


A 10 Gbps Wireline Transceiver Link

To Interface Future RF-DACs

Jun Feng
Master Thesis



August 2020

Department of Microelectronics

Faculty of Electrical Engineering, Mathematics and Computer Science



A 10 Gbps Wireline Transceiver Link

To Interface Future RF-DACs

by

Jun Feng

to obtain the degree of Master of Science

Responsible professor prof. dr. Leo de Vreede
Daily supervisor dr. Morteza Alavi

Research Group Electronic Circuits and Architectures (ELCA)
Department of Microelectronics

Defended on the 21st of August, 2020

Committee prof.dr. Leo de Vreede (chair, ELCA)
 dr. Morteza Alavi (ELCA)
 dr. Daniele Cavallo (Terahertz Sensing)



DELFT UNIVERSITY OF TECHNOLOGY

FACULTY OF ELECTRICAL ENGINEERING, MATHEMATICS
AND COMPUTER SCIENCE

MSC ELECTRICAL ENGINEERING

This thesis cannot be made public before the 1st of September, 2022

Abstract

This thesis presents the development of circuits and systems for fast wireline transceiver links that will enable a move towards highly integrable RF digital-to-analog converters.

A new perspective on the analysis of bit error rates in wireline links leads to the PAM spectral design space chart: a novel, visual system design analysis tool for PAM wireline circuit designers.

Moreover, a < 2 mW/Gbps/lane, 10 Gbps wireline transmitter has been designed and taped-out in 40 nm CMOS. The proposed inherently pipelined 16:1 multiplexer and current mode logic driver design procedure are the key enablers for this performance.

Finally, for development of a 10 Gbps wireline receiver, a novel self-synchronized receiver design is proposed that removes the need of a classical clock and data recovery loop. At its core, this receiver comprises the design of a high-speed two-tail comparator and an asynchronous metastability detection loop.

Acknowledgments

This challenging thesis project has been a real joy to carry out. First and foremost, I would like to thank dr. Morteza Alavi for fully trusting me with this new ELCA research topic. The supervision was excellent, from all perspectives.

I thank prof. dr. Leo de Vreede and dr. Daniele Cavallo for their willingness and interest to assess this work.

I would also like to thank PhD student Milad Mehrpoo for his committed guidance during the project. I most fondly look back at the thought provoking discussions we had while coming up with new solutions.

I am deeply indebted to PhD student Mohammadreza Beikmirza for his critical assistance during my sudden tape-out endeavor. His help allowed me to work in focus on my first piece of real silicon.

I am thankful for the financial support kindly arranged by Kave Kianush from Catena for this project and corresponding tape-out.

My gratitude also goes out to all microelectronics professors from whom I learned so much. I am happy to have discovered such exciting topics in the past two years.

I am most grateful that I could pursue this master while having the full support of my parents, brothers and close friends. From the bottom of my heart: thank you all!

Jun Feng
August 2020, Delft

Contents

Abstract	v
Acknowledgments	vii
1 Introduction	1
1.1 Problem and Purpose	1
1.2 Wireless or Wireline?	2
1.3 Thesis Overview	4
2 Design Space of PAM Wireline Links	5
2.1 High Data Rates in Wireline	5
2.1.1 Using The Bandwidth	5
2.1.2 Raw Digital Circuit Speed	7
2.2 PAM Spectral Design Space Chart	9
2.2.1 Constellation and Errors	9
2.2.2 Development and Demonstration: PAM Spectral Design Space Chart	11
2.3 Thesis Design Targets	15
2.4 Chapter Conclusion	16
3 10 Gbps Wireline Transmitter	19
3.1 Overview and Challenges	19
3.1.1 Core Circuit Styles	19
3.1.2 Functionalities	21
3.2 CMOS Multiplexer Design	22
3.2.1 MUX Tree Speed Limitations	22
3.2.2 10 Gbps Inherently Pipelined MUX Tree	25
3.3 Current-Mode Logic Driver Design	28
3.3.1 Current Mode Logic Chain Issues	30
3.3.2 Driver Design	33
3.3.3 CMOS-to-CML	36
3.4 Auxiliary Circuit Design	37
3.5 Physical Implementation and Performance Verification	38
3.5.1 Chip Layout	39
3.5.2 Core Simulations	41
3.6 Initial Measurement Results	43
3.7 Chapter Conclusion	46

