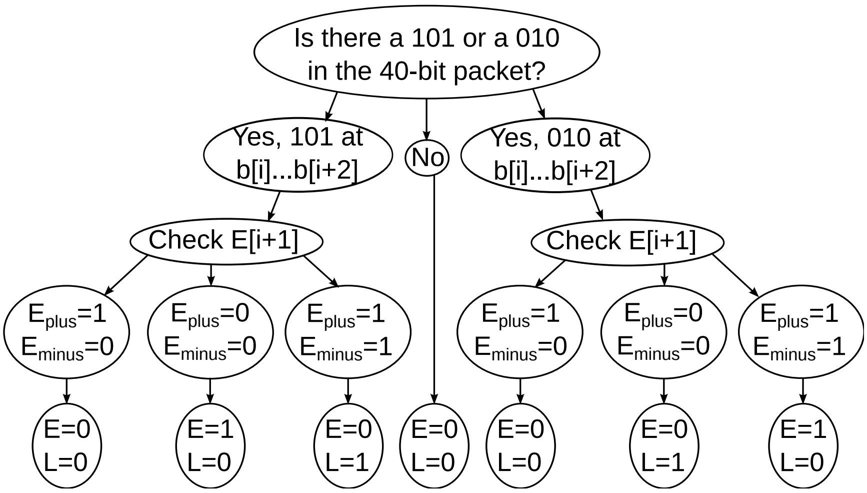
Is there a 101 or a 010 in the 40-bit packet?
Yes, 101 at b[i]...b[i+2]
No
Yes, 010 at b[i]...b[i+2]
Check E[i+1]
Check E[i+1]
E_{plus}=1 E_{minus}=0
E_{plus}=0 E_{minus}=0
E_{plus}=1 E_{minus}=1
E_{plus}=1 E_{minus}=0
E_{plus}=0 E_{minus}=0
E_{plus}=1 E_{minus}=1
E=0 L=0
E=1 L=0
E=0 L=1
E=0 L=0
E=0 L=0
E=0 L=1
E=1 L=0

Figure 5.10: Phase detection algorithm as implemented in this work.

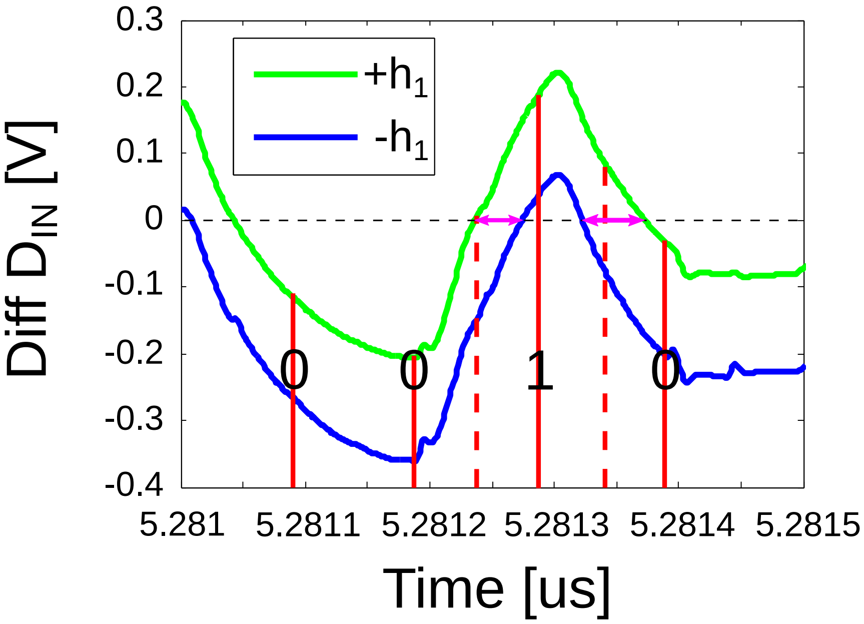

Figure 5.11: Results of circuit simulations (see Section 5.3). The differential outputs of the two summers in the even path, which is also the differential input of the samplers in the even data path, are shown, along with the data (full lines) and edge (dashed lines) sampling instants.

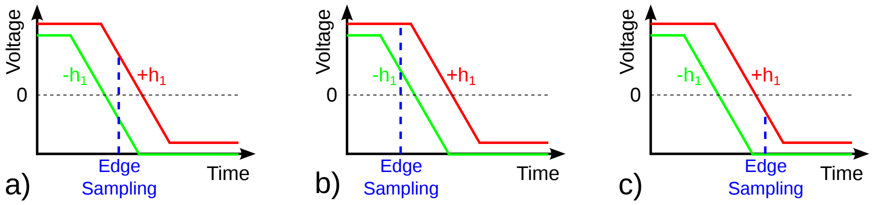

Figure 5.12: Position of the edge sampling point with respect to speculative data transitions $+h_1$ and $-h_1$. a)$+h_1$ edge sampler samples a 1, $-h_1$ a zero: the sampling point is optimally set; b)$+h_1$ edge sampler samples a 1, $-h_1$ too: the sampling point is early; c) $+h_0$ edge sampler samples a 0, $-h_1$ a one: the sampling point is late.

The output of the PD is not directly fed into the PI counter. An early/late decision is taken after the averaging: the voting block does a majority voting on *averaging* decisions of the PD, and the output to the PI counter increases by 1 if *Early* was the majority, -1 if there were more *Late* and 0 otherwise. Averaging is performed for two main reasons. First, because it is not guaranteed that a 3-bit clock pattern is present in a 40-bits packet, so without averaging the CDR loop may be open for some time, and CDR bandwidth would be undefined. Secondly, because, since the zero-crossing time depends on the random data pattern preceding the three-bits (more than one post-cursor may be relevant) filtered transition, the PD algorithm cannot give the correct results all the times.

In the frame of board-to-board communication with two different crystals used for transmitter and receiver, averaging (n_{avg}) and crystal quality (precision ϵ) can be related as:

$$2 \cdot \epsilon \cdot t_S < \frac{t_S}{n_{avg} \cdot n_{PIsteps} \cdot n_{bitsperpacket}}$$

t_S being the bit period, from where we find

$$n_{avg} < \frac{1}{2 \cdot \epsilon \cdot n_{PIsteps} \cdot n_{bitsperpacket}} \tag{5.1}$$

where in our case $n_{bitsperpacket}$=40 and the number of steps of the PI is 64. At the same time we can compute the CDR bandwidth as

$$BW_{CDR} \simeq \frac{f_{CK}}{n_{avg} n_{bitsperpacket}} \tag{5.2}$$

Fig. 5.13 is obtained directly from Eqs. (5.1)-(5.2) and shows that with reasonable crystal quality, very low number of averages is allowed (3 for ϵ =100ppm). In this respect, there is limited room for incorrect results coming out of the PD.

Finally, in order to get a better understanding of the mix between half-rate architecture, CDR and DFE, a Simulink model of the whole receiver has been implemented before designing the actual transistor level. A screen-shot of the whole Simulink setup is shown in Fig. 5.14. The functioning of the structure has been validated with a realistic input data sequence, in which ISI of a typical

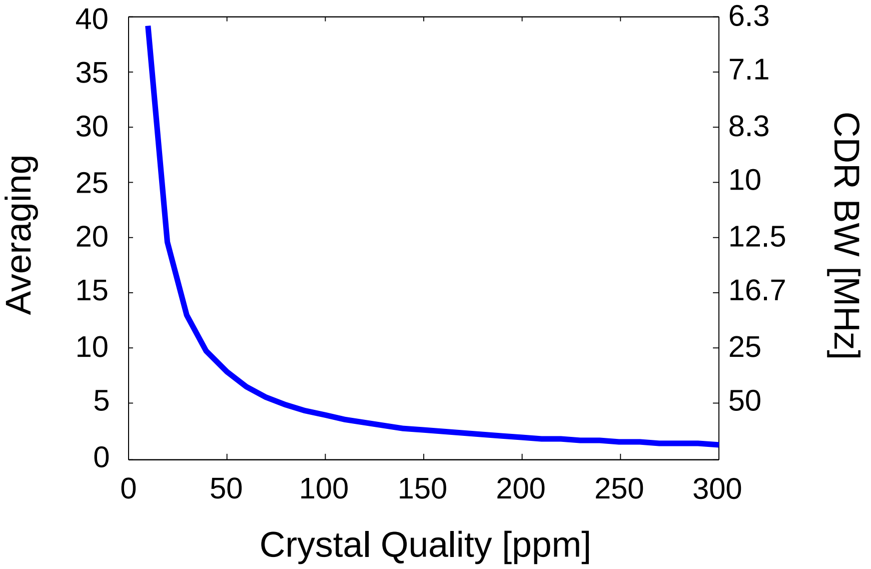

Figure 5.13: CDR Bandwidth vs quality of the crystal. The left axis reports the number of 40-bit packets to be elaborated before taking a decision, therefore giving the CDR bandwidth reported on the right axis.

automotive channel was accounted, showing the correct settling of the PI code and the error-less reconstruction of the received data (see Fig. 5.15).

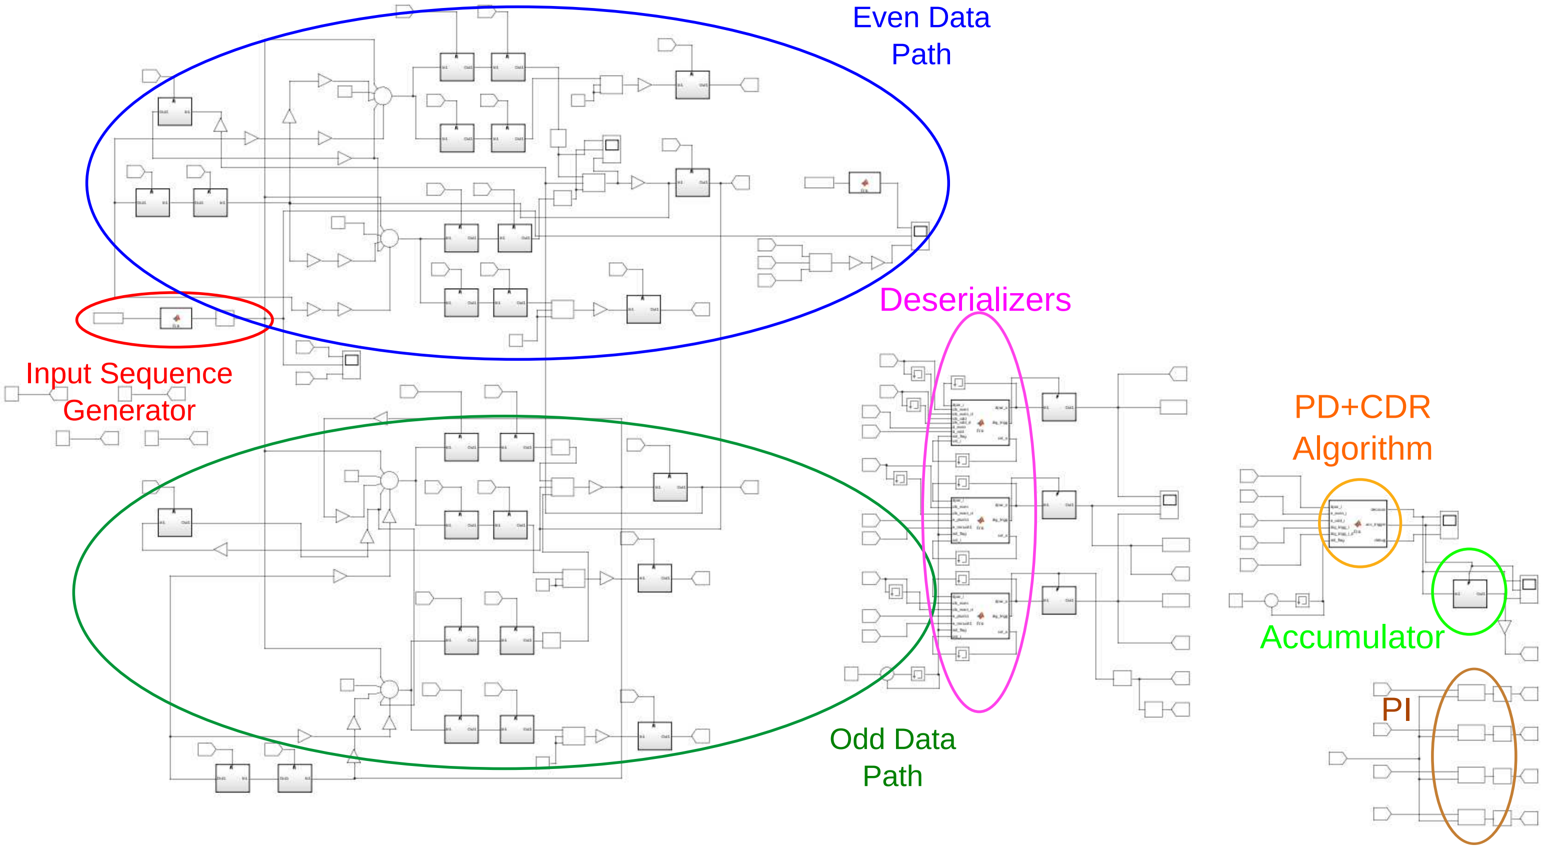

Figure 5.14: Receiver Structure as implemented in Simulink. From the left to the right we find the following blocks: red box, input sequence generator; blue box: even data path; dark green box: odd data path; pink box: deserializers; orange box: phase detector and CDR digital algorithm; light green box: accumulator; brown box: phase interpolator.

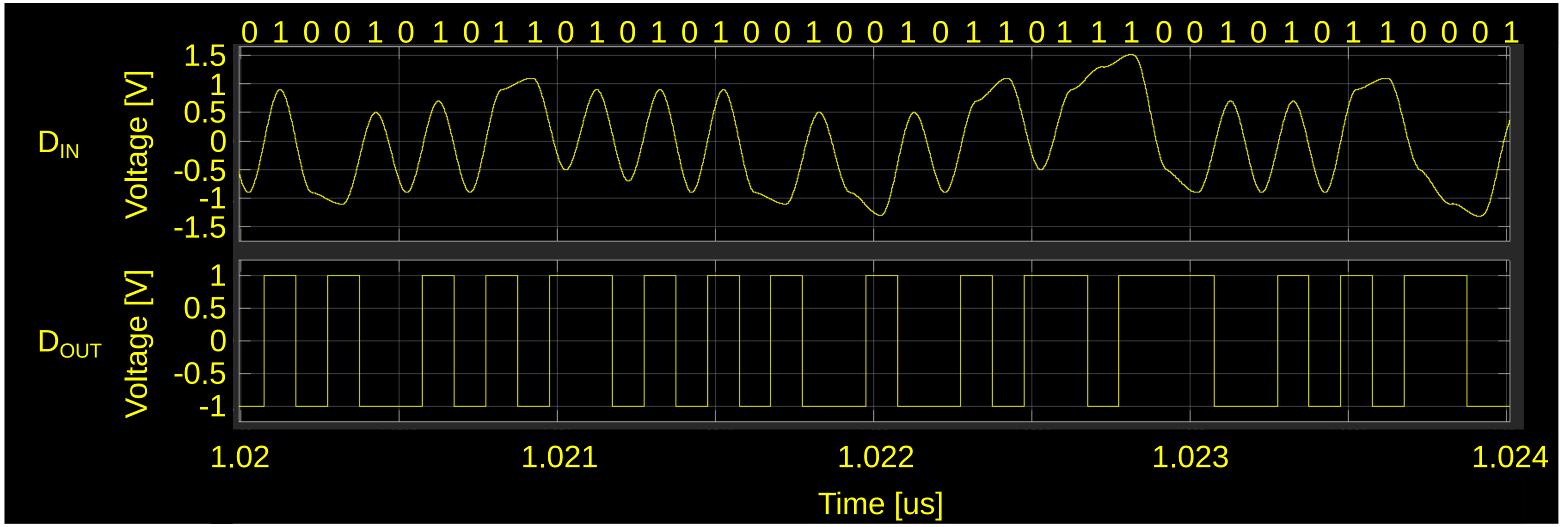

Figure 5.15: Data reconstruction for the Simulink model of Fig. 5.14. In the above panel, the input data D$_{IN}$ at the receiver, affected by ISI, are shown, whereas in the bottom panel there are the reconstructed data at the receiver output, D$_{OUT}$. The figure shows an error-less reconstruction.

5.3 Transistor level design and simulation results

In the following section, details on the transistor level implementation of the single blocks of the receiver will be given. The technology used is the same 28nm planar CMOS technology used for the transmitter. To clarify better the structure as explained until now, the timing diagram of the data and edge samples (so basically data and edge samples path as from Fig. 5.7, but for both odd and even path) along the path are shown in Fig. 5.16.

5.3.1 Input Stage

The structure of the input stage is reported in Fig. 5.17.

The operating point of the CTLE is set via a resistive divider. The value of the divider is trimmable in eight values, setting the biasing point from a minimum of $\frac{V_{DD}}{2}$ to a maximum of $\frac{2V_{DD}}{3}$, value that can be set depending on the differential peak-to-peak opening of the transmitted signal.

The DC biasing resistive divider is followed by a CTLE, realized as a RC source-degenerated differential pair [66]. The peak frequency and the boost gain can be tuned by means of a 3-bits trimmable capacitance and a 2-bits trimmable resistance, so to compensate for PVT variations. Along with these settings, also the current is 4-bits trimmable. For channels with small attenuation, it is possible to switch off the CTLE and have it working as a conventional differential pair. Since this equalization scheme does not aim to equalize channels having a very big loss (not expected for a Nyquist frequency of 5GHz), our main goal was not to have a high boost at high-frequency, whereas to have very similar transfer function for the CTLE over all PVT corners. The transfer function of the CTLE for PVT corners is reported in Fig. 5.18. The variations in the transfer function depending on the trimming bits are reported in Fig. 5.19.