4	Design of Comparators for Wireline	47
4.1	Considerations for Wireline Receivers	47
4.1.1	Clocked Operation	48
4.1.2	Hysteresis	49
4.1.3	Input Common-mode and Offset	53
4.2	State-of-the-Art Comparators	54
4.3	Two-tail Comparator Analysis and Design	56
4.3.1	Summary of Requirements	56
4.3.2	Analysis of Operation	56
4.3.3	Design Choices and Results	60
4.3.4	Fast Simulation Methods	61
4.4	Chapter Conclusion	61
5	Self-Synchronizing Wireline Receiver	63
5.1	Overview and Key Challenges	63
5.1.1	Considerations for Sampling Clock Alignment	63
5.2	Synchronization by Metastability	64
5.2.1	Motivation of Architecture	65
5.2.2	Concept Soundness	67
5.3	Self-Synchronizing Receiver Design	69
5.3.1	Metastability Detector	69
5.3.2	Clock Swapper	74
5.3.3	Simulation Results	74
5.4	Chapter Conclusions	76
6	Conclusion and Future Work	77
6.1	Conclusion	77
6.2	Future Work	77
	Bibliography	81
	Appendices	84
	Appendices	85
A	Wireline Transceiver Literature Review	87
A.1	Modern Wireline Transceivers	87
A.1.1	General Trends	87
A.1.2	High Data Rates at 65 and 40 nm	89
A.1.3	Low-Power Design	89
A.2	Original Data	90
A.3	Wireline Interface Standard Survey	94
B	Wireline Transmitter Auxiliary Circuit Designs	97
B.1	Aligned Divided Clock Generation	97
B.2	Timing Delay Trimmers	100

B.3	Pseudo-Random Bitstream Generator	102
B.4	Current Reference	102
C	Wireline Transmitter Layout Pin Plan	105
D	Wireline Transmitter Measurement Setup PCBs	107
E	BER Derivation for PAM-M Flash Comparators	113
F	Fast Comparator Simulation Methods	121
F.1	Noise Testbench	121
F.2	Offset Testbench	123

1 Introduction

It is remarkable how heavy the demands on state-of-the-art high-speed electronics are. Within a time span of 50 ps^1 , a ray of light can travel through our pupils to arrive at the retina, where optical receptors signal the available information to our brain. Meanwhile, standards on short-reach, single wireline connections in the next generation of ethernet applications foresee a single bit transmission that is given a maximum existence of only 10 ps until the next bit should arrive and be resolved.

1.1 Problem and Purpose

At the ELCA group, *all-digital RF-DACs* (radio frequency digital-to-analog converters) are an active research topic. To satisfy the demand for a continuously higher data throughput, these RF-DACs are foreseen to ramp up in the direction of welcoming a 12-bit resolution IQ input information stream up to 3 Gbps. While we are not at that stage yet, testing the current generation of RF-DAC ICs is cumbersome for the following reasons:

1. A conventional wired link to provide the baseband test data for the RF-DAC under test is limited in bit quantity (number of package pins) and data rate (full-swing CMOS, crosstalk).
2. As a solution, on-chip memory banks are currently employed to load in test signals:
 - As it stands, the speed of 40 nm CMOS^2 on-chip memory will not be sufficient to provide the RF-DACs their required data in future generations.
 - The pre-loaded test signal prohibits the use of real-time testing or long bit streams.
 - The on-chip memory bank is large in area and dominates the total tape-out cost.

This master's thesis will present our work for the use case of testing such high-speed ICs, to take away the above limitations and to replace the memory banks with on-chip *wireline receivers* and supply data from an external *wireline transmitter*.

Our purpose is to gain design experience of wireline transceiver design: from system level to circuit level and down to implementation at the transistor level. This one-year project is the first step towards this goal and on this specific topic at ELCA. The resulting insight and proposed designs on the transmit *and* receive side should accomplish a leap towards the realization of our envisioned serial wireline link for all-digital transmitters. The silicon design work is executed in 40 nm CMOS .

Before we give a precise overview of this thesis in Section 1.3, let us first briefly review the general context of wireline links to provide the readers a common point of reference.

¹ $50\times$ one millionth of one millionth of one second.

² 40 nm CMOS is currently one of ELCA's main transistor technology processes for RF-DAC design. Ultimately, our problem holds likewise for any other RF-DAC, regardless of technology.




Figure 1.1: Simplified traditional wireless communication system architecture.




Figure 1.2: Simplified functional decomposition of a RF-DAC. IQ_{BB} represents the baseband in-phase and quadrature data streams.

1.2 Wireless or Wireline?

While the topic of interest is *wireline transceivers*, we must start from trends seen in *wireless* communications to comprehend the setting of this thesis. Namely, modern *wireless* communication systems are experiencing an “all-digital revolution” for circuit design. A well-known development is the adoption of all-digital phase-locked loops (ADPLLs) to obtain stable high-frequency clocks for frequency translation or modulation. Traditionally, such (de)modulation operations would be performed on an analog signal as shown in Fig. 1.1.

This push for digitization is approaching the antenna: a multitude of successful endeavours are being reported in the realization of all-digital transmitters and receivers, leaving the only occasional amplifying stages at the front end in the pure analog domain. A concrete example is the emergence of the *direct-digital RF modulator* (DDRM or *RF-DAC*) as a standalone replacement of both digital-to-analog conversion and RF modulation [1].

Fig. 1.2 shows the RF-DAC from a functional point of view. At its input, the RF-DAC accepts a low-speed, digital data stream (usually constituting the bit words of in-phase and quadrature components). The RF-DAC then upconverts the “baseband” data stream by means of digital interpolation (upsampling and digital low-pass filtering). A simple AND-operator acts as a mixing stage preceding an array of transistors that generates the RF signal in the analog domain, carrying the baseband information.

The above approach comes with its own bottlenecks: the processor-to-modulation interface is now in the digital interface. With the continued increase of modulation complexity motivated by the need for more data throughput, the number of bits arriving at the modulator per clock cycle will increase to prohibitive amounts. A normal full-swing CMOS interface to such a modulator does not suffice anymore—we foresee the need for RF-DACs that modulate 2 streams of a width of 12-bit each at 3 Gbps ($\equiv 72$ Gbps).

Clearly, to arrive at economical footprints, standalone RF-DACs must employ serialized data interfaces as shown in Fig. 1.3. Even in measurement activities in the development of RF-DACs, a compatible (or co-designed) serialized interface must be in place to allow for real-time and extensive



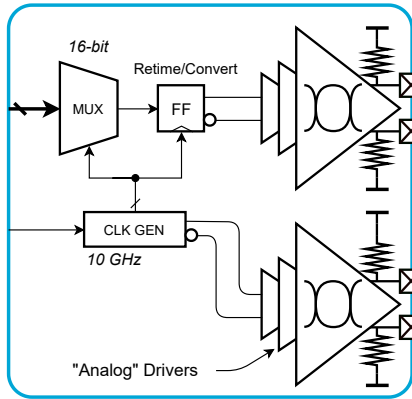


Figure 1.3: Proposed RF-DAC with a serialized input, that is, employing a SerDes wireline interface.

Chapter 2: Design Space of Wireline Links



Chapter 3: 10 Gbps Wireline Transmitter



Chapter 5: Self-Synchronizing Wireline Receiver

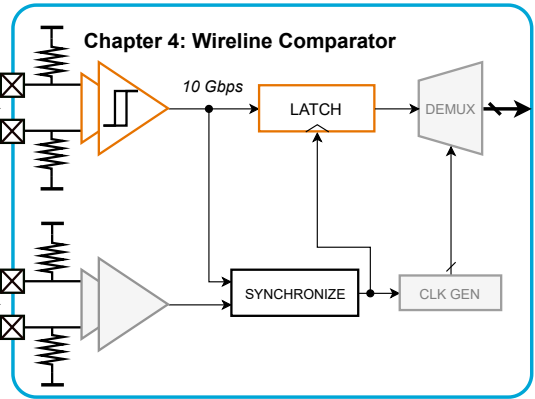


Figure 1.4: Organization of this thesis.

testing.