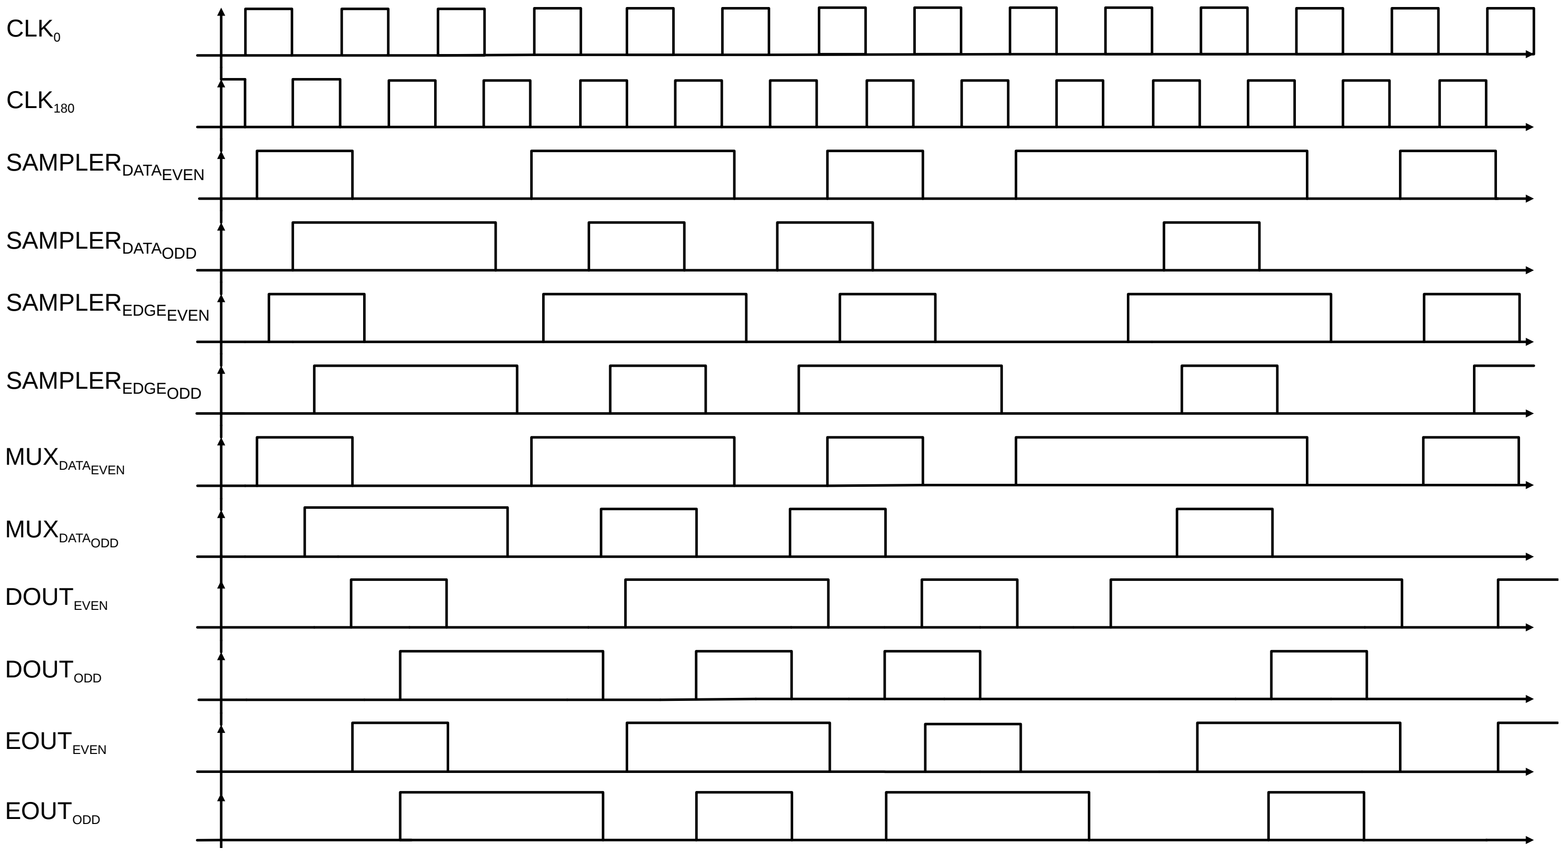

Figure 5.16: Timing for the half-rate receiver of which only the even path was shown in Fig. 5.7. For clarity, here only one sample between plus and minus has been shown. SAMPLER$_{DATA_{EVEN}}$, SAMPLER$_{DATA_{ODD}}$, SAMPLER$_{EDGE_{EVEN}}$, SAMPLER$_{EDGE_{ODD}}$ are the output after the first samplers the data and edge meet in Fig. 5.7. MUX$_{DATA_{EVEN}}$ and MUX$_{DATA_{ODD}}$ are the data samples after the two speculative samples have passed the mux. D$_{OUT_{EVEN}}$ and D$_{OUT_{ODD}}$ are the data after the final data sampler, ready to be sent to the deserializers of Fig. 5.9, and same goes for E$_{OUT_{EVEN}}$ and E$_{OUT_{ODD}}$, where again only one between the speculative minus and plus has been shown.

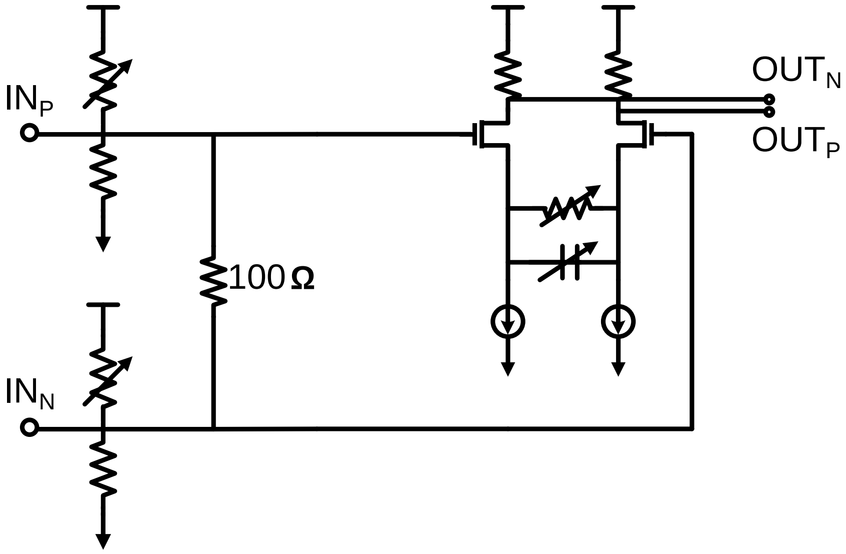

Figure 5.17: Structure of the receiver input stage, including DC biasing and CTLE implemented as an RC source degenerated differential pair.

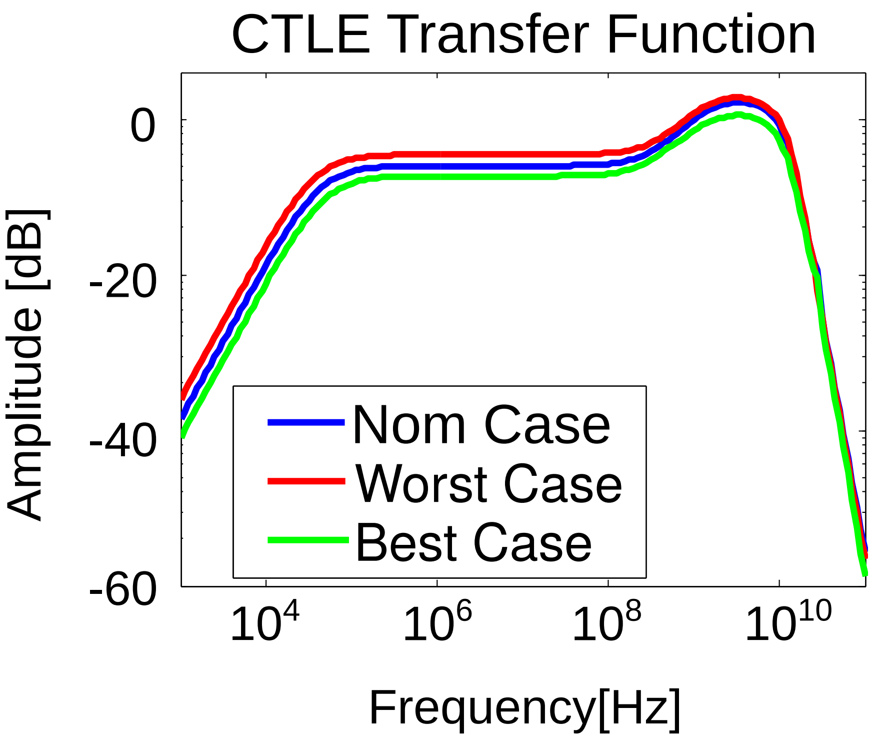

Figure 5.18: Transfer function of the CTLE for different PVT corners and fixed RC degeneration settings. Blue curve is the nominal case ($V_{DD} = 900mV$, 25°C, top-top technology corner), red curve is the worst case ($V_{DD} = 800mV$, -40°C, slow-slow technology corner) and the green curve is the best case ($V_{DD} = 1V$, 175°C, fast-fast technology corner).

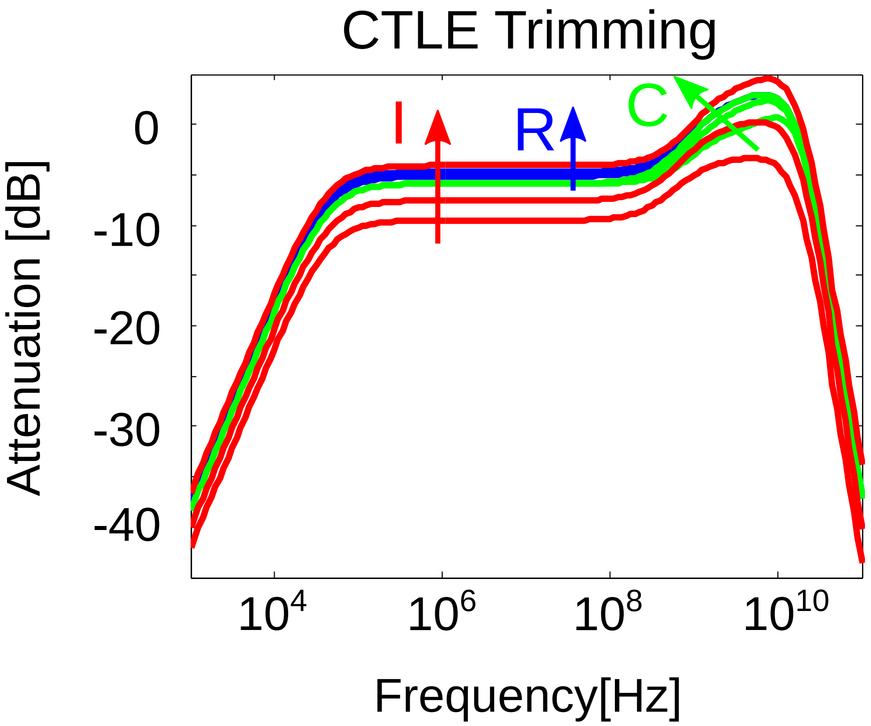

Figure 5.19: Transfer function of the CTLE for different configurations for biasing current or source resistance and capacitance. The curves are grouped into varying resistance trimming code (blue curves), varying capacitance trimming code (green curves) and varying current trimming code (red curves) while all others settings are fixed. The arrows indicate increasing values for resistance, capacitance and current. As expected, changes in the source resistance only influence the low-frequency behavior of the CTLE, capacitance changes shift the gain and the frequency of the boost and biasing current changes influence the whole transfer function.

5.3.2 DFE Timing

In this section, we will analyze the timing of the signals regarding DFE, in order to demonstrate that there is no issue with this part of the circuit. In all this subsection we will refer to Fig. 5.20 and we will analyze five time instants, named a, b, c, d and e in the figure. The analysis will be carried out for the even data path, but can conversely be done for the odd data path.

Let's start from time instant a): the transmitted bit b_0 that is visible at the differential input of the sampler (second panel) is a 1 and the previous one is a 0, therefore the positive $+h_1$ speculative differential input (violet curve) should be selected. The bits previously received are b_{-1}=0, b_{-2}=0, b_{-3}=1 and b_{-4}=1. At a), we only care of the signal entering h_2, h_3 and h_4 taps, since the h_1 tap will just decide which of the two speculative paths we will select out of the mux. The selected one then enters the output latch of the even data path and will be frozen at the next rising edge of the clock, so h_{1_p} has time until c) to settle. In fact, at a) we have that the input to h_{2_p}, h_{3_p} and h_{4_p} are respectively 0, 1 and 1, so it's fine. So h_{1_p} should be ready and correct at c), since at this moment in time b_0 is frozen at the output latch: h_{1_p} should be 0, and it is. So we know that all DFE taps are correctly set when it comes to help in sampling correctly b_0.

Now we can check the other aspect: does b_0 operate correctly to help with the equalization of the following bits? Does it arrive on time to be h_{1_p}, h_{2_p}, h_{3_p} and h_{4_p} for b_1, b_2, b_3 and b_4 respectively? Let's start from the easiest, h_{2_p}. At b), b_0 has been already shifted through the mux and the mux output is frozen since the previous data coming from the odd path, D$_{OUTP_{ODD}}$, is frozen: so b_0 is ready to be h_{2_p} for b_2 way before c), which is the timing instant at which b_2 arrives at the speculative samplers input. Let's now focus on h_{1_p} and h_{3_p}: we have seen that at c) b_0 is frozen at the output latch of the even data path: this means that it is ready to go also to the odd data path for playing the part of h_{1_p} and h_{3_p} for b_1 and b_3, respectively. In fact, b_1 is frozen at the output latch of the odd data path only at d), so b_0 is ready since 100ps for the right selection on the mux of the odd data path. At the same time d), b_3 is frozen at the sampler on the odd data path and b0 is again ready since 100ps at the input of a transparent latch which gets frozen at the same instant as the sampler (see Fig. 5.7), so as long as the D-to-Q delay of the latch is below 100ps we are fine, and we are well below this number (circa 30ps in the worst case). Finally, we check what happens for h_{4_p}: we have seen that at c) b_0 is frozen at the output latch, so it is ready at the h_4 latch input as well, which starts to be transparent at this moment. This latch will be frozen at d), so b_0 is ready to be h_{4_p} for b_4 100ps before this instant.

Finally, we take a look at the glitch that occurs on the signal h_{2_p} input slightly after e). At e), the output latch (the one at rightmost in Fig. 5.7) of the odd data stream starts to be transparent, so there might be a glitch in the even data mux output, since there is a change in its selection signal. Anyway, this glitch is not relevant because the mux output has to be frozen 100ps later and ready 200ps later. Its value at e) has no meaning, since the output latch input is open at this point in time. Concluding, this glitch does not hurt our DFE process.

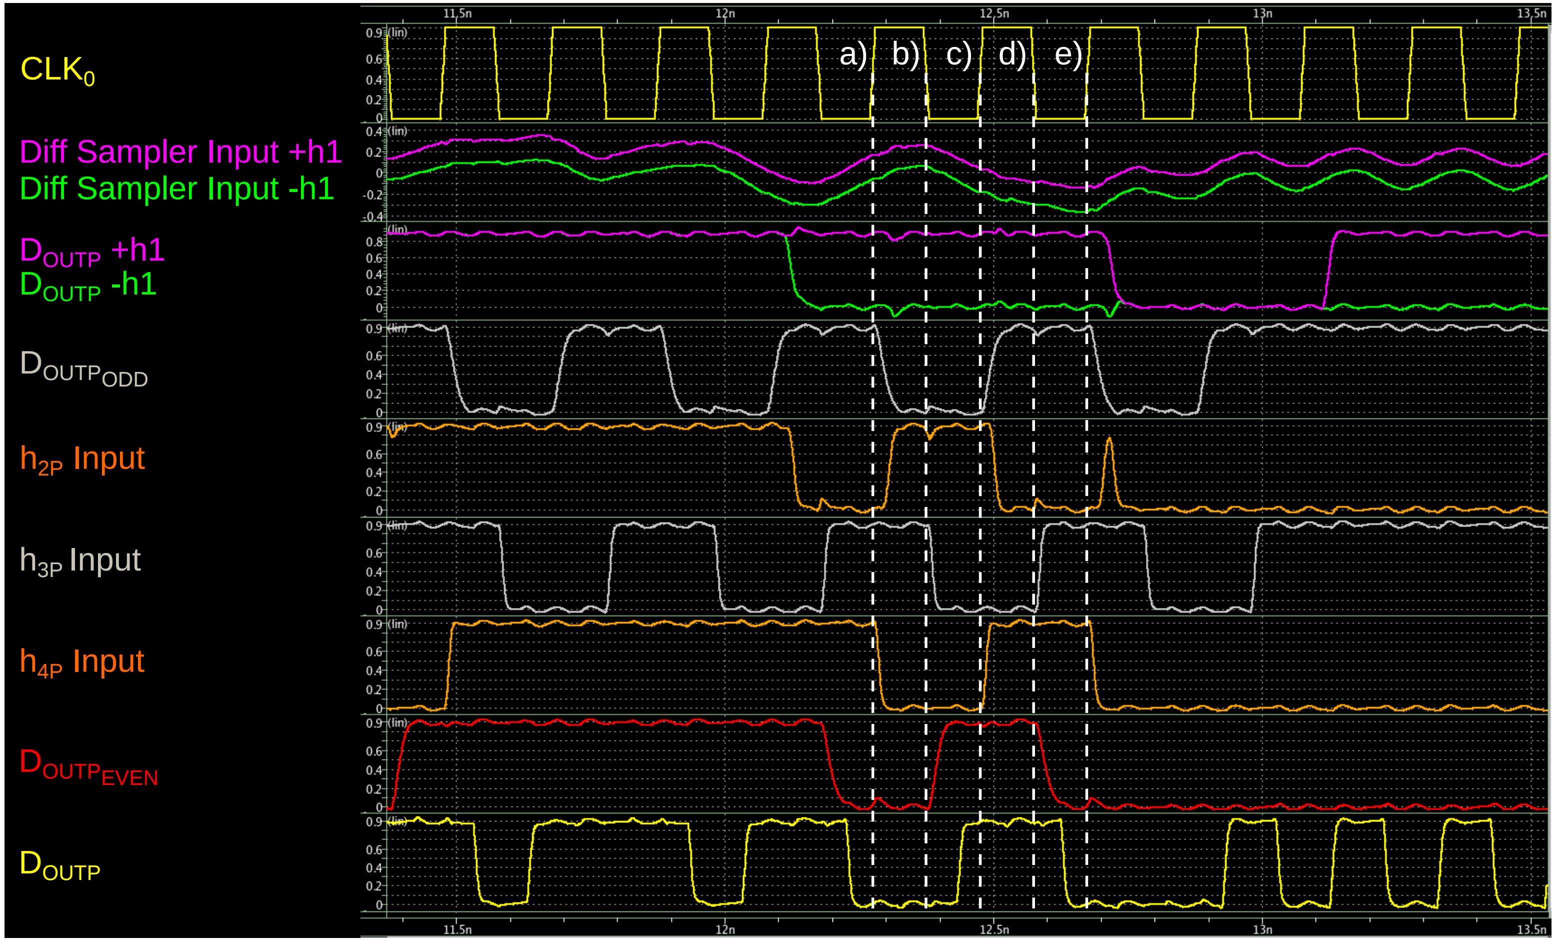

Figure 5.20: Timing for DFE signals. From top to bottom, the signals are the following: CLK$_0$, positive $+h_1$ speculative differential input for data samplers, negative $-h_1$ speculative differential input for data samplers, positive output of the positive speculative $+h_1$ sampler, positive output of the negative speculative $-h_1$ sampler, positive output data coming from the odd data path (so this is the mux selection signal in Fig. 5.7), positive input to h_{2_P} tap, positive input to h_{3_P} tap, positive input to h_{4_P} tap, positive output data of the even data path and positive output data at full rate (so to show the complete reconstruction of the data in the second panel).

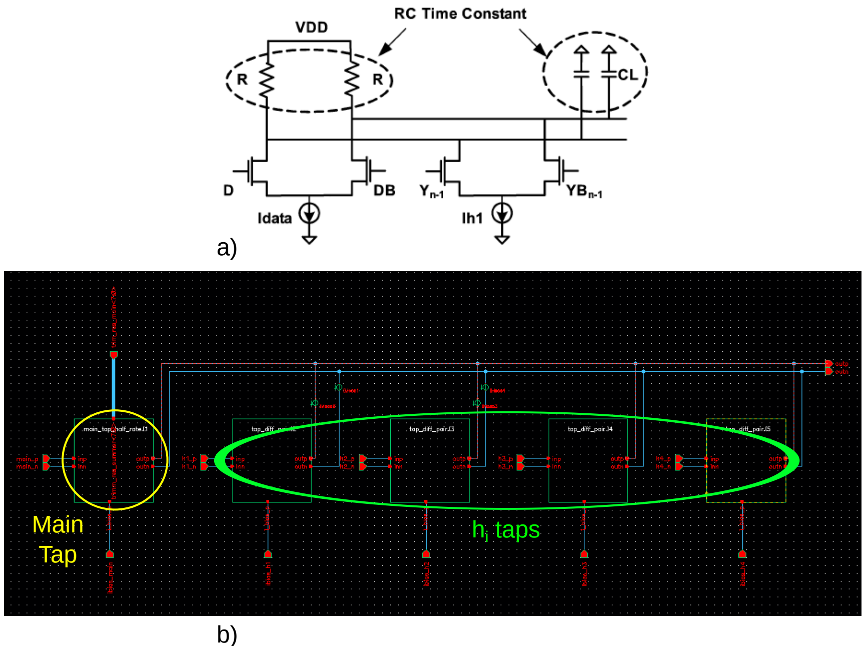

Figure 5.21: a) Schematic of a current-mode summer [74] and b) our transistor level implementation with 4-taps DFE.

5.3.3 Summers, Comparators and PI

The summers have been realized with a traditional current-mode architecture [74]. The basic schematic of this architecture is shown in Fig. 5.21 a), which is the same used in our transistor level implementation reported in Fig. 5.21 b), highlighting the four taps of DFE.

The important thing for the summer is that its output should reach its final value as fast as possible for ISI cancellation to work properly. This time is set by the RC constant highlighted in Fig. 5.21. For 10 Gbps half-rate RX, this RC delay must be around 17ps [74], which sets a mutual constraint between drain resistance of the main tap and output capacitance. Anyway, we are well below this time constant, since we have just four small differential pair connected at the output node, summing up to a C_L of circa 3fF.