In other markets, such as display connectivity or optical networking, wired serialized interfaces have readily been developed for similar high-speed, limited-pin problems. These serialized interface solutions (Serialized to Deserialized, SerDes) have generally been developed to cater a wide range of applications possibly facing highly lossy channels and having different bit error tolerances.

Thus, the presented knowledge and proposed design efforts in this thesis should drive the development of a SerDes TRX IC with the specific application of enabling real-time RF-DAC testing. This envisioned future SerDes interface for RF-DACs (or other forms of digital-RF circuits) must not limit the overall communication system performances.

In the remainder of this thesis, we will refer to SerDes interfaces as “wireline links” or “wireline transceivers” to be able to emphasize that this work concerns high-speed *baseband* data transmission over a wired medium. Specifically, fast communication from chip to chip over a printed circuit board.

As part of an early literature study, we examined a large number of wireline papers and wireline commercial standards. We should preface the results of this literature study by saying that the observed developments are not representative for the starting point of this thesis project. In this project, we are cultivating design experience by building a wireline link from the ground up. The literature study is available for the reader in Appendix A.

1.3 Thesis Overview

This master thesis reports on our various efforts on system, circuit and silicon development during this first run at ELCA towards the realization of the envisioned serial wireline link for all-digital transmitters. We put emphasis on the resulting critical functions, considerations and design insights. To this end, the thesis is organized into three parts as illustrated in Fig. 1.4.

Chapter 2 explores a wireline link design from the system level perspective. This first part shows how the data rate of a wireline link is limited and develops a design tool that clarifies the trade-offs between spectral efficiency, bit error rate and circuit level limitations.

A combination of this understanding and the constraints of this project lead to a formulation of the design targets for parts two and three. Our aim is to design and implement a 10 Gbps PAM-2 wireline transceiver link with a bit error rate $< 10^{-8}$.

Chapter 3 presents the successful design and implementation of a complete, testable 10 Gbps wireline transmitter IC in 40 nm CMOS. We took the challenge of designing for fabrication after the rise of a sudden tape-out opportunity. Among others, the design of an inherently pipelined multiplexer tree, efficient current mode logic drive chain and other important auxiliary circuits are showcased.

Chapter 4 and Chapter 5 together establish the last, third part of this thesis that focuses on the receiving end in a wireline link. We first examine the challenging design of a 10 Gbps two-tail comparator for wireline applications. An analysis of BER degradation effects and development of fast simulation methods in part allow the proper design of such a high-speed comparator.

Chapter 5 covers the exciting development a self-synchronized wireline receiver. In view of the finite design time left over, we decided to fully focus on the the remaining critical challenge in a wireline receiver: sampling clock alignment. As shown in Fig. 1.4, the receiver utilizes the designed two-tail comparator and a new synchronization scheme based on metastability detection.

We draw a conclusion in Chapter 6 and end with a future outlook.

2 Design Space of PAM Wireline Links

Making architectural choices and setting up requirements for wireline links is a true challenge; it is evident from the long list of authors in leading wireline transceiver papers, that the overall complexity grows super-linear with the list of functions to consider. Unlike more traditional data-converter design or RF transceiver design that seek certain SNR and/or linearity, wireline links deal directly with the effects of bit errors. As such, wireline links call for a BER-centric design intent.

This chapter embodies the first part of this thesis work, in which we explore the design task at hand from a high-level perspective. We first discuss the limitations of data rate: the primary function of a wireline link is to pump as much data as possible with little resources (i.e. physical pins). Then, we will develop critical design insights regarding trade-offs between BER and data rates for wireline link *system* design. These insights were found to be scarcely or never described before, so we provide intuitive and concise explanations. This chapter ends by motivating the design targets for the remaining two parts of this thesis.

2.1 High Data Rates in Wireline

The data rate R_b of a wireline link is the most prominent specification for the end user. At first sight, the achievable data rate depends on the available channel bandwidth and technology performance. Let us clarify these aspects from the perspective of a wireline link circuit designer.

2.1.1 Using The Bandwidth

Channel Capacity for Wireline Shannon's theoretical limit (the channel capacity C in bits per second) for *error-less* transmission of information states

$$C = BW \log_2(1 + SNR) \quad (2.1)$$

where BW is the channel bandwidth and SNR the signal to noise power ratio (SNR). The point of Shannon's limit is that for a given available bandwidth and total SNR at the detector, the spectral efficiency $\eta_{BW} = R_b/BW$ will never exceed $\eta_{BW,MAX} = C/BW$. The limit on spectral efficiency can only be improved by spending power to gain SNR .

In reality, the circuit designer does not design for Shannon's case. Reaching C/BW at a given SNR required the implementation of a so-called "golden code", which is not the task of the circuit designer. Also, wireline interfaces are not *error-less* as in Shannon's theory. The error probability, i.e. the *bit error rate* BER, is always tolerable under a specified maximum level¹.

Nevertheless, the point of Shannon's relation still stands firm for a wireline circuit designer. Namely, for a given bandwidth, the data rate is limited by a maximum spectral efficiency. Lifting

¹For example, the BER should be low enough to ensure that the mean time until failure does not exceed the expected lifetime of a commercial system. In our application, it is the maximum (non-zero) specification on the error vector magnitude for wireless communication that should be maintained.

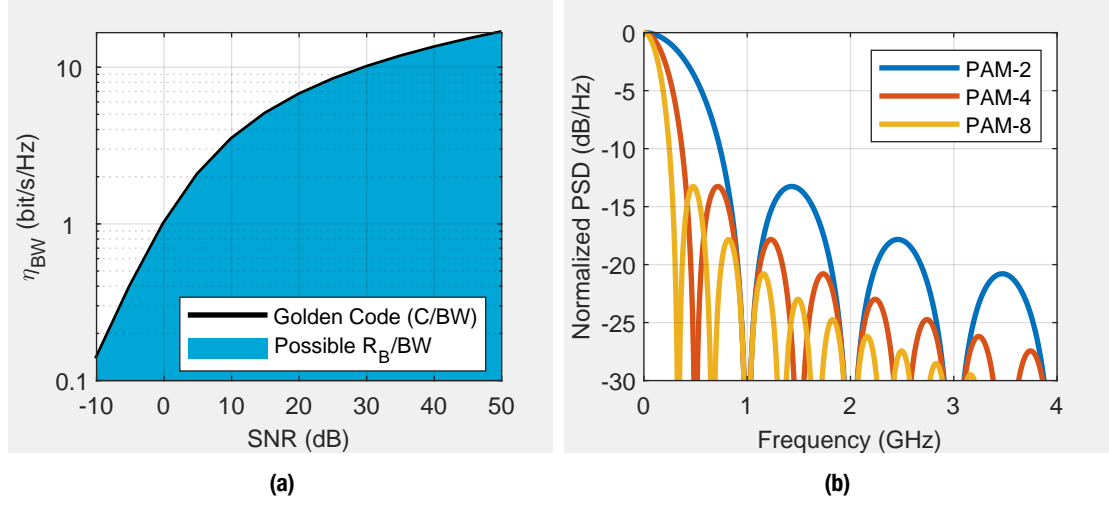


Figure 2.1: Spectral efficiency η_{BW} and PAM power spectral densities (PSD). (a) Any (SNR, η_{BW}) combination above Shannon-Hartley limit C/BW is not realistic. (b) Here R_B (bit/s) is fixed at 1 Gbps: the *symbol* rate (symbol/s) reduces for higher PAM order, thus BW_{NULL} reduces.

this limit requires an expenditure of power for higher SNR.

Modulation Complexity For the wireline circuit designer, increasing the SNR lifts the spectral efficiency limit, but does not immediately imply a higher data rate! For example, if the designer lowers the noise floor in the circuits, there will be less bit errors² but no change of the effective spectral efficiency. In other words, the BER dropped but the data rate remains unchanged.

If the designer wants to make use of the lifted spectral efficiency limit at an increase of SNR, the signal modulation complexity of the circuit design must increase. Wireline links are inherently baseband systems, so more information must be packed in the DC signal levels of the transmitted symbols.