In this stage, we added a high level of programmability in order to offer a variety of solutions: the current in the main tap of the summers can be trimmed with 16 levels, and so does the current flowing in the DFE taps. The drain resistances of the main summer feature 8 different trimmable levels, so that the finest DFE correction that can be programmed is 6mV.

Comparators have been implemented in a classical double-tail topology, whereas latches have been designed in the same pseudo-differential architecture used in the transmitter [57]-[53]. The schematic of a latch is reported in Fig. 5.22.

As already mentioned, the PI is implemented in a traditional fashion [21] (see Fig. 5.23) and is divided in 64 steps. A DPLL, not described in this work,

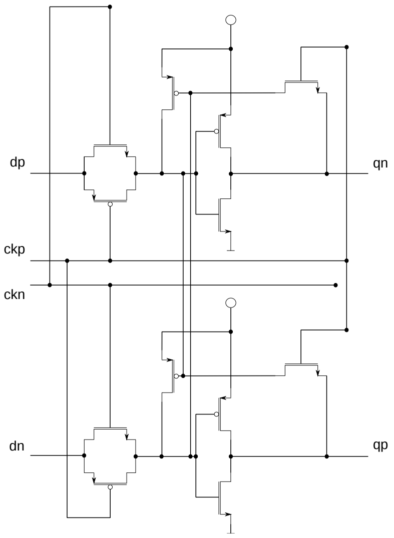

Figure 5.22: Schematic of a pseudo-differential latch used in the receiver.

provides a 10GHz clock, that then is divided down to 5GHz in four phases, as already described in Fig. 5.8.

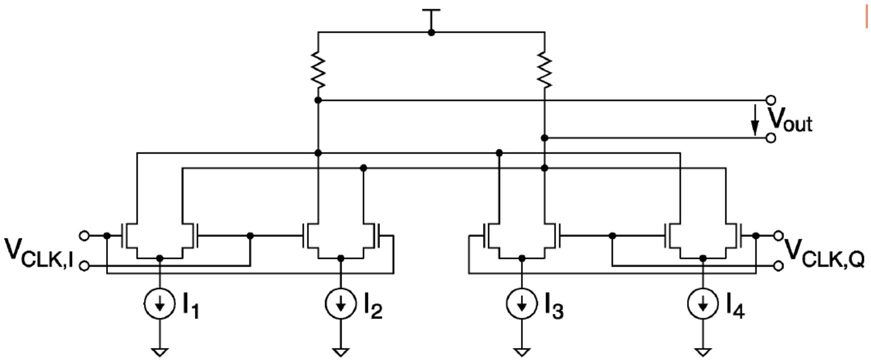

Figure 5.23: Schematic of the core of a traditional phase interpolator [21]. Currents I_1-I_4 are controlled via the 6-bits PI code set by the CDR algorithm.

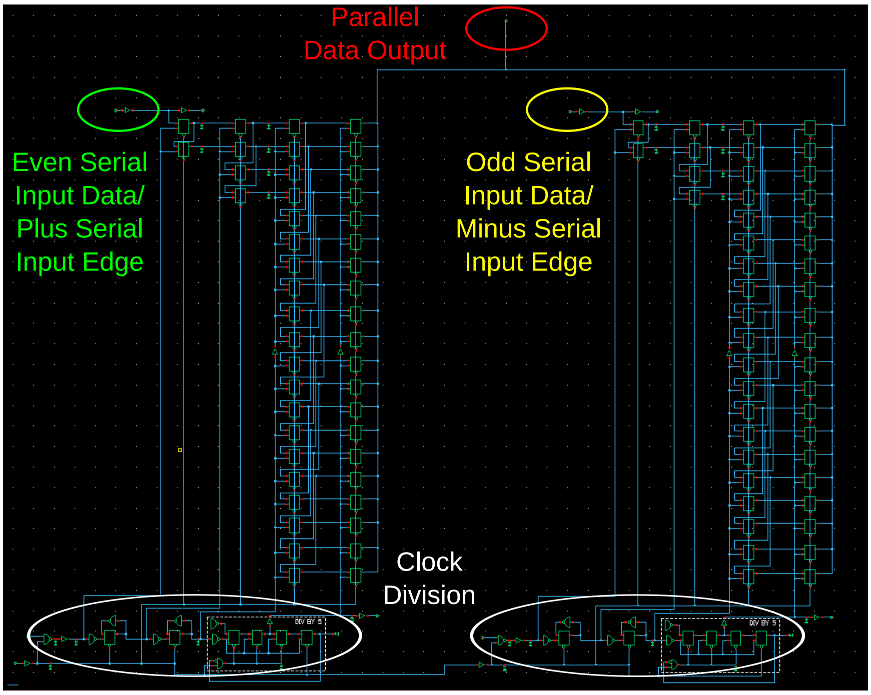

Figure 5.24: Schematic of the 2:40 deserializer used in the receiver.

5.3.4 Deserializers

All three deserializers have been implemented as 2:40 deserializers and their schematic is reported in Fig. 5.24. Two serial data streams are taken as inputs to be interleaved at the parallel data output. As can be seen in Fig. 5.25, the data are parallelized in four stages with four different clocks with decreasing frequency, generated in the clock division part of the circuit as shown in Fig. 5.26. Since the deserializer is taking as input two distinct serial data streams, two different clock phases have to be used to latch these serial data in.

The data deserializer takes two serial data stream as inputs, even and odd: for this reason, both CLK_0 and CLK_{180} must be provided to the data deserializer. On the contrary, the two edge deserializers take as input two serial streams that can be latched with the same clock: the even edge deserializer takes the plus and minus speculative even edges and interleaves them, and same goes for the odd ones and the odd deserializer. In fact, both even edge streams can be latched in with CLK_{90} and both odd edge streams with CLK_{270}. For symmetry and timing reasons, we don't change the deserializer architecture of the data deserializer to have just one clock as an input, but we provide the same clock two times to the edge deserializers. Fig. 5.27 briefly summarizes the timing of this block, which represents the interface between the analog and the digital world: in fact, one of the 250MHz clocks that are generated inside the deserializer is also used as a clock for the digital part.

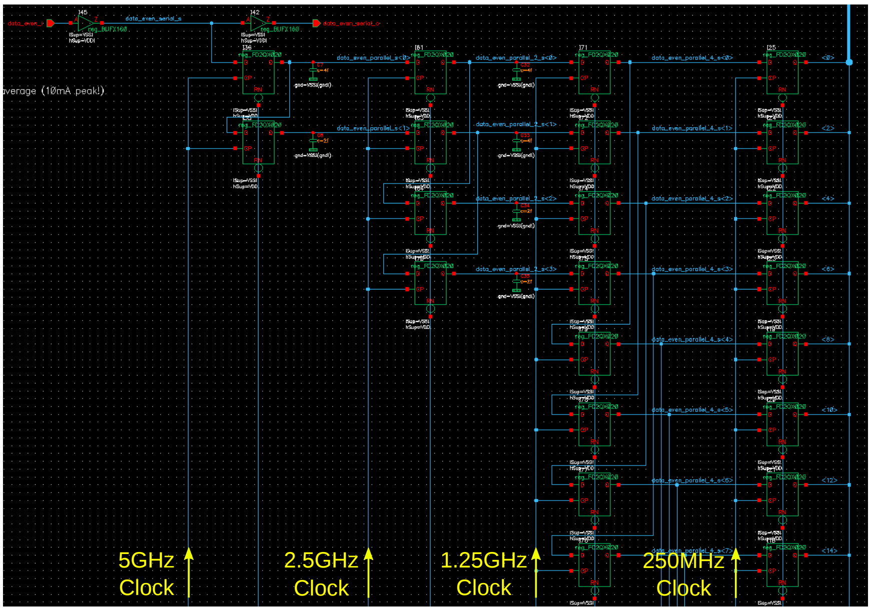

Figure 5.25: Zoom of Fig. 5.24 on the parallelization of the input data streams.

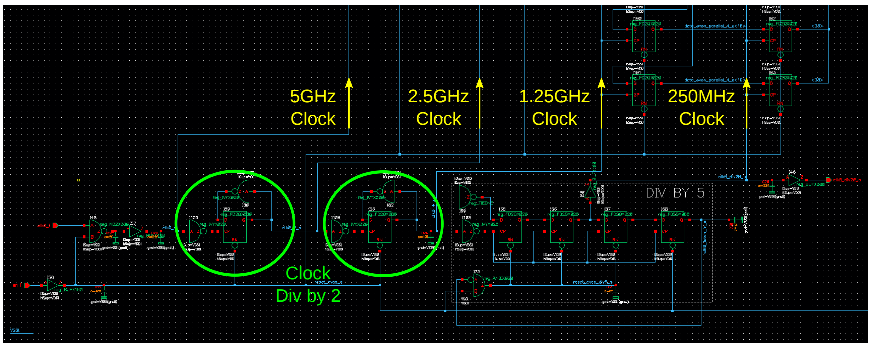

Figure 5.26: Zoom of Fig. 5.24 on the divided clocks generation.

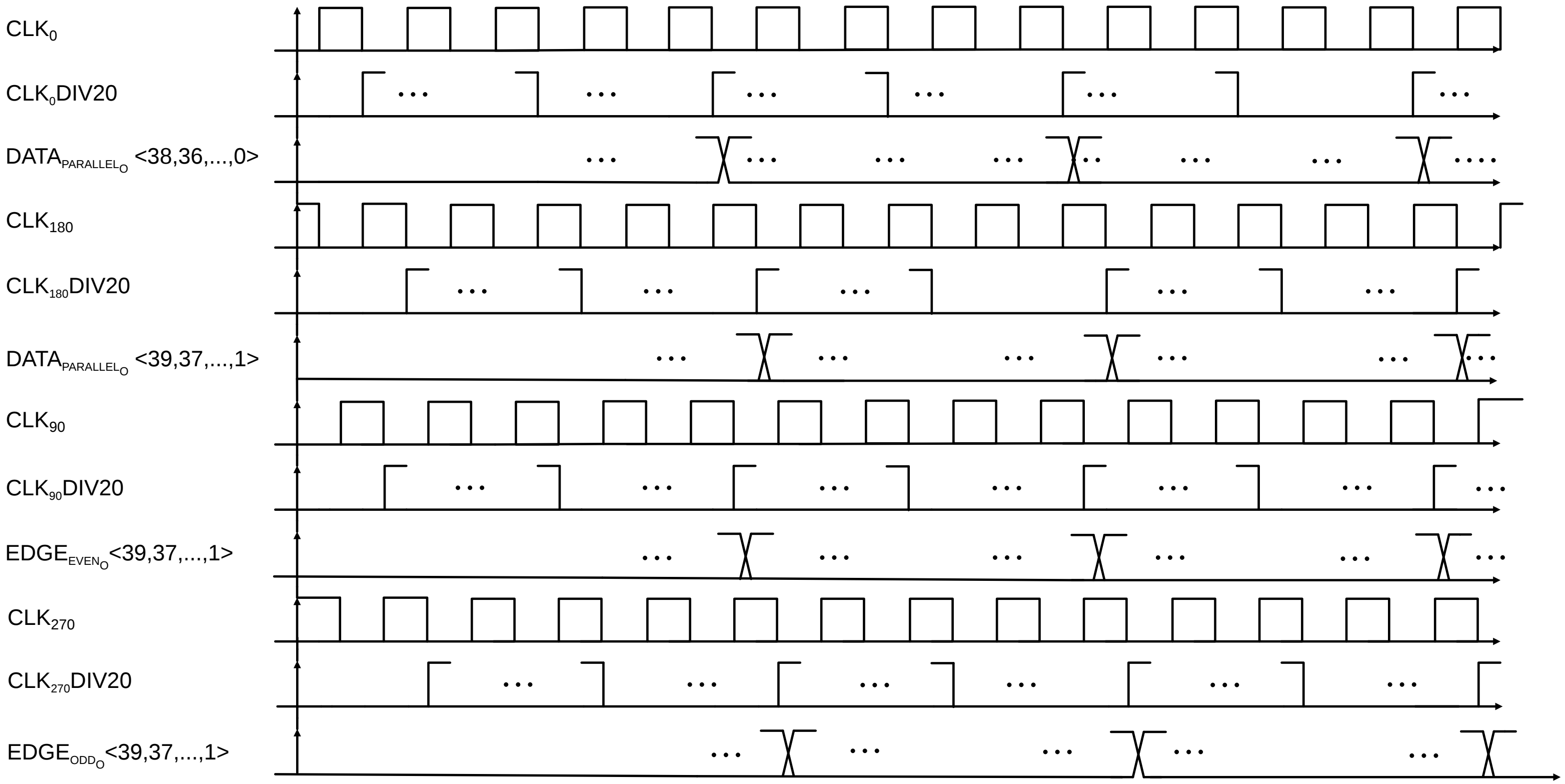
CLK_0
$\mathrm{CLK}_0\mathrm{DIV20}$
$\mathrm{DATA}_{\mathrm{PARALLEL}_O}$ <38,36,...,0>
CLK_{180}
$\mathrm{CLK}_{180}\mathrm{DIV20}$
$\mathrm{DATA}_{\mathrm{PARALLEL}_O}$ <39,37,...,1>
CLK_{90}
$\mathrm{CLK}_{90}\mathrm{DIV20}$
$\mathrm{EDGE}_{\mathrm{EVEN}_O}$<39,37,...,1>
CLK_{270}
$\mathrm{CLK}_{270}\mathrm{DIV20}$
$\mathrm{EDGE}_{\mathrm{ODD}_O}$<39,37,...,1>

Figure 5.27: Timing of clocks and outputs of the deserializer shown in Fig. 5.24.

5.3.5 Timing of the CDR Algorithm

This section aims to add some transistor level insight on the CDR algorithm already introduced in Section 5.2. In Fig. 5.28 we show some transistor level waveforms that, along with the following explanation, should clarify the concepts previously expressed.

From D_{OUTP} stream, we see that the received data sequence is 00010101010. In this analysis we focus on the fifth, sixth and seventh bit of this sequence, so on a 010 pattern. At the timing instant c), we find the third bit of a 010 sequence in D_{OUTP}: this is the first step of our algorithm. We will refer to this three bit as b_0, b_1 and b_2. We see that, at c), b_2, the second zero in the 010 sequence is frozen at the rising edge of CLK_0, so this is an even bit. This zero that we see now frozen in D_{OUTP} has actually been frozen at instant a) (200ps earlier than c)) at the two speculative samplers, of which we see the differential inputs in the second panel. In the second step of the algorithm, we have to check the edge samples to understand whether the sampling clock is well positioned or not. We have just said that b_2 is an even bit, so since we want to look at the transition between b_1 and b_2, we have to check the edges sampled with CLK_{270}, so at the odd edges before b_2, thus at $EDGE_{ODD+h_1}$ and $EDGE_{ODD-h_1}$ at timing instant b). At this point in time, the two speculative odd edges $EDGE_{ODD+h_1}$ and $EDGE_{ODD-h_1}$ are different between each other, therefore the sampling clock is correctly positioned. We can actually check this by looking at a), where you can see that the rising edge of CLK_0 is correctly set with respect to the two speculative differential inputs of the samplers.

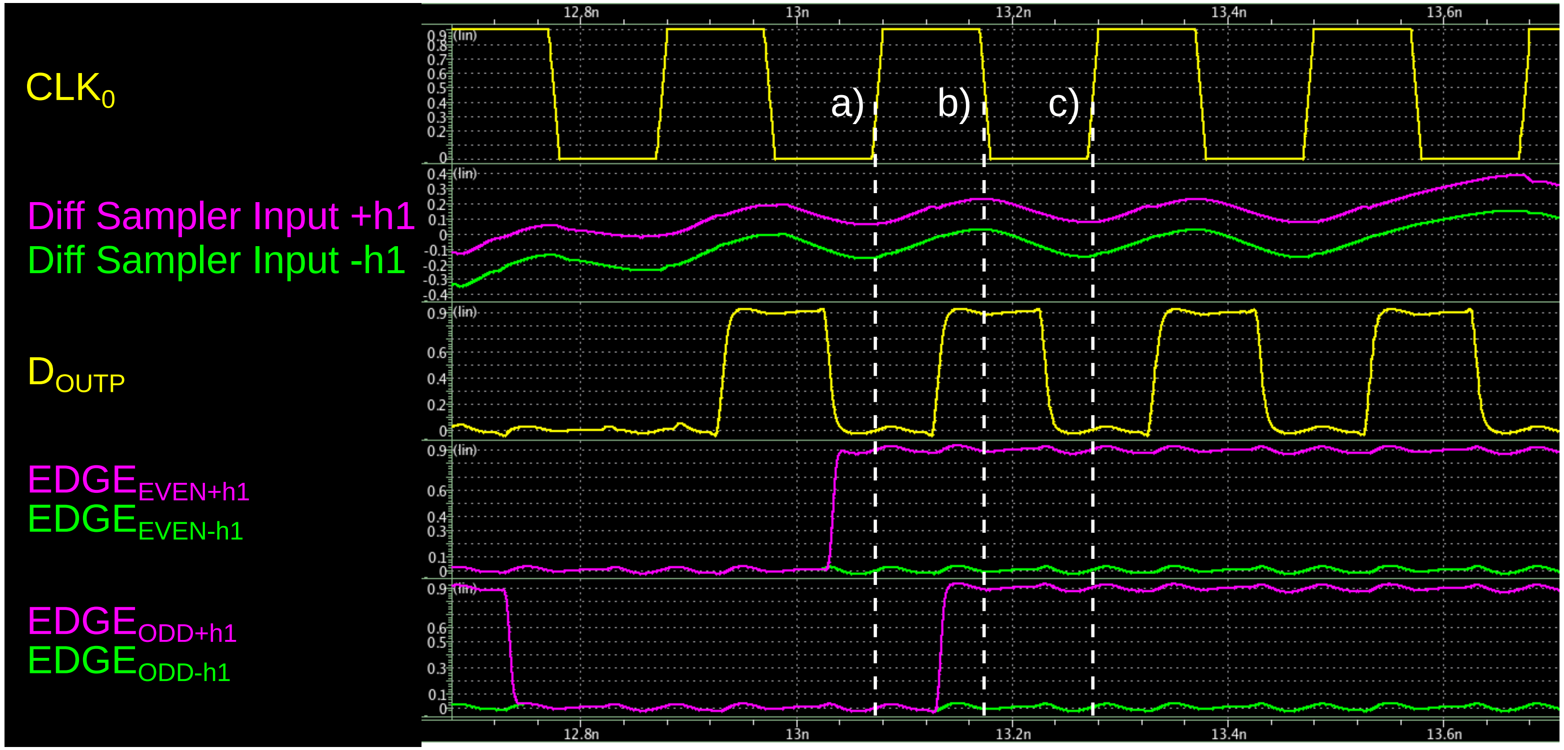

Figure 5.28: Transistor level waveforms to show the functioning of the CDR algorithm. From top to bottom, the signals are the following: CLK$_0$, positive $+h_1$ speculative differential input for data samplers, negative $-h_1$ speculative differential input for data samplers, positive output data at full rate, positive $+h_1$ even edge samples, negative $-h_1$ even edge samples, positive $+h_1$ odd edge samples and negative $-h_1$ odd edge samples.

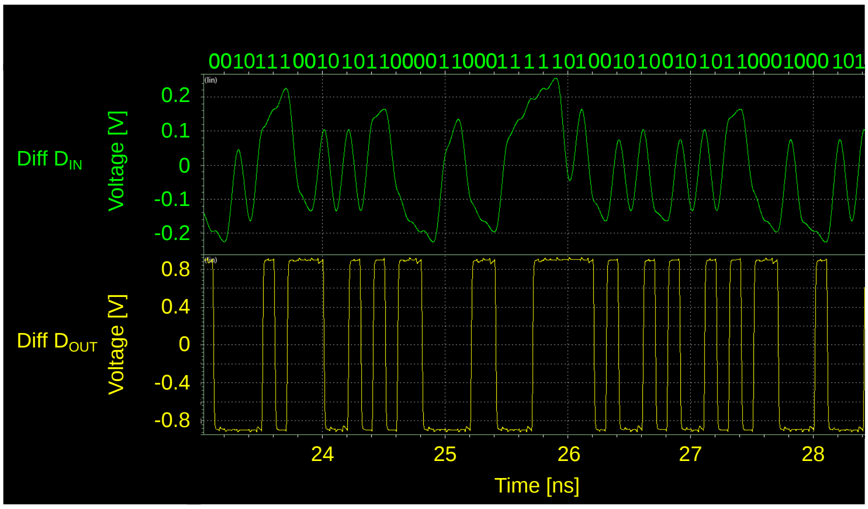

Figure 5.29: Data reconstruction of the receiver. The green curve shows the differential input data at the receiver, the yellow curve the differential reconstructed data. On top, the transmitted bits are shown.