We commonly refer to this as pulse amplitude modulation (PAM). PAM-2 refers to two-level signaling, PAM-4 to four-level signaling etc. For the same (single-sided) occupied null bandwidth BW_{NULL} ³, the data rate of PAM-2^M grows as

$$R_B = M \cdot BW_{NULL} \quad (2.2)$$

where M is the number of bits per symbol.

Thus, an increase of spectral efficiency for data rate, when close to the current $\eta_{BW,MAX}$, requires an expenditure of power *and* an increase the PAM complexity by means of a significant change of the circuits.

Data Rate Matching for BER Having this clear, the following question can easily arise from the circuit designer: what is the effect if the data rate R_B (i.e. B_{NULL}) does not perfectly align with channel bandwidth? The reality is that the signal spectrum is not guaranteed to match the channel bandwidth well in practice, such that η_{BW} is not guaranteed to be optimum after choosing a modulation complexity.

A dangerous case is when R_B is set so high that the channel significantly attenuates the high-frequency parts of the signal spectrum. From a wireline transmission perspective, it means that

²This basic understanding is sufficient for now. We arrive at the precise BER versus SNR relationships in Section 2.2.

³The spectrum of PAM signals is sinc-like: a series of bulges centered around 0 Hz. The null bandwidth is defined as the first frequency where the signal power spectrum drops to zero.

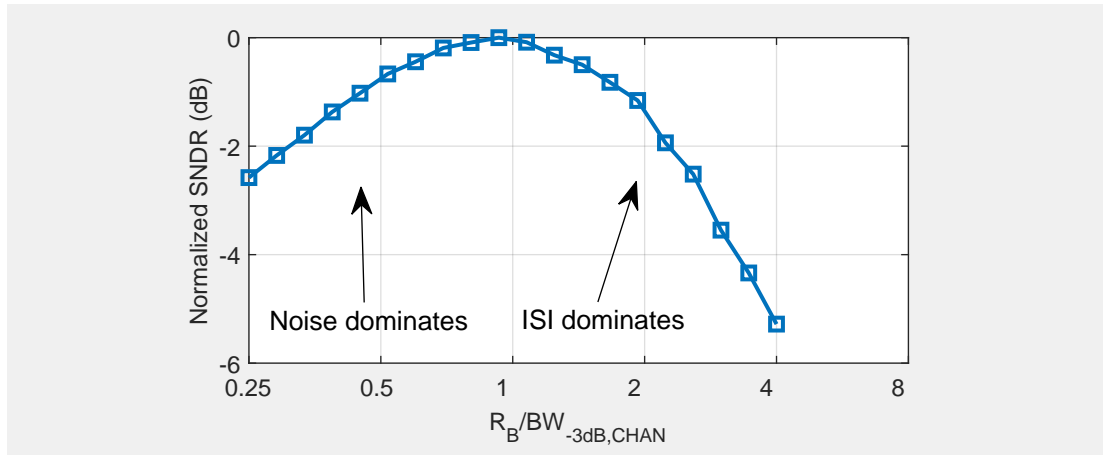


Figure 2.2: Simulation of a noise and ISI trade-off, here for PAM-2 with a steep channel filter roll-off. Therefore, the ISI starts to affect the signal to noise and distortion ratio already when $\eta_{BW} = 1$.

each symbol is smeared out over time and affects a future transmission. This phenomenon is known as *inter-symbol interference* (ISI) and increases the achievable BER. From the perspective of the spectrum it is clear that some of the transmitted signal power is lost, incurring a loss of SNR.

On the flip side, just having a conservative R_B to prevent signal filtering does not help this case. As BW_{NULL} drops, the relative portion of unfiltered noise by the channel increases, again incurring more bit errors at the moment of detection.

These two trends suggest that there is an optimum arrangement of the signal spectrum versus channel bandwidth (channel profile). [2] suggests that for PAM-2 signaling it is wise for the BER to have roughly $R_{B,PAM2} = 1.5 \cdot BW_{CHAN}$ to balance noise and ISI. That is, $\eta_{BW} = 1.5$ to maximize SNR by just adjusting the relative position of signal and channel profiles.