5.3.6 Results with fixed PI code

All blocks have been singularly validated post-layout in all PVT corners. The functioning of the receiver has been first tested by fixing the PI code to an optimal value. The transmitter and the receiver have been connected by an ideal channel, in order to check the basic functioning of the receiver structure. Fig. 5.29 shows an error-less data reception, that at this point has been demonstrated in all PVT corners (from $V_{DD} = 800mV$, -40°C, slow-slow technology corner to V_{DD}=1V, 175°C, fast-fast technology corner).

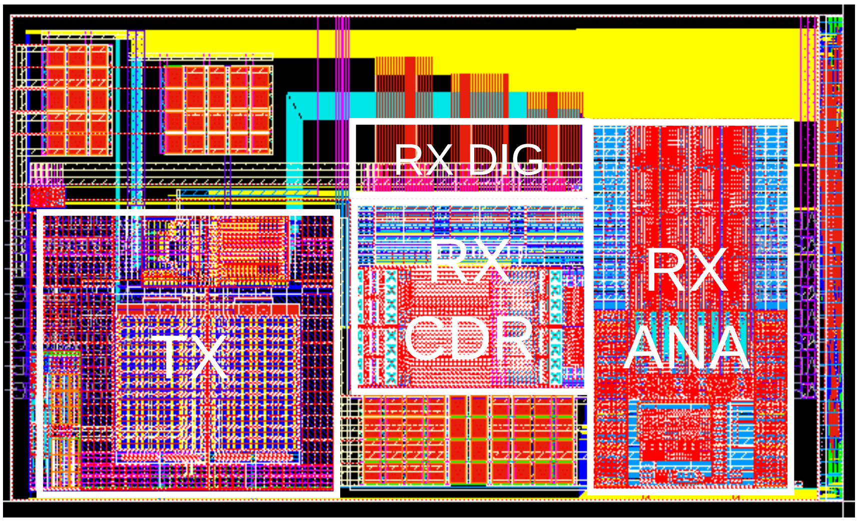

Figure 5.30: Final layout of the whole system. The total area is 0.125mm^2, whereas the area occupied by the receiver amounts to 0.05mm^2.

5.3.7 Results with XA-VCS

To verify the proposed receiver topology and CDR algorithm, the whole RX was assembled (Analog part, Digital part including CDR algorithm and JTAG interface and Simplified Timing Shell, STS, which re-times the interface signals between digital and analog) along with the post-layout transistor level model of a TX [57] and the whole system was simulated using the simulator XA-VCS [81]. The final layout of the structure is reported in Fig.5.30. These simulations show a power consumption for the RX of 2.05mW/Gbps in nominal PVT conditions (V_{DD}=900mV, typical technology corner, 25°C): the analog part of the receiver (bias distribution, CTLE, DFE, summers, comparators) consumes 0.77mW/Gbps, whereas the CDR circuitry (PI, PI Decoder, 2:40 Deserializers) 1.28mW Gbps. In the least-consuming PVT corner (V_{DD}=800mV, slow-slow technology corner, -40°C), the analog part consumes 0.5mW/Gbps, whereas the CDR circuitry 0.8mW/Gbps, which gives a total 1.3mW/Gbps.

Example of simulation results for different PD algorithms, different channels and different DFE settings are reported in Figs. 5.31-5.32. As from plot a), between our CDR algorithm and a pure Alexander PD [82] (no transition filtering) there is a difference in the clock sampling position of roughly eight steps (i.e., 25ps). It can be seen from plot b that our PD algorithm is very well settled with respect to the input data of the sampler, meaning therefore that the clock as set by the Alexander PD may have a displacement from the correct sampling position, thus reducing the jitter budget for the PLL. Furthermore, in our algorithm the oscillations of the PI are smaller than when using an Alexander PD without transition filtering.

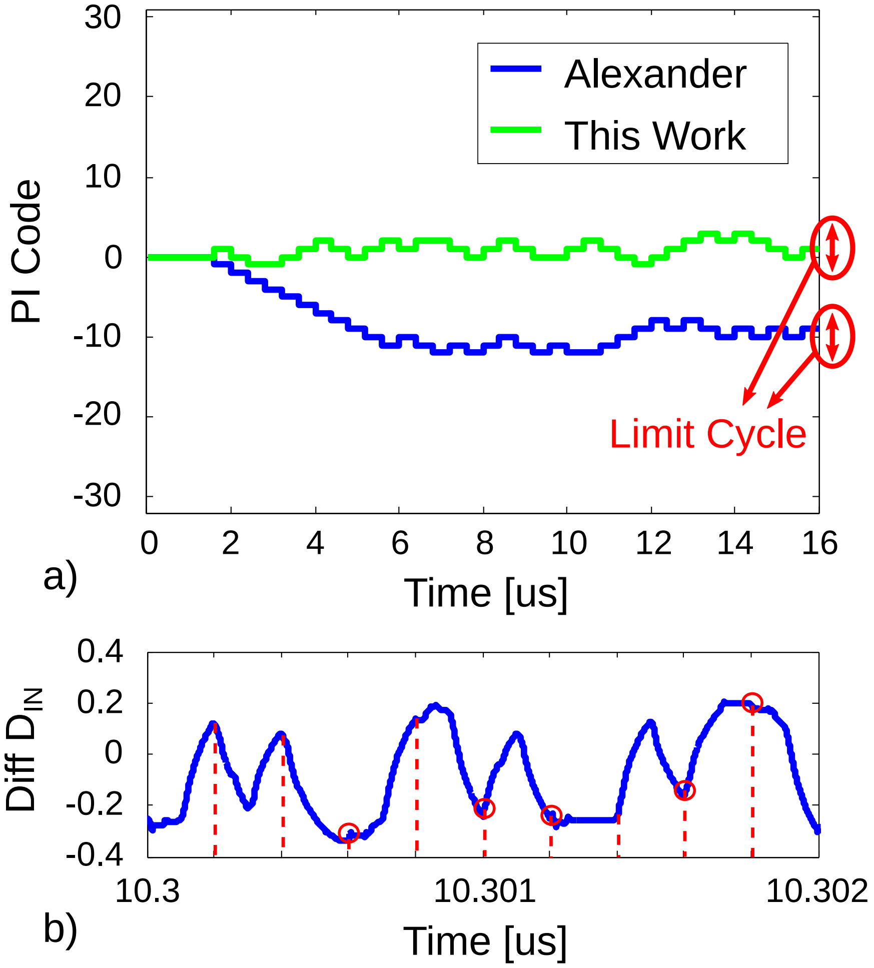

Figure 5.31: a) Phase Interpolator selection code for Alexander PD and for the PD of this work (Fig.5.9) when using $n_{avg} = 100$. b) Data sampling for even path in the case of the PD of this work (red dotted lines). Only $-h_1$ speculative even path is here shown for clarity. The samples chosen among the two speculative even paths are circled in red: since here $-h_1$ speculative even path is shown, only samples preceded by a '1' are chosen, whereas the others are taken from $+h_1$ speculative even path (here not shown). Both plots have been obtained for a template channel consisting of an RC low-pass filter showing 5dB of loss at 5GHz. The peak-to-peak output swing of the transmitter is 380mV. The values used for h_1, h_2, h_3 and h_4 are 78mV, 39mV, 20mV and 20mV respectively and are not optimized to show the robustness of the algorithm in case of sub-optimal equalization. In both cases, the PI limit cycle has been obtained from 16μs long simulations as in Fig. 8a, after having removed the initial transient of the PI code (in these cases, 6μs).

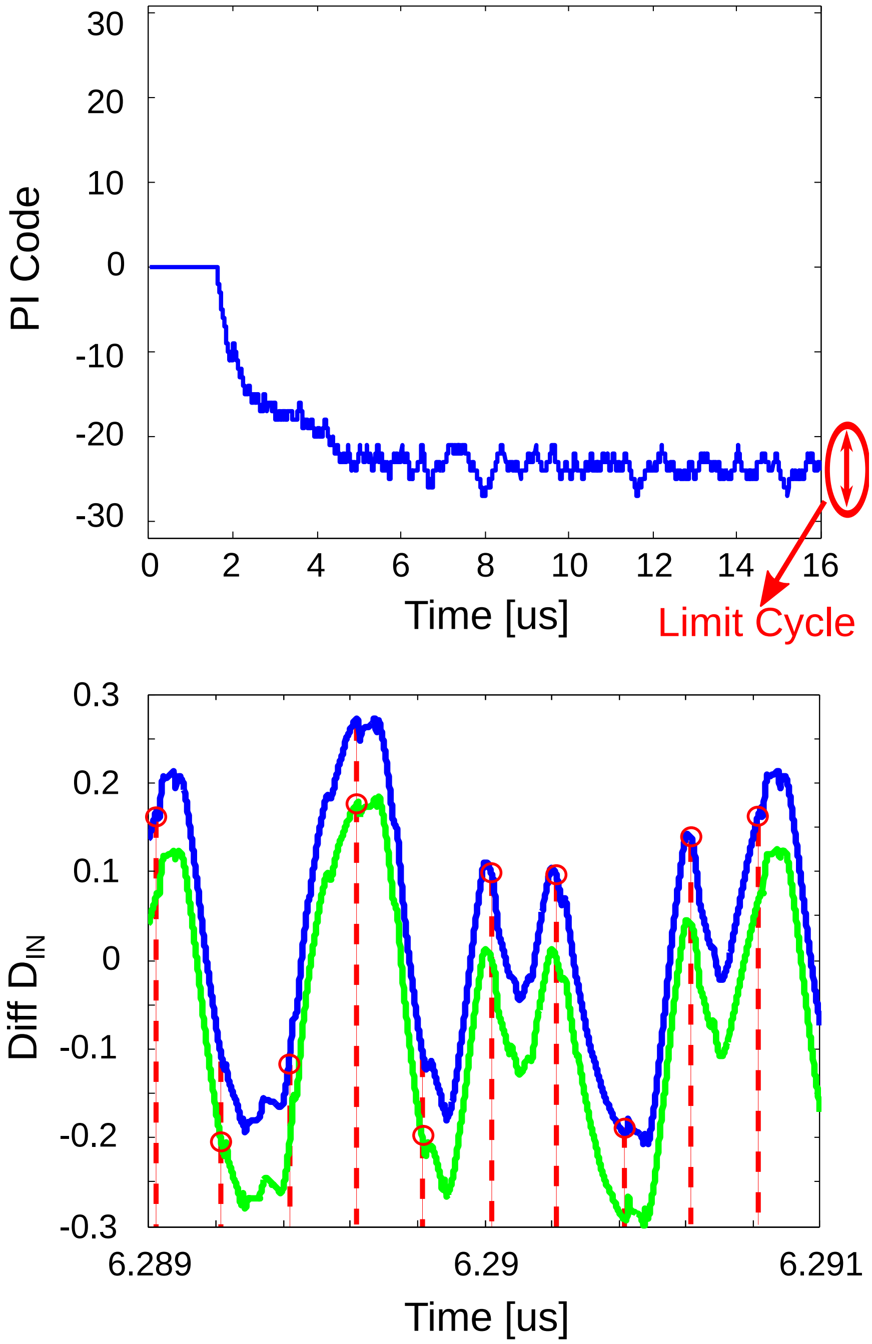

Figure 5.32: a) Phase Interpolator selection code for the PD of this work (Fig.5.9) when using $n_{avg} = 7$. b) Data sampling for even path in the case of the PD of this work (red dotted lines). Both $+h_1$ and $-h_1$ speculative path are here shown. The samples chosen among the two speculative even paths are circled in red: when preceded by a '1', h_1 samples are chosen, whereas the others are taken from $+h_1$ speculative even path. The plot has been obtained for a template channel consisting of an RC low-pass filter showing 10dB of loss at 5GHz. The peak-to-peak output swing of the transmitter is 380mV. The values used for h_1, h_2, h_3 and h_4 are 39mV, 20mV, 0mV and 0mV respectively and are not optimized to show the robustness of the algorithm in case of sub-optimal equalization. The PI limit cycle has been obtained from 16μs long simulations as in Fig. 5.31a), after having removed the initial transient of the PI code (in these cases, 6μs).

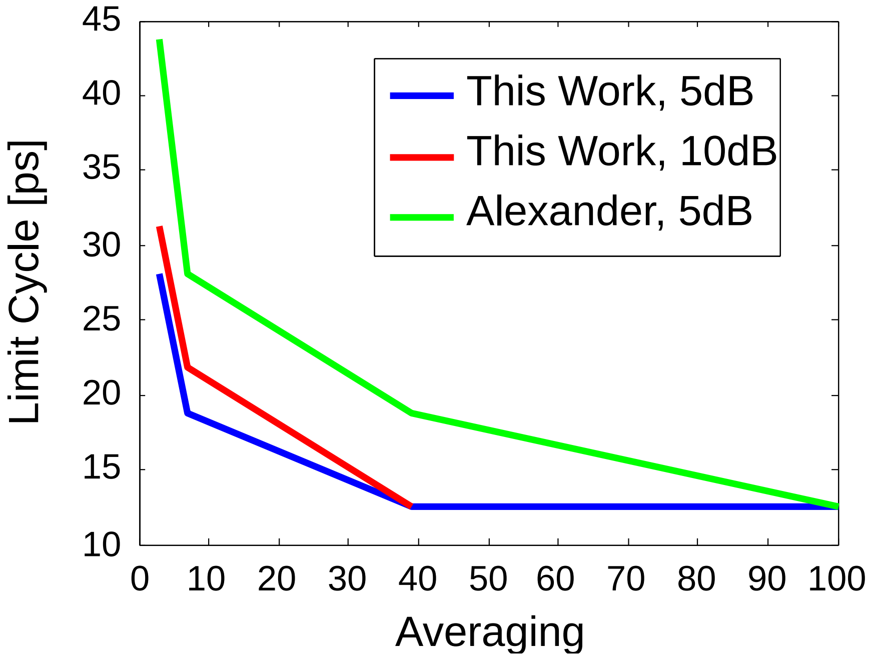

Figure 5.33: Amplitude of the PI limit cycle for different n_{avg}. Two different pairs of channels and DFE settings have been used. One, 5dB, has the same channels and DFE settings as in Fig. 5.31, whereas the other, 10dB, has a a template channel consisting of an RC low-pass filter showing 10dB of loss at 5GHz and the values used for h_1, h_2, h_3 and h_4 are 39mV, 20mV, 0mV and 0mV respectively and are not optimized to show the robustness of the algorithm in case of sub-optimal equalization. In both cases, the PI limit cycle has been obtained from 16μs long simulations as in Fig. 5.31a, after having removed the initial transient of the PI code (in these cases, 6μs).

To further analyze this latter aspect, the limit cycle of the PI code waveform is plotted vs n_{avg} in Fig.5.33: a high number of averages is required by the Alexander PD w/o transition filtering to keep under control the limit cycle. It must be noted that our PI algorithm depends on the value of h_1 set in the DFE: if no DFE is applied, our algorithm is equivalent to an Alexander PD with a transition filtering of three bits. By increasing h_1, the sampling time for which both *early* and *late* outputs stay low increases.

Finally, Figs. 5.34-5.38 show the diagram at different point of transceiver, namely at the output of the transmitter, at the input of the receiver, at the output of the CTLE and at the input of the two speculative samplers. The plots have been obtained for a template channel consisting of an RC low-pass filter showing 10dB of loss at 5GHz. The peak-to-peak output swing of the transmitter is 380mV. The values used for h_1, h_2, h_3 and h_4 are 78mV, 20mV, 0mV and 0mV respectively

Fig.5.39 show the eye mask in different points of the system with the same settings as in Fig. 5.31. The channel and DFE settings are the same as Fig.

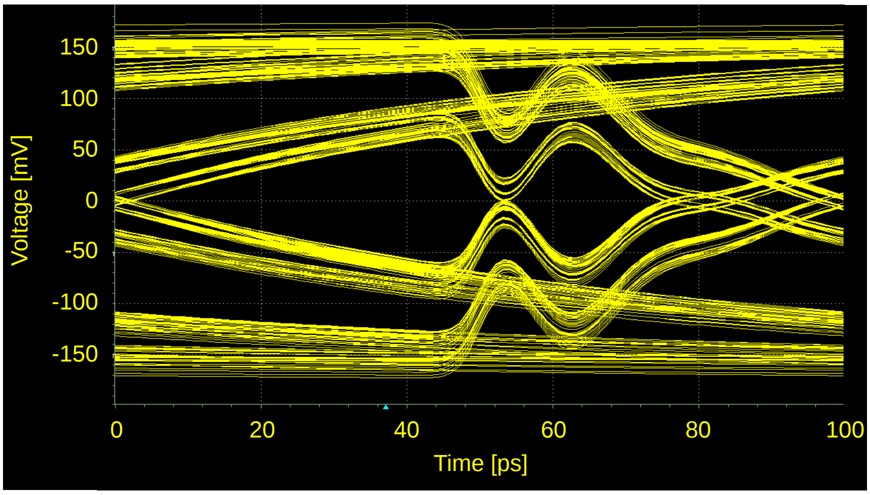

Figure 5.34: Eye diagram at the output of the transmitter. In this figure, we see the typical effect of reflections due to the channel.

5.31. It must be noted that, in order to build the eye mask at the input of the samplers, one must take the transitions 0 → 1 and 0 → 0 from the positive speculative sampler (see Fig. 5.37) and the transitions 1 → 0 and 1 → 1 from the negative speculative sampler (see Fig. 5.38).

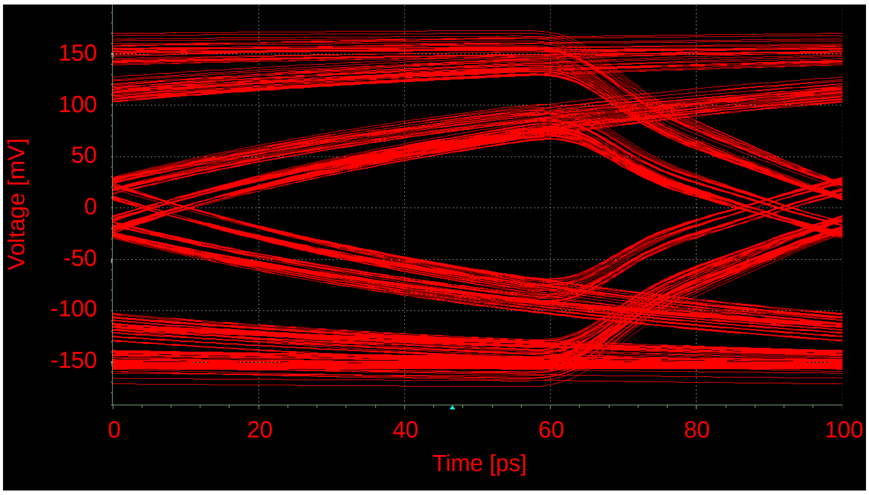

Figure 5.35: Eye diagram at the receiver. From this figure, the low-pass nature of the channel can be understood.

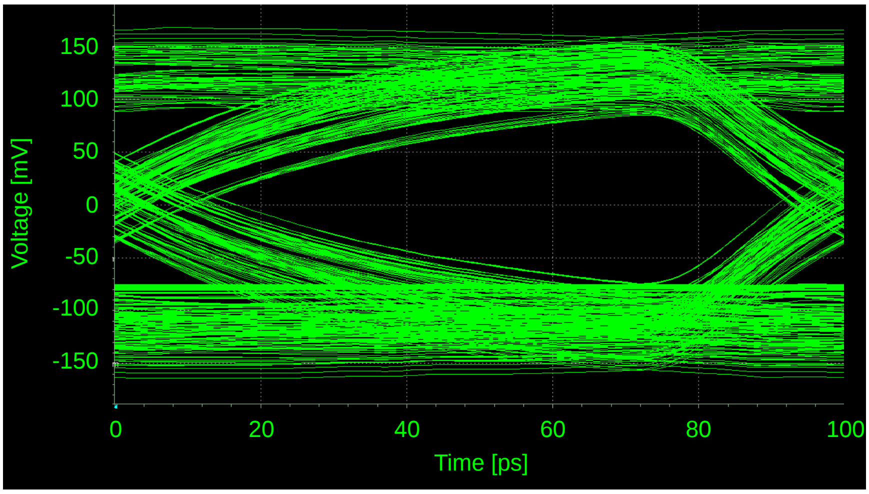

Figure 5.36: Eye diagram at the output of the CTLE. As can be seen, the high and low levels of the eye diagram are now smaller, but also the rise and fall times are smaller than the ones in Fig. 5.35.

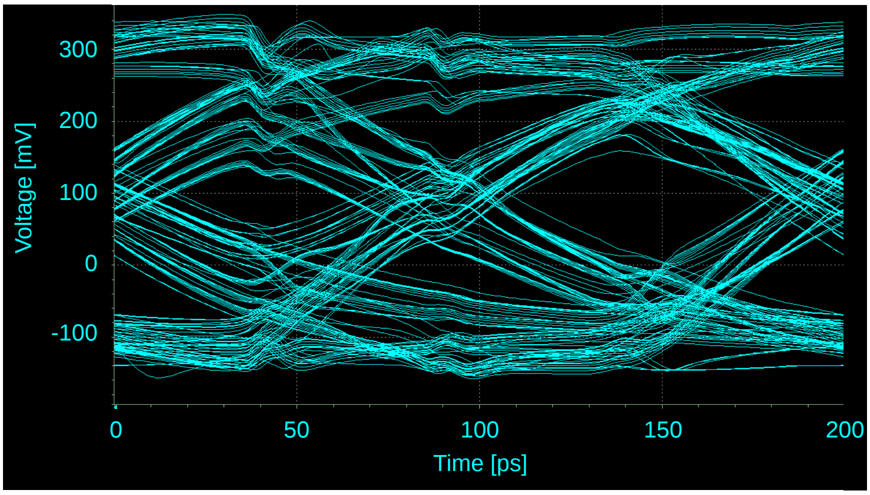

Figure 5.37: Eye diagram at the input of the positive even speculative sampler. Here a 200ps eye diagram is shown in order to highlight the fact that the frequency of the clock used to sample this eye is 5GHz. In fact, the only eye diagram that matters is the one on the right. As said in the text, the only transitions that matter in this eye diagram are the 0 → 1 and 0 → 0.

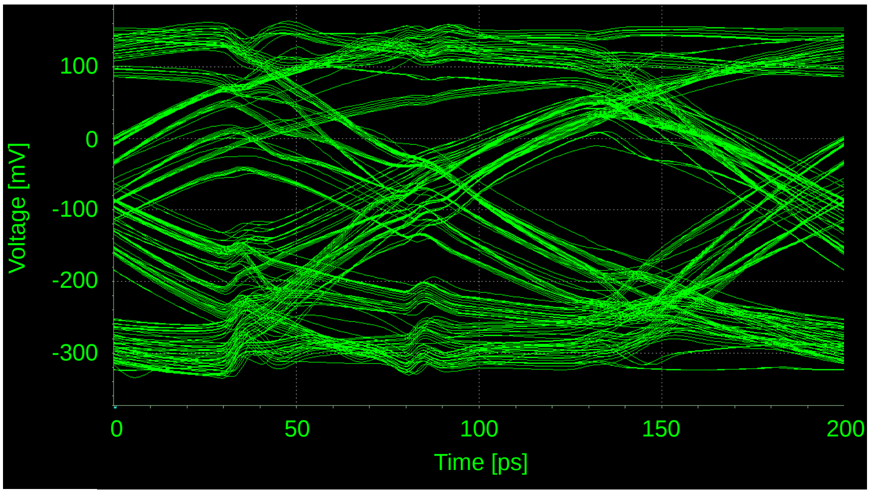

Figure 5.38: Eye diagram at the input of the negative even speculative sampler. Here a 200ps eye diagram is shown in order to highlight the fact that the frequency of the clock used to sample this eye is 5GHz. In fact, the only eye diagram that matters is the one on the right. As said in the text, the only transitions that matter in this eye diagram are the 1 → 0 and 1 → 1.

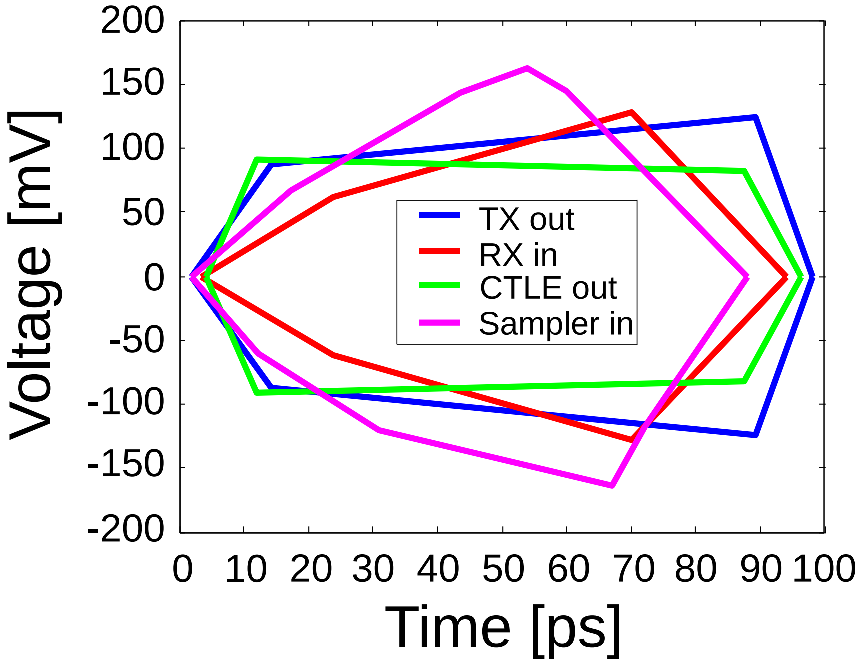

Figure 5.39: Eye masks in different points of the signal path for the same channel and settings as in Fig. 5.31.

5.4 Conclusions

We have designed a 10Gbps HSSI half-rate receiver with power efficiency aligned with the state-of-the-art. An original CDR algorithm is proposed that works in the presence of loop-unrolled DFE with a limited amount of transition filtering, allowing for a large CDR bandwidth. The algorithm is robust and does not require specific optimization of the DFE taps. The correctness of the proposed algorithm has been verified by simulating a full transceiver architecture at transistor level with post-layout parasitics. Experimental data on the fabricated system will be available by the end of the year.

Chapter 6

Characterization of the Full Link

In this chapter we analyze the results of the measurements of InnoTC of the full link as described in the previous chapters. We will start with a quick overview of the test-chip, focusing in particular on the high-speed transceiver, then pass to evaluate the BER of the link and the obtained bathtubs and how do the two phase detection algorithms described in Chapter 5 perform.

6.1 InnoTC Overview

Fig. 6.1 show the layout of InnoTC along with the bondplan to the package as realized for the measurements of the high-speed link. The total area of InnoTC is roughly 9mm^2, and on top of it several other circuits have been placed. The total area of the high-speed interface, including decoupling capacitance, bias and clock distribution, PRBS generator, PRBS checker and JTAG interface is roughly 0.5mm^2, with the decoupling capacitance occupying almost two thirds of it. The high-speed interface analyzed here is located in the left-side of the chip, near to the digital PLL providing the interface with the clock. The DPLL itself takes the clock from a DCO which is centered around 17.4GHz, tunable from 16.4GHz to 18.4GHz. The DPLL then divides this clock by a factor of two and delivers it to the high-speed interface: for this reason, only transmission speeds between 8.2Gbps and 9.2Gbps are available. Thanks to the division by two and by four implemented inside our transceiver, also 2.1-3.2Gbps and 4.1-4.6Gbps transmission speed can be tested.

It can be seen from Fig. 6.1 that, as explained in the previous chapter, we have distinct power domains, two for the transmitter and two for the receiver. As already stated in Chapter 4, for each one of these power domain two pads are available on chip, so to reduce the parasitic inductance on the supply by a

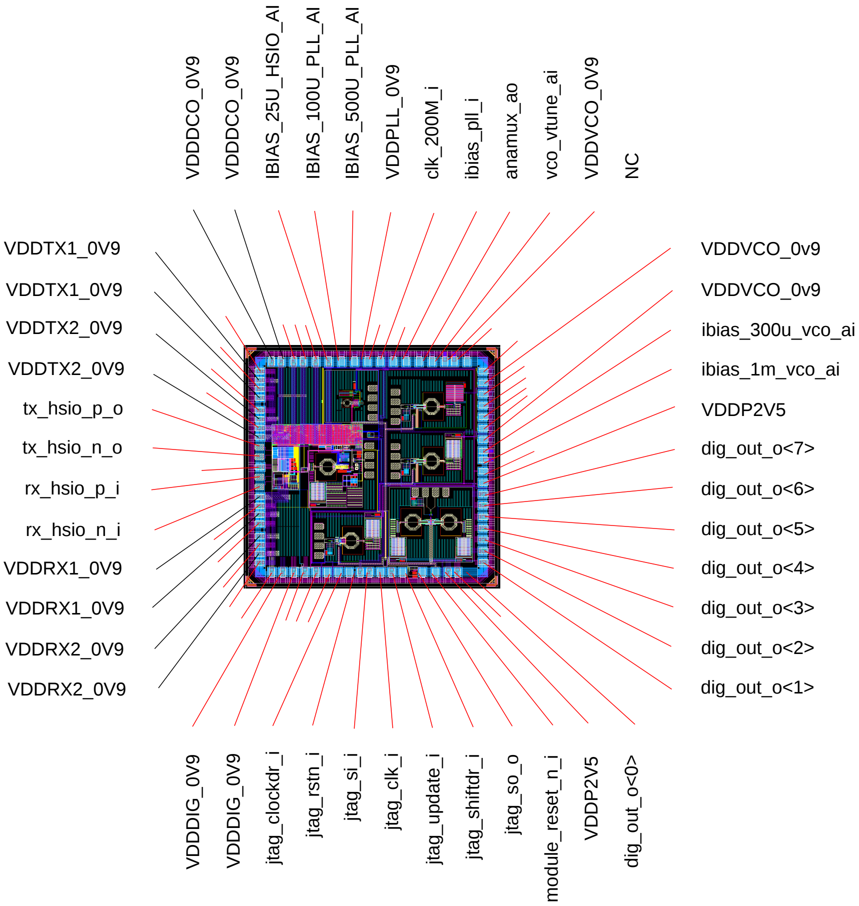

Figure 6.1: Bondplan of InnoTC as realized for the measurements of the high-speed transceiver. The proportions between the link, the bonding and the package respect the real ones.

factor of 2. The bondwires to the power supplies are roughly 3nm long each, so that the effective parasitic inductance on each power domain is roughly 1.5nH.

6.2 InnoTC Measurement Setup

Fig. 6.2 shows the measurement setup used for evaluating the performance of the transceiver. The chip is bonded to a package and then mounted on a PCB, called daughterboard. This board is itself mounted on another board, called daughterboard, which brings four different supplies to the chip: one for 2.5V digital pads, one core voltage to the digital circuit, one for the DCO and the PLL and one for the HSIO. This means that all four power domains are connected together at motherboard level, and separated only at daughterboard level. Two different versions of the daughterboard have been tried, one with a crystal reference mounted on top and one with an external clock reference of 200MHz. Since there was no significant difference between the two in measurements in loopback mode (to be defined in the next paragraph), only measurements with an external reference clock will be reported here.

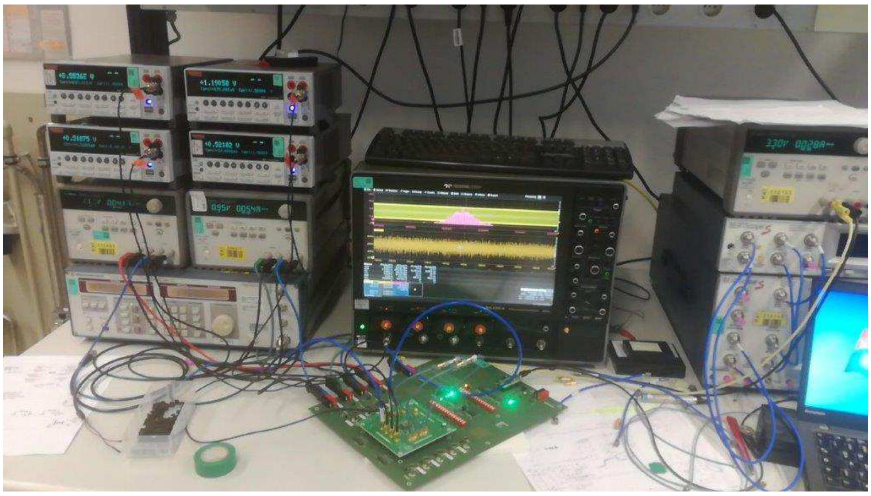

Figure 6.2: Measurement Setup for InnoTC. In the picture, motherboard, daughterboard and all voltage and current sources are visible, both for JTAG programming and for performance measurement.

Fig. 6.3 show the connections in order to operate the full link in loopback mode: the outputs of the transmitter are directly fed into the receiver inputs, after having passed through a channel composed by roughly 40cm of cables, a SMB connector and an AC coupler, which mimic the effect of the on-board capacitance needed to separate the common mode levels of transmitter and receiver. No S-parameters measurements were performed on this channel, but by comparing the vertical opening of the eye diagram at the output of the transmitter when fed into the scope with the same channel, a loss of circa 3dB at 2.5GHz can be estimated.

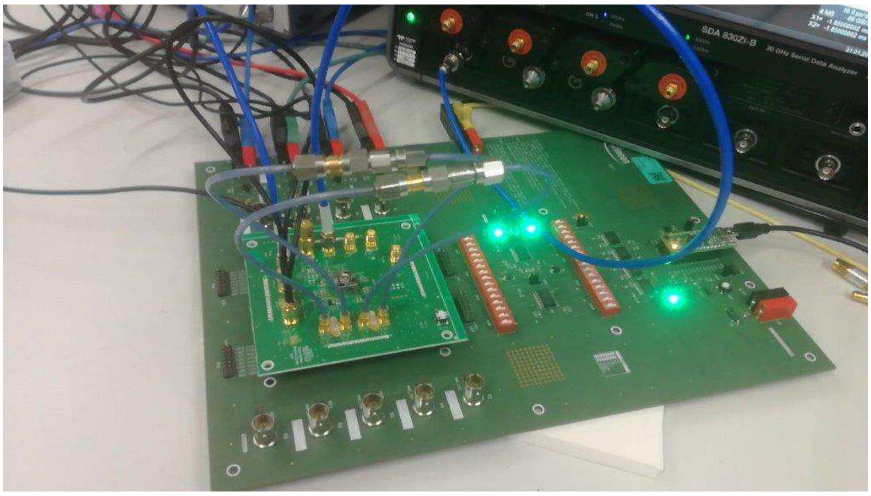

Figure 6.3: Measurement Setup for the loopback functioning mode, in which the outputs of the transmitter are directly fed into the receiver inputs of the same chip after having passed roughly 40cm of cables, a SMB connector and an AC coupler.

6.3 Full Link Measurements Results

Unfortunately, it was not possible to obtain a BER at 9.2Gbps because the PRBS checker was unable to start at this speed due to a timing issue between the receiver and the Simplified Timing Shell. This error has already been demonstrated in simulation and one sample of the chip has already undergone a FIB to solve this issue. Further measurements will be carried out on this sample to have the PRBS checker working. Despite this issue, the correctness of the analog design at this speed can be demonstrated with the correct transmission and reception of a static pattern along with the CDR locking at 1V for both transmitter and receiver supply. By fixing the PI code from the outside and sweeping it, a range of roughly 15 PI codes (circa 50ps) where the static pattern is correctly detected can be found, denoting a clean opening of the eye diagram at the receiver sampler. Tough, it would make little to no sense to draw a bathtub for such case, since the eye opening strongly depends on the transmitted bits, therefore the BER at a certain distance from the eye center would also depend on this.

The aforementioned timing issue doesn't show up at 4.6Gbps, therefore BER measurements can be performed. For all the measurements reported in the figures below, except when differently stated, the settings reported in tab. 6.1 have been used. TX Replicas indicates the number of active replicas at transmitter side, LDO Current is the bias current fed into the LDO and CTLE Res and Cap Code are indicating the amount of RC source degeneration in the CTLE (we preferred to indicate the codes and not the dB of equalization since these strongly depend on the technology corner for the same settings). The CDR bandwidth is set thanks the averaging: for a transmission speed of 4.6Gbps, the digital clock has a frequency of 115MHz, therefore an averaging of 115 has to be set.

Fig. 6.4 shows the bathtub curves for three different V_{DD}, highlighting

Table 6.1: Settings used for TX-RX loopback measurements.

TX Replicas	LDO Current[μA]	CTLE Res, Cap Code	DFE	CDR PD, BW [MHz]
12	32	3, 7	h_1=30mV, h_2=15mV	Alexander, 1

closer bathtub with decreasing supply voltage. This curves have been obtained with the following method: first of all, we run the link with the CDR enabled in order to find the PI code at which the CDR settles. Once this has been found, we sweep all PI codes up to ±16 values away from the center one and we run the BER counter for 30 seconds. This is enough for finding actual BER up to 10^{-7}, since in 30 seconds we are transmitting circa $1.5 \cdot 10^{11}$ bits so for a BER of 10^{-7} we find roughly 10000 errors. After this sweep has been done, we run the BER counter for 10 minutes setting the PI code to the values adjacent to those where we found a BER of 10^{-7}, until we don't find errors anymore. In fact, in 10 minutes we transmit roughly $3 \cdot 10^{12}$ bits, so if we don't find an error on such amount of time it's safe to say that the BER for this PI codes is 10^{-12} or lower, and for accurately measuring this BER points one should let the interface run for several hours for each one of this PI codes. Anyway, as highlighted from the bathtub curves, in all three cases an open eye diagram is clearly visible at the input of the receiver sampler. Fig. 6.5 reports the corresponding power consumption of the whole high-speed link for each power supply value, along with the power consumed by the single components of the link. For comparison's sake, in fig. 6.6 we report the power consumption of the whole link for different power supply values also for 9.2Gbps transmission speed.

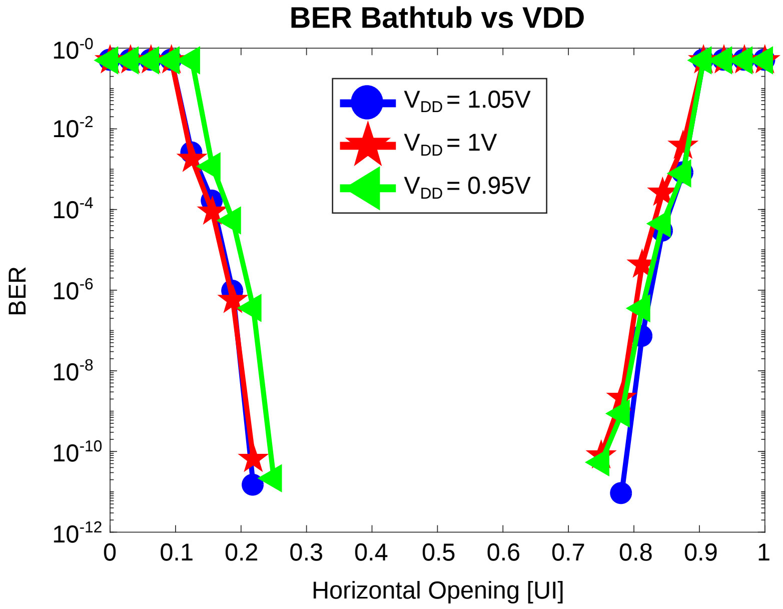

Figure 6.4: Bathtub curves at 4.6Gbps for various V_{DD} in loopback mode.

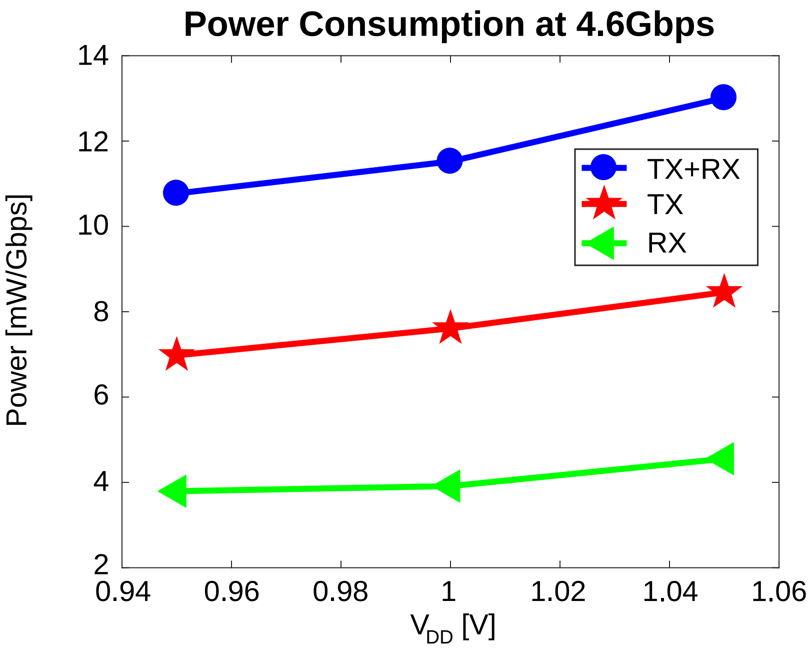

Figure 6.5: Power consumption of the whole link for different power supplies at 4.6Gbps. The power budget for the PI and the digital part (PRBS generator, PRBS checker, PD and CDR algorithm) are reported under the receiver power consumption.

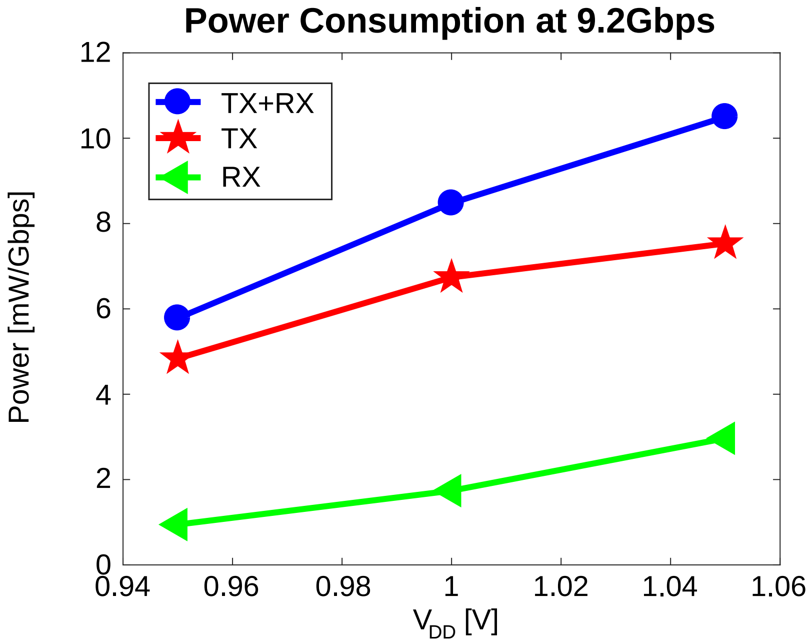

Figure 6.6: Power consumption of the whole link for different power supplies at 9.2Gbps. The power budget for the PI and the digital part (PRBS generator, PRBS checker, PD and CDR algorithm) are reported under the receiver power consumption.

Finally, some measurements have been performed in order to evaluate the performances of the two implemented phase detection methods, Alexander and the one proposed in this work. In order to evaluate the two methods, we run the link with active CDR and active BER, in order to monitor the overall performance as well. During this run, we read the PI code 200 times and count the occurrences of each PI code among these 200 reads. The results of these measurements are reported in figg. 6.7-6.11. Comparing fig. 6.7 and fig. 6.9, we can say that the two phase detection algorithms perform in a similar way when using low CDR bandwidth: in fact, both PD settle around three PI codes, with rare occurrences in two adjacent ones. The only difference is related to the center PI code at which the CDR settles: without having an eye monitor before the sampling at receiver side, it's difficult to say which one of the two is more centered. Further measurements will be done in order to understand this point.

When the bandwidth of the CDR increases, then the spread of the PI codes increases, as shown in fig. 6.9-6.11. This is due to the fact that, when increasing the bandwidth of the CDR, noise at higher frequencies will also be tracked, whereas at low CDR bandwidth the deterministic jitter seen is the one given by the CDR limit cycle itself. At the same time, when increasing the CDR bandwidth, also the BER was increasing: no error was seen when reading the codes reported in figg. 6.7-6.9, whereas for figg. 6.9-6.11 we found BER of roughly 10^{-10} and 10^{-6}, respectively. This is most probably due to the fact that when going at higher CDR bandwidth, also the number of voting averaging is reduced, therefore the outliers are not smoothed out by the voting process. The fact that the BER is higher for our phase detector than for the Alexander may indicate that the PI code at which our algorithm settles is not so well centered as the Alexander one. In order to understand at which bandwidth our algorithm stops to perform well for this settings, different CDR bandwidths have been tried out: the BER for a CDR bandwidth of 11.5MHz is roughly 10^{-10}, and as shown in fig. 6.10 the PI codes are not so spread. In fact, as already seen, the performances of the two PD algorithms should depend on the DFE settings: the more DFE is used, the more robust our PD algorithm should become and the less centered the Alexander PD should be. To check this, we should run our interface when connecting transmitter and receiver with channels where higher amount of DFE is needed.

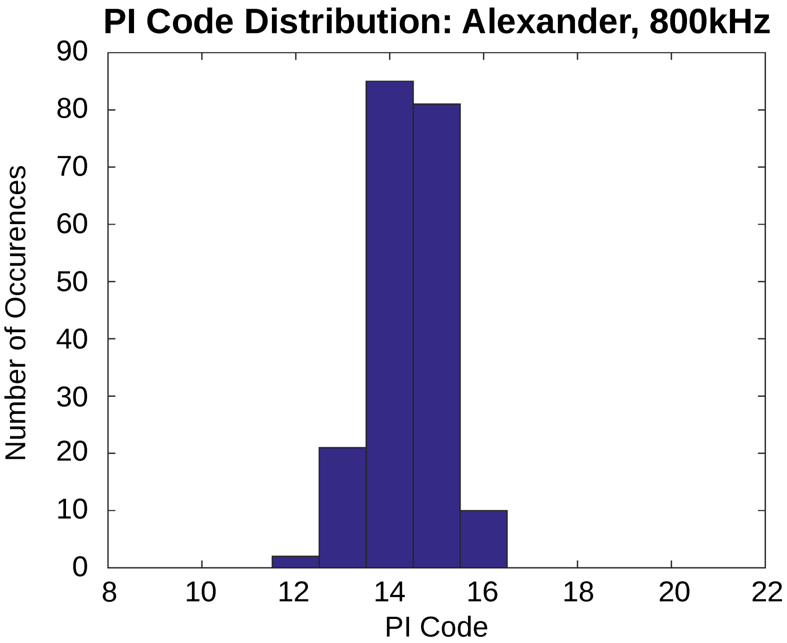

Figure 6.7: Distribution of the PI codes for Alexander phase detector with CDR bandwidth of 800kHz. The transmission speed is 4.6Gbps. The settings used for the link are the same as reported in tab. 6.1.

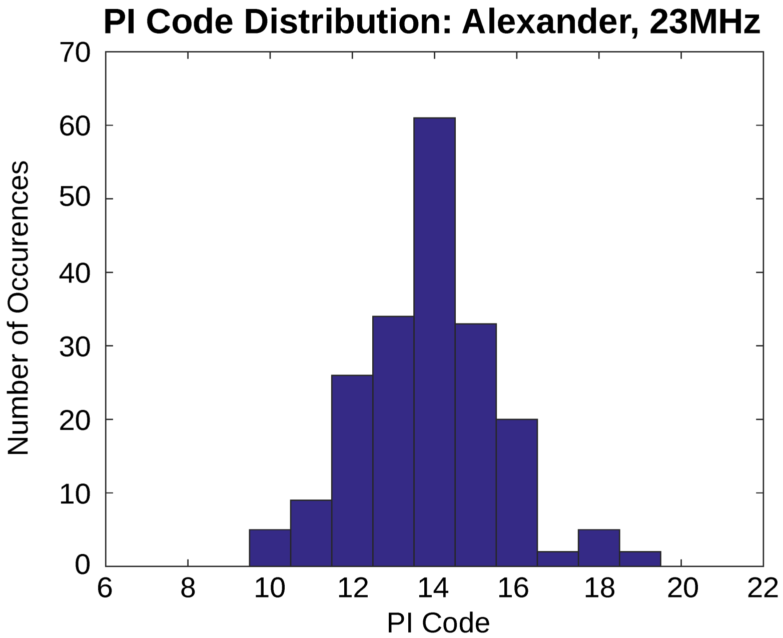

Figure 6.8: Distribution of the PI codes for Alexander phase detector with CDR bandwidth of 23MHz. The transmission speed is 4.6Gbps. The settings used for the link are the same as reported in tab. 6.1.

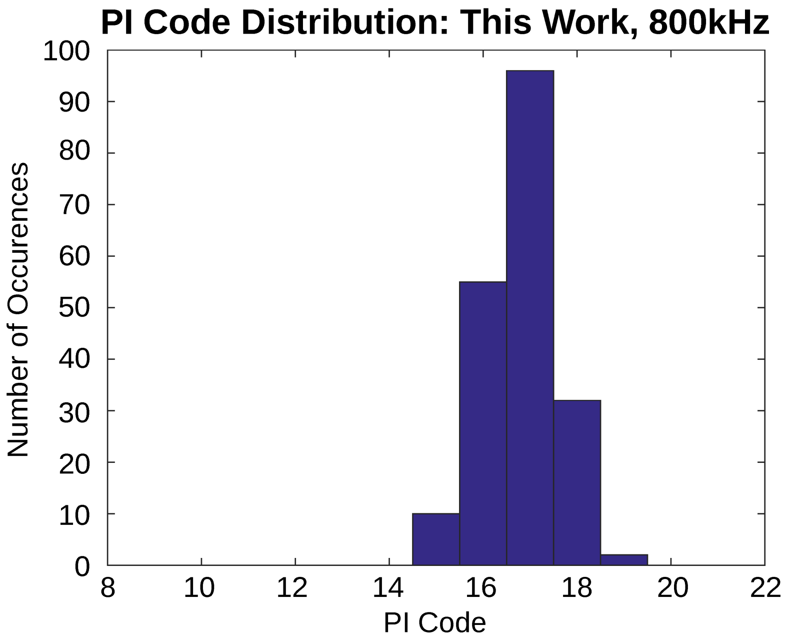

Figure 6.9: Distribution of the PI codes for phase detector proposed in this work with CDR bandwidth of 800kHz. The transmission speed is 4.6Gbps. The settings used for the link are the same as reported in tab. 6.1.

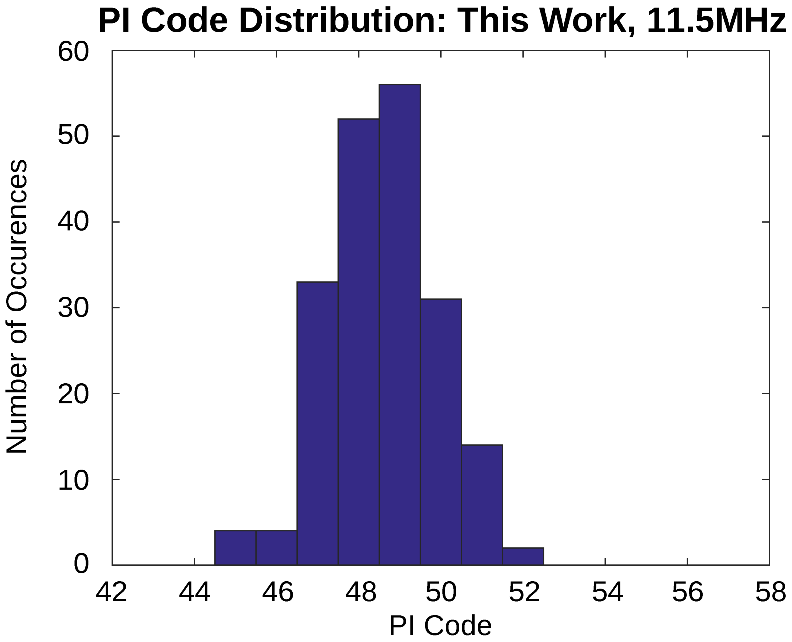

Figure 6.10: Distribution of the PI codes for phase detector proposed in this work with CDR bandwidth of 11.5MHz. The transmission speed is 4.6Gbps. The settings used for the link are the same as reported in tab. 6.1.

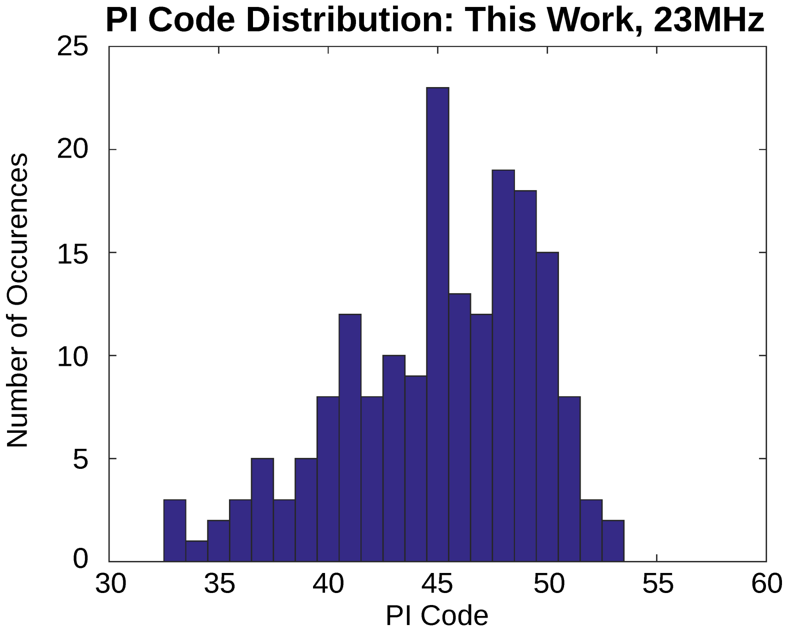

Figure 6.11: Distribution of the PI codes for phase detector proposed in this work with CDR bandwidth of 23MHz. The transmission speed is 4.6Gbps. The settings used for the link are the same as reported in tab. 6.1.

6.4 HSIO Evaluation Board

In order to address the need of channels with higher losses stated at the end of the previous chapter so to compare the performances of the two phase detection algorithms and to better investigate the equalization capabilities of our interface, a Printed Circuit Board with several channels on top has been designed and fabricated (fig. 6.12). Another goal of the measurements that will be performed with this board is to understand the influence of many bad habits in PCB design (90 degree angles, high speed signals through vias, aggressors on high speed lines) on the link performance. The channels on this board, if combined, can combine for a loss up to 20dB at 2.5GHz. The PCB consists of four layers and its dimensions are 40cmX40cm. The dielectric used is standard FR-4, which is the material used for typical automotive mass production applications due to its lower cost.

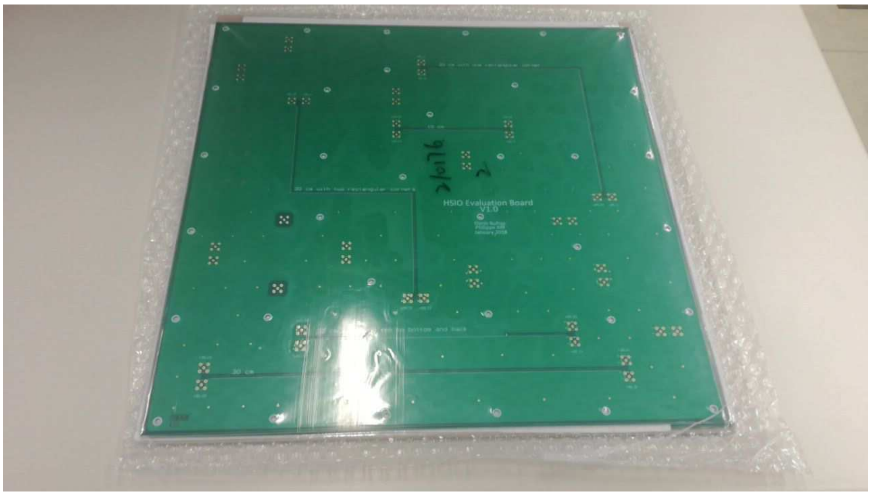

Figure 6.12: HSIO Evaluation Board.

Conclusions

In this thesis, after having introduced the frame of high-speed serial interfaces for chip-to-chip communication, we have presented the whole design flow of a 10Gbps transceiver, from system level analysis to transistor level and layout implementation. Finally, the whole transceiver has been characterized through measurements.

We started by showing procedures to define an equalization strategy at the transmitter and to evaluate its effect on the transmitted data, moving then to the transistor level design of the transmitter architecture supporting such a FFE scheme. The transmitter has been designed with a voltage-mode driver in order to save power and features FFE with 8 taps (one pre-cursor, main tap and six post-cursor taps), programmable with 16 levels. The whole architecture is divided into two power domains and the driver is supplied by an LDO with a selectable output voltage among eight different values, giving an eye diagram peak-to-peak differential height that can be tuned from 250mV to 380mV in nominal PVT conditions (V_{DD}=900mV, typical technology corner, 25°C). With a decoupling capacitance of 50pF for each power domain, the transmitter is correctly operating over a broad range of V_{DD} (from 800mV to 1.05V), technology corners (all possible combinations of slow, typical and fast technological corners) and temperature (from -40°C to 175°C).

The functioning of the transmitter has been demonstrated with transistor level simulations and with measurements performed on two separated test-chips. On the transmitter placed on the second test-chip, a whole characterization has been carried out. Moreover, an extended analysis on the effect of parasitic inductance on V_{DD} and V_{SS} domains has been illustrated, highlighting the necessity to keep the values of these inductances under control in the chip-bonding process.

We have then focused our attention on the design of the receiver, starting from a system level analysis aiming at understanding the inter-dependence between half-rate architecture, loop-unrolled DFE and CDR. To do so, we have built a Simulink model. We have then proposed a novel phase detection algorithm dedicated to CDR in systems with DFE and validated it using the aforementioned Simulink model.

The receiver architecture features a source-degenerated differential pair CTLE programmable in the source resistance (2 bits), source capacitance (3 bits) and biasing current (4 bits). Also the operating point of this CTLE can be tuned (3 bits) to find the optimal value depending on the PVT conditions and the transmitted eye opening. The CTLE can boost the frequency components around Nyquist (5GHz) of the transmitted bits up to 10dB with respect to low-frequency ones. The receiver features also a 4 post-cursor taps DFE, the first one being loop-unrolled. The current in the main tap of the summers can be trimmed with 16 levels, and so does the current flowing in the DFE taps. The drain resistances of the main summer feature 8 different trimmable levels, so that the finest DFE correction that can be programmed is 6mV. Also the receiver is divided into two power domains. With a decoupling capacitance of 50pF for each power domain, simulations show that the receiver is correctly operating over a broad range of V_{DD} (from 800mV to 1.05V), technology corners (all possible combinations of slow, typical and fast technological corners) and temperature (from -40°C to 175°C).

The functioning of the receiver and its integration with the already taped-out transmitter has then been demonstrated with transistor level and mixed-signal simulations, including also the digital part of the transceiver. These simulations show that the proposed phase detection algorithm is suited in case of low-quality crystals for frequency generation (100ppm or above). Both transmitter and receiver are also capable of working at three different transmission speeds: 2.5Gbps, 5Gbps and 10Gbps.

Simulated data point out a power consumption of 22.5mW in the least consuming corner for the transmitter, corresponding to an efficiency of 2.25 pJ/bit, in line with the state-of-the-art in the literature. As for the receiver, the power consumption is 20.5mW in nominal PVT conditions (V_{DD}=900mV, typical technology corner, 25°C), corresponding to a power efficiency of 2.05pJ/bit. The analog part of the receiver (bias distribution, CTLE, DFE, summers, comparators) consumes 0.77mW/Gbps, whereas the CDR circuitry (PI, PI Decoder, 2:40 Deserializers) 1.28mW Gbps. In the least-consuming PVT corner (V_{DD}=800mV, slow-slow technology corner, -40°C), the analog part consumes 0.5mW/Gbps, whereas the CDR circuitry 0.8mW/Gbps, which gives a total 1.3mW/Gbps.

Finally, the functioning of the whole transceiver has been demonstrated with measurements. A static pattern in loopback mode at 9.2Gbps with settling CDR has been demonstrated. A whole characterization of the link at 4.6Gbps in loopback mode has been carried out, highlighting BER lower than 10^{-12} and good agreement between simulations and measurements in terms of power consumption. An in depth analysis of the performance of the two phase detection algorithms has just started and the first results are already available in this thesis. In the next steps, more measurements will be performed using as a channel different traces placed on an already designed and produced Printed Circuit Board.

List of Publications

Results of this research work have been presented at international conferences with peer-review:

- Bandiziol, A., Grollitsch, W., Brandonisio, F., Nonis, R., Palestri, P. (2016, May). Design of a transmitter for high-speed serial interfaces in automotive micro-controller. In Information and Communication Technology, Electronics and Microelectronics (MIPRO), 2016 39th International Convention on (pp. 84-88). IEEE.
- Bandiziol, A., Grollitsch, W., Brandonisio, F., Nonis, R., Palestri, P. (2016, October). Design of a 8-taps, 10Gbps transmitter for automotive micro-controllers. In Circuits and Systems (APCCAS), 2016 IEEE Asia Pacific Conference on (pp. 321-324). IEEE.
- Bandiziol, A., Grollitsch, W., Brandonisio, F., Nonis, R., Palestri, P. (2017, June). System and transistor level analysis of an 8-taps FFE 10Gbps serial link transmitter with realistic channels and supply parasitics. In Ph. D. Research in Microelectronics and Electronics (PRIME), 2017 13th Conference on (pp. 297-300). IEEE.

and at national level in:

- Bandiziol, A., Grollitsch, W., Brandonisio, F., Nonis, R., Palestri, P., Selmi L. (2016, June). Design of a transmitter for high-speed serial interfaces in automotive micro-controller. Proceedings of GE2016, 48th Conference, Brescia, Italy.

Moreover, the following articles have been accepted and will be presented at international conferences with peer-review:

- Bandiziol, A., Grollitsch, W., Brandonisio, F., Bassi M., Nonis, R., Palestri, P. Design of a half-rate receiver for a 10Gbps automotive serial interface with 1-tap-unrolled 4-taps DFE and custom CDR algorithm. Circuits and Systems (ISCAS), 2018 International Symposium on. IEEE.
- Dazzi M., Palestri P., Rossi D., Bandiziol A., Loi I., Bellasi D., Benini L. Sub-mW multi-Gbps chip-to-chip communication Links for Ultra-Low Power IoT end-nodes (Invited Paper). Circuits and Systems (ISCAS), 2018 International Symposium on. IEEE.

Finally, research activities results have been presented at:

- Infineon Innovation Days, June 2015, Villach, Austria
- PhD Expo, 18th June 2015, Udine, Italy
- Infineon Innodays, November 2015, Villach, Austria
- CAS Day, 6th March 2016, Villa del Grumello, Italy
- PhD Expo, 26th May 2016, Udine, Italy
- PhD Expo, 25th May 2017, Udine, Italy

Acknowledgments

My first thanks goes to my tutor at University of Udine, Prof. Pierpaolo Palestri. It is hard for me to imagine ways to be a better guide than what he has been to me in these three years. His help has just been invaluable, his curiosity for all aspects of electronics has been a constant stimulus for me in trying to become a better researcher everyday.

An enormous thanks goes also to my supervisor at Infineon Technology Design Center, Dr Roberto Nonis. On top of the technical contents that he has transmitted to me in these years, the thing that I have tried to learn at most from his daily behavior is to being able to take time to listen to people even in the busiest days. Along with him, also Werner Grollitsch has taught me a lot more than high-speed serial interfaces design: humility is a much bigger virtue, especially when shown towards a guy that knows nearly nothing of a field in which you have several years of experience.

The fact that I have been able to fully concentrate on this work along this three year span has been possible only thanks to my family: nothing of this would have been possible without your continuous support. Examples are much more powerful instruments than words, and you all have given me plenty of them in my whole life. My biggest hope is to be every day even just one tenth as inspirational as you have been, and are, to me. A big thank goes also to Ginevra: it would have been easy for you to say that I was a little bit too focused on work in the last months, but instead you decided to fit to my time as water in a bowl. I will remember it when your bowl will be smaller.

Having one office with a great atmosphere is not for everybody. Having two of them has been just amazing. I have found all the support and good vibes that I needed in these years at CIS in Infineon Technologies, and when I talk to my friends or family about how I feel when I enter my office they almost don't believe me. It's hard to find such a mix of experience, love for the job, dedication and easy-going attitude. I want to particularly thank Fabio, Francesco, Andrea, Matteo, Dmytro and Thomas for the patience they have demonstrated to have with me every day. A big thanks goes also to Stefan and Andrea for their work done with the layout. Obviously, without them there would be nothing to measure in the lab. And so thanks also to Florin for the first round of measurements, and sorry for the next ones that will come. At the same time, I will never forget my time in Selmilab: Paolo, Stefano, Enrico, Giulia, Owes, Tommaso, Andrea, Alessandro, Julian, Martino, Patrick, Daniel, Federico, I have memories for each one of you. Thanks for the laughs, for the help, for the tips. If these four years are gone so fast and it seems to me that I entered that lab for the first time yesterday, it's mainly because of you.

Bibliography

[1] Altera, "FPGAs at 40 nm and >10 Gbps: Jitter-, Signal Integrity-, Power-, and Process-Optimized Transceivers," Tech. Rep. April, 2013.

[2] I. C. Society, *IEEE Std 802.3^{TM}-2015 (Revision of IEEE Std 802.3-2012), IEEE Standard for Ethernet*. 2015.

[3] D. Law, D. Dave, J. D'Ambrosia, M. Hajduczenia, M. Laubach, and S. Carlson, "Evolution of Ethernet Standards in the IEEE 802.3 Working Group," *IEEE Communications Magazine*, no. 8, pp. 88–96, 2013.

[4] PCI, "PCI Express ® Base Specification Revision 3.0," 2010.

[5] Hypertransport, *HyperTransport 3.1*. 2008.

[6] CPRI, *CPRI Specification V7.0*. 2015.

[7] Interlaken, *Interlaken Protocol Definition*. 2008.

[8] OIF, *Common Electrical I/O (CEI) - Electrical and Jitter Interoperability agreements for 6G+bps, 11G+bps and 25G+bps I/O*. 2014.

[9] SATA-IO, *Serial ATA 3.0*. 2009.

[10] RapidIO, *RapidIO Interconnect Specification 4.0*. 2016.

[11] K. Chang, G. Zhang, and C. Borrelli, "Evolution of Wireline Transceiver Standards," *IEEE Solid-State Circuits Magazine*, no. 4, pp. 47–52, 2015.

[12] R. Mellitz, A. Ran, M. P. Li, and V. Ragavassamy, "Channel Operating Margin (COM): Evolution of Channel Specifications for 25 Gbps and Beyond," in *DesignCon*, pp. 1–20, 2013.

[13] X. D. M. M. F. R. W. J. G. Zhang, "Relating COM to Familiar S-Parameter Parametric to Assist 25Gbps System Design," in *DesignCon*, 2014.

[14] T. C. Carusone, "Introduction to Digital I/O," *IEEE Solid-State Circuits Magazine*, no. 4, pp. 14–22, 2015.

[15] B. Razavi, "Historical Trends in Wireline Communications," *IEEE Solid-State Circuits Magazine*, no. 4, pp. 42–46, 2015.

[16] S. Palermo, "High-Speed Serial I/O Design for Channel- Limited and Power-Constrained Systems," in *CMOS Nanoelectronics Analog and RF VLSI Circuits*, 2011.

[17] A. X. Widmer and P. A. Franaszek, "A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code," *IBM J. Res. Develop*, vol. 27, no. 5, pp. 440–451, 1983.

[18] M. P. Li, *Jitter, Noise, and Signal Integrity at High-Speed*. 2007.

[19] B. Razavi, *Design of integrated circuits for optical communications*. 2012.

[20] B. Casper, "Clocking Wireline Systems," *IEEE Solid-State Circuits Magazine*, no. 4, pp. 32–41, 2015.

[21] R. Kreienkamp, U. Langmann, C. Zimmermann, T. Aoyama, and H. Siedhoff, "A 10-Gb/s CMOS Clock and Data Recovery Circuit With an Analog Phase Interpolator," *IEEE Journal of Solid State Circuits*, vol. 40, no. 3, pp. 736–743, 2005.

[22] S. Hu, C. Jia, K. Huang, C. Zhang, X. Zheng, and Z. Wang, "A 10Gbps CDR based on Phase Interpolator for Source Synchronous Receiver in 65nm CMOS," in *IEEE International Symposium on Circuits and Systems*, pp. 309–312, 2012.

[23] S. S. Iyer and T. Kirihata, "Three-Dimensional Integration," *IEEE Solid-State Circuits Magazine*, no. 4, pp. 63–74, 2015.

[24] J. Fan, X. Ye, J. Kim, B. Archambeault, and A. Orlandi, "Signal Integrity Design for High-Speed Digital Circuits: Progress and Directions," *IEEE Trabsactions on Electromagnetic Compatibility*, vol. 52, no. 2, pp. 392–400, 2010.

[25] P. K. Hanumolu, G.-y. Wei, and U.-k. Moon, "Equalizers for High-Speed Serial Links," *International Journal of High Speed Electronics and Systems*, vol. 15, no. 2, pp. 429–458, 2005.

[26] J. F. Bulzacchelli, "Equalization for Electrical Links: Current Design Techniques and Future Directions," *IEEE Solid-State Circuits Magazine*, no. 4, pp. 23–31, 2015.

[27] T. Kawamoto, "Multi-Standard 185fs 0.3-to-28Gb/s 40dB Backplace Signal Conditioner with Adaptive Pattern-Match 36-Tap DFE and Data-Rate-Adjustment PLL in 28nm CMOS," in *IEEE International Solid-State Circuits Conference*, pp. 54–56, 2015.

[28] E. Mammei, F. Loi, F. Radice, A. Dati, M. Bruccoleri, M. Bassi, and A. Mazzanti, "Analysis and Design of a Power-Scalable Continuous-Time FIR Equalizer for 10 Gb/s to 25 Gb/s Multi-Mode Fiber EDC in 28 nm LP CMOS," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 12, pp. 3130–3140, 2014.

[29] J. F. Bulzacchelli, T. J. Beukema, D. W. Storaska, D. R. Hanson, P.-h. Hsieh, S. V. Rylov, D. Furrer, D. Gardellini, A. Prati, T. Morf, V. Sharma, R. Kelkar, H. A. Ainspan, W. R. Kelly, L. R. Chieco, G. A. Ritter, J. A. Sorice, J. D. Garlett, R. Callan, P. Buchmann, M. Kossel, and D. J. Friedman, "A 28-Gb/s 4-Tap FFE/15-Tap DFE Serial Link Transceiver in 32-nm SOI CMOS Technology," *IEEE Journal of Solid State Circuits*, vol. 47, no. 12, pp. 3232–3248, 2012.

[30] A. Agrawal, J. F. Bulzacchelli, T. O. Dickson, Y. Liu, J. A. Tierno, and D. J. Friedman, "A 19-Gb/s Serial Link Receiver With Both 4-Tap FFE and 5-Tap DFE Functions in 45-nm SOI CMOS," *IEEE Journal of Solid State Circuits*, vol. 47, no. 12, pp. 3220–3231, 2012.

[31] T. Toifl, P. Buchmann, T. Beukema, M. Beakes, P. A. Francese, C. Menolfi, M. Kossel, L. Kull, and T. Morf, "A 3.5pJ/Bit 8-Tap-Feed-Forward 8-Tap-Decision Feedback Digital Equalizer for 16Gb/s I/Os," in *IEEE European Solid State Circuits Conference (ESSCIRC)*, pp. 455–458, 2014.

[32] H. Yueksel, L. Kull, A. Burg, M. Braendli, P. Buchmann, P. A. Francese, C. Menolfi, M. Kossel, T. Morf, T. M. Andersen, D. Luu, and T. Toifl, "A 3.6 pJ/b 56 Gb/s 4-PAM Receiver with 6-Bit Ti-SAR ADC and Quarter-Rate Speculative 2-Tap DFE in 32 nm CMOS," in *IEEE European Solid State Circuits Conference (ESSCIRC)*, pp. 148–151, 2015.

[33] S. Rylov, T. Beukema, Z. Toprak-deniz, T. Toifl, Y. Liu, A. Agrawal, P. Buchmann, A. Rylyakov, M. Beakes, B. Parker, and M. Meghelli, "A 25Gb/s ADC-Based Serial Line Receiver in 32nm CMOS SOI," in *IEEE International Solid-State Circuits Conference*, pp. 56–58, 2016.

[34] S. Gondi and B. Razavi, "Equalization and Clock and Data Recovery Techniques for 10-Gb/s CMOS Serial-Link Receivers," *IEEE Journal of Solid-State Circuits*, vol. 42, no. 9, pp. 1999–2011, 2007.

[35] A. Healey, "Noise considerations for 40/100GBASE-CR4/10," in *IEEE P802.3ba Task Force Meeting*, no. July, 2009.

[36] Infineon, "Annual Report 2016," tech. rep., 2016.

[37] R. Tummala, "New Era of Automotive Electronics as the most complex Electronic System of Our Life Time," in *IEEE Electrical Design of Advanced Packaging and Systems (EDAPS)*, 2016.

[38] M. Traub and A. Maier, "Future Automotive Architecture and the Impact of IT Trends," *IEEE Software*, no. May/June, pp. 27–32, 2017.

[39] J. Endo, "Wireless Communication In and Around the Car: Status and Outlook - ES3: High-Speed Communications on 4 Wheels: What's in Your Next Car?," in *IEEE International Solid-State Circuits Conference*, 2013.

[40] C. Schmidt, "Automotive Electronics - Enabling the future of individual mobility," in *IEEE International Electronic Devices Meeting (IEDM)*, pp. 3–8, 2007.

[41] A. Cristofoli, P. Palestri, L. Selmi, and N. Da Dalt, "Efficient Statistical Simulation of Intersymbol Interference and Jitter in High-Speed Serial Interfaces," *IEEE Transactions on Components, Packaging and Manifacturing Technology*, vol. 4, no. 3, pp. 480–489, 2014.

[42] M. Kossel, J. Weiss, C. Menolfi, P. Buchmann, T. Morf, T. Toifl, and M. Schmarz, "A T-Coil-Enhanced 8.5Gb/s High-Swing SST Transmitter in 65nm Bulk CMOS With <-16dB Return Loss Over 10GHz Bandwidth," *IEEE Journal of Solid State Circuits*, vol. 43, no. 12, pp. 2905–2919, 2008.

[43] J. G. Proakis, *Digital Communications*. 2008.

[44] Ansys, "Ansys Electronic Desktop 16.0 Manual," 2016.

[45] Ansys, "Ansys SIWave 16.0 Manual," 2016.

[46] A. Cossettini, A. Cristofoli, W. Grollitsch, L. Alves, R. Nonis, L. D. Ricca, P. Palestri, and L. Selmi, "Design, characterization and signal integrity analysis of a 2.5 Gb/s high-speed serial interface for automotive applications overarching the chip/PCB wall," in *2015 IEEE 1st International Forum on Research and Technologies for Society and Industry Leveraging a better tomorrow (RTSI)*, 2015.

[47] J. F. Bulzacchelli, M. Meghelli, S. V. Rylov, W. Rhee, A. Rylyakov, H. A. Ainspan, B. D. Parker, M. P. Beakes, A. Chung, T. J. Beukema, P. K. Pepeljugoski, L. Shan, Y. H. Kwark, S. Gowda, and D. J. Friedman, "A 10-Gb/s 5-tap DFE/4-tap FFE transceiver in 90-nm CMOS technology," *IEEE Journal of Solid State Circuits*, vol. 41, no. 12, pp. 2885–2900, 2007.

[48] T. Beukema, M. Sorna, K. Selander, S. Zier, B. L. Ji, P. Murfet, J. Mason, W. Rhee, H. Ainspan, B. Parker, and M. Beakes, "A 6.4-Gb/s CMOS SerDes core with feed-forward and decision-feedback equalization," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 12, pp. 2633–2644, 2005.

[49] S. Saxena, R. K. Nandwana, and P. K. Hanumolu, "A 5 Gb/s Energy-Efficient Voltage-Mode Transmitter Using Time-Based De-Emphasis," *IEEE Journal of Solid State Circuits*, vol. 49, no. 8, pp. 1827–1836, 2014.

[50] C. Menolfi, T. Toifl, P. Buchmann, M. Kossel, T. Morf, J. Weiss, and M. Schmatz, "A 16Gb/s Source-Series Terminated Transmitter in 65nm CMOS SOI," in *IEEE International Solid-State Circuits Conference*, pp. 446–448, 2007.

[51] W.-s. Choi, G. Shu, M. Talegaonkar, Y. Liu, D. Wei, L. Benini, and P. K. Hanumolu, "A 0.45-to-0.7V 1-tp-6Gb/s 0.29-to-0.58pJ/b Source-Synchronous Transceiver Using Automatic Phase Calibration in 65nm CMOS," in *IEEE International Solid-State Circuits Conference*, pp. 66–68, 2015.

[52] T. Shibasaki, T. Danjo, Y. Ogata, Y. Sakai, H. Miyaoka, F. Terasawa, M. Kudo, H. Kano, A. Matsuda, S. Kawai, T. Arai, H. Higashi, N. Naka, H. Yamaguchi, T. Mori, Y. Koyanagi, and H. Tamura, "A 56Gb/s NRZ-Electrical 247mW/lane Serial-Link Transceiver in 28nm CMOS," in *IEEE International Solid-State Circuits Conference*, pp. 64–66, 2016.

[53] M. J. Myjak, J. G. Delgado-Frias, and S. Kwang Jeon, "An Energy-Efficient Differential Flip-Flop for Deeply Pipelined Systems," in *Midwest Symposium on Circuistemsts and Systems*, no. September 2006, 2006.

[54] G. Shu, W.-s. Choi, S. Saxena, S.-j. Kim, M. Talegaonkar, R. Nandwana, A. Elkholy, T. Nandi, and P. K. Hanumolu, "A 16Mb/s-to-8Gb/s 14.1-to-5.9pJ/b Source Synchronous Transceiver Using DFVS and Rapid On/Off in 65nm CMOS," in *IEEE International Solid-State Circuits Conference*, pp. 398–400, 2016.

[55] H. Hentzell, S.-t. Persson, H. Hesselbom, B. Lofstedt, and M. Hansen, "Techniques For Reducing Switching Noise In High Speed Digital Systems," in *IEEE International ASIC Conference and Exhibit*, pp. 21–24, 1995.

[56] G. Khanna, R. Chandel, and A. K. Chandel, "Impact of Skew and Jitter on the Performance of VLSI Interconnects," in *IEEE Asia-Pacific Conference on Circuits and Systems (APCCAS)*, pp. 1223–1226, 2010.

[57] A. Bandiziol, W. Grollitsch, F. Brandonisio, R. Nonis, and P. Palestri, "Design of a 8-taps, 10Gbps transmitter for automotive micro-controllers," in *IEEE Asia-Pacific Conference on Circuits and Systems (APCCAS)*, pp. 6–9, 2016.

[58] H. Kimura, P. M. Aziz, T. Jing, A. Sinha, S. P. Kotagiri, R. Narayan, H. Gao, P. Jing, G. Hom, A. Liang, E. Zhang, A. Kadkol, R. Kothari, G. Chan, Y. Sun, B. Ge, J. Zeng, K. Ling, M. C. Wang, A. Malipatil, L. Li, C. Abel, and F. Zhong, "A 28 Gb/s 560 mW Multi-Standard SerDes With Single-Stage Analog Front-End and 14-Tap Decision Feedback Equalizer in 28 nm CMOS," *IEEE Journal of Solid State Circuits*, vol. 49, no. 12, pp. 3091–3103, 2014.

[59] S. Shahramian, B. Dehlaghi, and A. C. Carusone, "A 16Gb/s 1 IIR + 1 DT DFE Compensating 28dB Loss with Edge-Based Adaptation Converging in 5us," in *IEEE International Solid-State Circuits Conference*, pp. 410–412, 2016.

[60] B. Leibowitz, R. Farjad, T. Greer, and V. Stojanovic, "A 7.5Gb/s 10-Tap DFE Receiver with First Tap Partial Response, Spectrally Gated Adaptation, and 2nd-Order Data-Filtered CDR," in *IEEE International Solid-State Circuits Conference*, pp. 228–220, 2007.

[61] T. Norimatsu, T. Kawamoto, K. Kogo, N. Kohmu, F. Yuki, N. Nakajima, T. Muto, J. Nasu, T. Komori, H. Koba, T. Usugi, T. Hokari, T. Kawamata, Y. Ito, S. Umai, M. Tsuge, T. Yamashita, M. Hasegawa, and K. Higeta, "A 25Gb/s Multistandard Serial Link Transceiver for 50dB-Loss Copper Cable in 28nm CMOS," in *IEEE International Solid-State Circuits Conference*, pp. 60–62, 2016.

[62] S. Shahramian and A. C. Carusone, "A 0.41 pJ/bit 10 Gb/s Hybrid 2-IIR and 1 Discrete-Time DFE Tap in 28 nm-LP CMOS," *IEEE Journal of Solid State Circuits*, vol. 50, no. 7, pp. 1722–1735, 2015.

[63] S. Shahramian, H. Yasotharan, and A. C. Carusone, "Decision Feedback Equalizer Architectures With Multiple Continuous-Time Infinite Impulse Response Filters," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 59, no. 6, pp. 326–330, 2012.

[64] G. R. Gangasani, C.-m. Hsu, J. F. Bulzacchelli, T. Beukema, W. Kelly, H. H. Xu, D. Freitas, A. Prati, D. Gardellini, R. Reutemann, G. Cervelli, J. Hertle, M. Baecher, J. Garlett, P.-a. Francese, J. F. Ewen, D. Hanson, D. W. Storaska, and M. Meghelli, "A 32 Gb/s Backplane Transceiver With On-Chip AC-Coupling and Low Latency CDR in 32 nm SOI CMOS Technology," *IEEE Journal of Solid State Circuits*, vol. 49, no. 11, pp. 2474–2489, 2014.

[65] F. Zhong, S. Quan, W. Liu, P. Aziz, T. Jing, J. Dong, C. Desai, H. Gao, M. Garcia, G. Hom, T. Huynh, H. Kimura, R. Kothari, L. Li, C. Liu, S. Lowrie, K. Ling, A. Malipatil, R. Narayan, T. Prokop, C. Palusa, A. Rajashekara, A. Sinha, C. Zhong, and E. Zhang, "A 1.0625-14.025 Gb/s Multi-Media Transceiver With Full-Rate Source-Series-Terminated Transmit Driver and Floating-Tap Decision-Feedback Equalizer in 40nm CMOS," *IEEE Journal of Solid State Circuits*, vol. 46, no. 12, pp. 3126–3139, 2011.

[66] S. Parikh, T. Kao, Y. Hidaka, J. Jiang, A. Toda, S. Mcleod, W. Walker, Y. Koyanagi, T. Shibuya, and J. Yamada, "A 32Gb/s Wireline Receiver with a Low-Frequency Equalizer, CTLE and 2-Tap DFE in 28nm CMOS," in *IEEE International Solid-State Circuits Conference*, pp. 28–30, 2013.

[67] J. Han, Y. Lu, N. Sutardja, K. Jung, and E. Alon, "A 60Gb/s 173mW Receiver Frontend in 65nm CMOS Technology," in *2015 VLSI Circuits Digest of Technical Papers*, pp. 230–231, 2015.

[68] A. Manian and B. Razavi, "A 40-Gb/s 9.2-mW CMOS Equalizer," in *2015 VLSI Circuits Digest of Technical Papers*, pp. 226–227, 2015.

[69] U. Singh, A. Garg, B. Raghavan, N. Huang, H. Zhang, Z. Huang, A. Momtaz, and J. Cao, "A 780mW 4X28Gb/s nsceiver for 100GbE Gearbox PHY in 40nm CMOS," *IEEE Journal of Solid State Circuits*, vol. 49, no. 12, pp. 3116–3129, 2014.

[70] P. A. Francese, T. Toifl, M. Braendli, C. Menolfi, M. Kossel, T. Morf, L. Kull, T. Meyer Andersen, H. Yueksel, A. Cevrero, and D. Luu, "Continuous-Time Linear Equalization with Programmable Active-Peaking Transistor Arrays in a 14nm FinFET 2mW/Gb/s 16Gb/s 2-Tap Speculative DFE Receiver," in *IEEE International Solid-State Circuits Conference*, pp. 186–188, 2015.

[71] P. A. Francese, C. Menolfi, M. Kossel, T. Morf, L. Kull, A. Cevrero, H. Yueksel, I. Oezkaya, D. Luu, and T. Toifl, "A 30Gb/s 0.8pJ/b 14nm FinFET Receiver Data-Path," in *IEEE International Electronic Devices Meeting (IEDM)*, pp. 408–410, 2016.

[72] B. Kim, Y. Liu, T. O. Dickson, J. F. Bulzacchelli, and D. J. Friedman, "A 10-Gb/s Compact Low-Power Serial I/O With DFE-IIR Equalization in 65-nm CMOS," *IEEE Journal of Solid State Circuits*, vol. 44, no. 12, pp. 3526–3538, 2009.

[73] O. Elhadidy and S. Palermo, "A 10 Gb/s 2-IIR-Tap DFE Receiver with 35 dB Loss Compensation in 65-nm CMOS," in *2013 VLSI Circuits Digest of Technical Papers*, pp. 272–273, 2013.

[74] A. Emami-neyestanak, A. Varzaghani, J. F. Bulzacchelli, A. Rylyakov, C.-k. K. Yang, and D. J. Friedman, "A 6.0-mW 10.0-Gb/s Receiver With Switched-Capacitor Summation DFE," *IEEE Journal of Solid State Circuits*, vol. 42, no. 4, pp. 889–896, 2007.

[75] Y.-c. Huang and S.-i. Liu, "A 6Gb/s Receiver with 32.7dB Adaptive DFE-IIR Equalization," in *IEEE International Solid-State Circuits Conference*, pp. 68–69, 2011.

[76] S. Son, H.-s. Kim, M.-j. Park, K. Kim, E.-h. Chen, B. Leibowitz, and J. Kim, "A 2.3-mW, 5-Gb/s Low-Power Decision-Feedback Equalizer Receiver Front-End and its Two-Step, Minimum Bit-Error-Rate Adaptation Algorithm," *IEEE Journal of Solid State Circuits*, vol. 48, no. 11, pp. 2693–2704, 2013.

[77] H. Miyaoka, F. Terasawa, M. Kudo, H. Kano, A. Matsuda, N. Shirai, S. Kawai, T. Arai, Y. Ide, K. Terashima, H. Higashi, T. Higuchi, and N. Naka, "A 28-Gb/s 4.5-pJ/bit Transceiver With 1-Tap Decision Feedback Equalizer in 28-n CMOS," in *IEEE Asian Solid-State Circuits Conference*, pp. 1–4, 2015.

[78] A. Bandiziol, W. Grollitsch, F. Brandonisio, R. Nonis, and P. Palestri, "Design of a transmitter for high-speed serial interfaces in automotive microcontroller," in *IEEE International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, pp. 90–94, 2016.

[79] R. Navid, E.-h. Chen, M. Hossain, B. Leibowitz, J. Ren, C.-h. A. Chou, B. Daly, B. Su, S. Li, M. Shirasgaonkar, F. Heaton, J. Zerbe, and J. Eble, "A 40 Gb/s Serial Link Transceiver in 28 nm CMOS Technology," *IEEE Journal of Solid State Circuits*, vol. 50, no. 4, pp. 814–827, 2015.

[80] J. Han, N. Sutardja, Y. Lu, and E. Alon, "Design Techniques for a 60-Gb/s 288-mW NRZ Transceiver With Adaptive Equalization and Baud-Rate Clock and Data Recovery in 65-nm CMOS Technology," *IEEE Journal of Solid State Circuits*, no. 99, pp. 1–12, 2017.

[81] Synopsis, *Synopsis Verilog Compiler Simulator 2017.03 Manual*. 2017.

[82] B. Razavi, "Challenges in the Design of High-Speed Clock and Data Recovery Circuits," *IEEE Communications Magazine*, vol. 40, no. 8, pp. 94–101, 2002.

[83] J. Poulton, R. Palmer, A. M. Fuller, T. Greer, J. Eyles, W. J. Dally, and M. Horowitz, "A 14-mW 6.25-Gb/s Transceiver in 90-nm CMOS," *IEEE Journal of Solid State Circuits*, vol. 42, no. 12, pp. 2745–2757, 2007.

[84] T. Musah, J. E. Jaussi, G. Balamurugan, S. Hyvonen, T.-c. Hsueh, G. Keskin, S. Shekhar, J. Kennedy, S. Sen, R. Inti, M. Mansuri, M. Leddige, B. Horine, C. Roberts, R. Mooney, and B. Casper, "A 4-32 Gb/s Bidirectional Link With 3-Tap FFE/6-Tap DFE and Collaborative CDR in 22nm CMOS," *IEEE Journal of Solid State Circuits*, vol. 49, no. 12, pp. 1–12, 2014.

[85] V. Balan, O. Oluwole, G. Kodani, C. Zhong, R. Dadi, A. Amin, and A. Ragab, "A 15–22 Gbps Serial Link in 28 nm CMOS With Direct DFE," *IEEE Journal of Solid State Circuits*, vol. 49, no. 12, pp. 3104–3115, 2014.

[86] M.-s. Chen and C.-k. K. Yang, "A 50–64 Gb/s Serializing Transmitter With a 4-Tap, LC-Ladder-Filter-Based FFE in 65 nm CMOS Technology," *IEEE Journal of Solid State Circuits*, vol. 50, no. 8, pp. 1903–1916, 2015.

[87] T. Anand, M. Talegaonkar, A. Elkholy, S. Saxena, A. Eishazly, and P. K. Hanumolu, "Ultra-High-Speed Wireline Transceivers and Energy-Efficient Links," in *IEEE International Solid-State Circuits Conference*, pp. 64–66, 2015.

[88] M.-s. Chen, Y.-n. Shih, C.-l. Lin, H.-w. Hung, and J. Lee, "A Fully-Integrated 40-Gb/s Transceiver in 65-nm CMOS Technology," *IEEE Journal of Solid State Circuits*, vol. 47, no. 3, pp. 627–640, 2012.

[89] A. Cristofoli, *Analysis and Design of High Speed Serial Interfaces for Automotive Applications*. PhD thesis, 2014.

[90] T. O. Dickson, Y. Liu, A. Agrawal, J. F. Bulzacchelli, H. A. Ainspan, Z. Toprak-deniz, B. D. Parker, M. P. Beakes, M. Meghelli, and D. J. Friedman, "A 1.8pJ/bit 16X16Gb/s Source-Synchronous Parallel Interface in 32nm SOI CMOS with Receiver Redundancy for Link Recalibration," *IEEE Journal of Solid State Circuits*, vol. 51, no. 8, pp. 1744–1755, 2016.

[91] A. A. Hafez, M.-s. Chen, and C.-k. K. Yang, "A 32–48 Gb/s Serializing Transmitter Using Multiphase Serialization in 65 nm CMOS Technology," *IEEE Journal of Solid State Circuits*, vol. 50, no. 3, pp. 763–775, 2015.

[92] W. Jia, B. Nikolic, V. Stojanovic, J. K.-s. Chiu, and M. M.-t. Leung, "Improved Sense-Amplifier-Based Flip-Flop: Design and Measurements," *IEEE Journal of Solid State Circuits*, vol. 35, no. 6, pp. 876–884, 2000.

[93] J. W. Jung and B. Razavi, "A 25-Gb/s 5-mW CMOS CDR/Deserializer," *IEEE Journal of Solid State Circuits*, vol. 48, no. 3, pp. 684–697, 2013.

[94] K. Kaviani, A. Amirkhany, C. Huang, P. Le, C. Madden, K. Saito, K. Sano, V. Murugan, W. Beyene, K. Chang, and C. Yuan, "A 0.4mW/Gb/s 16Gb/s Near-Ground Receiver Front- End with Replica Transconductance Termination Calibration," in *IEEE International Solid-State Circuits Conference*, vol. 6, pp. 152–153, 2012.

[95] J. Lee, P. Chiang, and C. Weng, "56Gb/s PAM4 and NRZ SerDes Transceivers in 40nm CMOS," in *2015 VLSI Circuits Digest of Technical Papers*, pp. 118–119, 2015.

[96] A. Manian and B. Razavi, "A 40Gb/s 14mW CMOS Wireline Receiver," in *IEEE International Solid-State Circuits Conference*, 2016.

[97] M. Muller and K. Mulller, "Timing Recovery in Digital Synchronous Data Receivers," *IEEE Transactions on Communications*, vol. C, no. 5, pp. 516–531, 1976.

[98] B. Raghavan, D. Cui, U. Singh, D. Pi, A. Vasani, Z. C. Huang, A. Momtaz, and J. Cao, "A Sub-2 W 39.8-44.6 Gb/S Transmitter and Receiver Chipset With SFI-5.2 Interface in 40 nm CMOS," *IEEE Journal of Solid State Circuits*, vol. 48, no. 12, pp. 3219–3228, 2013.

[99] W. Sansen, "Minimum Power in Analog Amplifying Blocks," *IEEE Solid-State Circuits Magazine*, no. 4, pp. 83–89, 2015.

[100] J. Savoj, S. Member, and B. Razavi, "A 10-Gb/s CMOS Clock and Data Recovery Circuit with a Half-Rate Linear Phase Detector," *IEEE Journal of Solid State Circuits*, vol. 36, no. 5, pp. 761–767, 2001.

[101] S. Saxena, G. Shu, R. K. Nandwana, M. Talegaonkar, A. Elkholy, T. Anand, W.-s. Choi, and P. K. Hanumolu, "A 2.8 mW/Gb/s, 14 Gb/s Serial Link Transceiver," *IEEE Journal of Solid State Circuits*, vol. 52, no. 5, pp. 1399–1411, 2017.

[102] H. Tamura, "Looking to the Future," *IEEE Solid-State Circuits Magazine*, no. 4, pp. 53–62, 2015.

[103] Y. Wang, B. Afshar, L. Ye, V. C. Gaudet, and A. M. Niknejad, "Design of a Low Power, Inductorless Wideband Variable-Gain Amplifier for High-Speed Receiver Systems," *IEEE Transactions on Circuits and Systems I: Reguar Papers*, vol. 59, no. 4, pp. 696–707, 2012.

List of Publications

Results of this research work have been presented at international conferences with peer-review:

- Bandiziol, A., Grollitsch, W., Brandonisio, F., Nonis, R., Palestri, P. (2016, May). Design of a transmitter for high-speed serial interfaces in automotive micro-controller. In Information and Communication Technology, Electronics and Microelectronics (MIPRO), 2016 39th International Convention on (pp. 84-88). IEEE.
- Bandiziol, A., Grollitsch, W., Brandonisio, F., Nonis, R., Palestri, P. (2016, October). Design of a 8-taps, 10Gbps transmitter for automotive micro-controllers. In Circuits and Systems (APCCAS), 2016 IEEE Asia Pacific Conference on (pp. 321-324). IEEE.
- Bandiziol, A., Grollitsch, W., Brandonisio, F., Nonis, R., Palestri, P. (2017, June). System and transistor level analysis of an 8-taps FFE 10Gbps serial link transmitter with realistic channels and supply parasitics. In Ph. D. Research in Microelectronics and Electronics (PRIME), 2017 13th Conference on (pp. 297-300). IEEE.

and at national level in:

- Bandiziol, A., Grollitsch, W., Brandonisio, F., Nonis, R., Palestri, P., Selmi L. (2016, June). Design of a transmitter for high-speed serial interfaces in automotive micro-controller. Proceedings of GE2016, 48th Conference, Brescia, Italy.

Moreover, the following articles have been accepted and will be presented at international conferences with peer-review:

- Bandiziol, A., Grollitsch, W., Brandonisio, F., Bassi M., Nonis, R., Palestri, P. Design of a half-rate receiver for a 10Gbps automotive serial interface with 1-tap-unrolled 4-taps DFE and custom CDR algorithm. Circuits and Systems (ISCAS), 2018 International Symposium on. IEEE.
- Dazzi M., Palestri P., Rossi D., Bandiziol A., Loi I., Bellasi D., Benini L. Sub-mW multi-Gbps chip-to-chip communication Links for Ultra-Low Power IoT end-nodes (Invited Paper). Circuits and Systems (ISCAS), 2018 International Symposium on. IEEE.

Finally, research activities results have been presented at:

- Infineon Innovation Days, June 2015, Villach, Austria
- PhD Expo, 18th June 2015, Udine, Italy
- Infineon Innodays, November 2015, Villach, Austria
- CAS Day, 6th March 2016, Villa del Grumello, Italy
- PhD Expo, 26th May 2016, Udine, Italy
- PhD Expo, 25th May 2017, Udine, Italy

Acknowledgments

My first thanks goes to my tutor at University of Udine, Prof. Pierpaolo Palestri. It is hard for me to imagine ways to be a better guide than what he has been to me in these three years. His help has just been invaluable, his curiosity for all aspects of electronics has been a constant stimulus for me in trying to become a better researcher everyday.

An enormous thanks goes also to my supervisor at Infineon Technology Design Center, Dr Roberto Nonis. On top of the technical contents that he has transmitted to me in these years, the thing that I have tried to learn at most from his daily behavior is to being able to take time to listen to people even in the busiest days. Along with him, also Werner Grollitsch has taught me a lot more than high-speed serial interfaces design: humility is a much bigger virtue, especially when shown towards a guy that knows nearly nothing of a field in which you have several years of experience.

The fact that I have been able to fully concentrate on this work along this three year span has been possible only thanks to my family: nothing of this would have been possible without your continuous support. Examples are much more powerful instruments than words, and you all have given me plenty of them in my whole life. My biggest hope is to be every day even just one tenth as inspirational as you have been, and are, to me. A big thank goes also to Ginevra: it would have been easy for you to say that I was a little bit too focused on work in the last months, but instead you decided to fit to my time as water in a bowl. I will remember it when your bowl will be smaller.

Having one office with a great atmosphere is not for everybody. Having two of them has been just amazing. I have found all the support and good vibes that I needed in these years at CIS in Infineon Technologies, and when I talk to my friends or family about how I feel when I enter my office they almost don't believe me. It's hard to find such a mix of experience, love for the job, dedication and easy-going attitude. I want to particularly thank Fabio, Francesco, Andrea, Matteo, Dmytro and Thomas for the patience they have demonstrated to have with me every day. A big thanks goes also to Stefan and Andrea for their work done with the layout. Obviously, without them there would be nothing to measure in the lab. And so thanks also to Florin for the first round of measurements, and sorry for the next ones that will come. At the same time, I will never forget my time in Selmilab: Paolo, Stefano, Enrico, Giulia, Owes, Tommaso, Andrea, Alessandro, Julian, Martino, Patrick, Daniel, Federico, I have memories for each one of you. Thanks for the laughs, for the help, for the tips. If these four years are gone so fast and it seems to me that I entered that lab for the first time yesterday, it's mainly because of you.