MATLAB simulations confirm (like [2]) that the actual optimum η_{BW} depends on the channel profile (e.g. order of roll-off) and noise level. Fig. 2.2 shows one of the results: in all cases, it was observed that the SNR degradation is moderate (few dBs). Moreover, the ISI-limited region is more sensitive to choice of η_{BW} .

Overall, the circuit designer should be aware that a change of modulation complexity can dangerously move the shape of the signal spectrum with respect to an existing channel, incurring loss of SNR. This may prohibit the intended promotion of data rate. Thus, it is good to note that if the channel bandwidth large and fixed, the receiving side can employ additional filtering (or even attempt to implement a matched filter).

2.1.2 Raw Digital Circuit Speed

Ultimately, a wireline link is an analog/mixed-signal design effort. For analog circuits, the gain-bandwidth product is commonly used to assess the allowable circuit speed⁴.

However, a good wireline circuit designer must be aware of the limits of the digital blocks too. Even if there is a large available channel bandwidth, the inherent technology performance may not allow the designer to just turn up the clock frequency.

Digital Clock Frequency The most intuitive way to assess the speed limitations of digital circuits is by investigating the highest practical clock frequency.

⁴The gain-bandwidth product is a simplified metric that estimates the maximum frequency at which a transistor (arrangement) can still generate gain, i.e. operate without loss

From the perspective of digital circuit speed, it is important to know the typical rise and fall times⁵ of the technology. $t_{R/F}$ symbolizes the average of rise and fall times. The digital clock must fully swing between 0 V and the supply voltage in order to guarantee system functionality. Because one clock period consists of two clock edges, it is impossible to work with a clock running with a clock period faster than $2t_{R/F}$. At this point, the clock signal is not a proper digital signal anymore, but more resembles a sine wave between $0.1V_{DD}$ and $0.9V_{DD}$. A more realistic requirement is $T_{CLK} > 4t_{R/F}$.

$t_{R/F}$ on its own, depends on the technology at hand and the willingness to burn power. For digital circuits, the *fan-out* measures the capacitive load seen by a logic gate (e.g. buffer) divided by its own drive strength. $t_{R/F}$ is at minimum when there is no load: only the parasitics of the device are loading the gate by itself. $t_{R/F, UNLOADED}$ is a characteristic of the technology.

In reality, $t_{R/F, UNLOADED}$ is impractical. The clock must be distributed around the (wireline transceiver) circuitry, requiring a clock buffer network. $t_{R/F, UNLOADED}$ is only attained when the driving stage is sized much larger (power hungry) than the driven stage. So, to practically distribute the clock, at least a fan-out of one (FO1) with rise/fall time $t_{R/F, FO1}$ should be maintained.

Unfortunately, the clock buffer network must always size up at some point to drive some sub-circuit along then chain, i.e. the maximum fan-out in the chain is greater than one. As Fig. 2.3a illustrates, $T_{CLK} > 4t_{R/F, FO1}$ may not be sufficient. In fact, to minimize overall propagation delay in a driver chain, a fan-out of four (FO4) [3] is a common design guideline.

Let us look at the digital clock frequency in 40 nm as an design example. At 1.1 V supply voltage, the nominal rise and fall times measure about 15 ps with a fan-out of one (FO1). To be able to create just a functioning FO1 buffer chain, the clock frequency cannot exceed $1/(4 \cdot 15 \text{ ps}) \simeq 16.6$ GHz. If the provided clock (reference) is too weak to drive the system or we would like to lower power consumption, the clock must have a lower frequency to allow for a higher fan-out distribution network! For example, a compact, power-efficient FO4 design cannot distribute clocks much faster than ~ 4.4 GHz.

Clock Edge Creation A possible remedy to circumvent this limit is the use of multiple precisely spaced clock phases. In effect, more clock edges are available in the system in the timespan of one clock period. For example [4] shows how to control the MUXes and detectors with *multiphase clocking*. An example of a quarter-rate TX multiplexer is shown in Fig. 2.3b.

In this case, the circuit designer should measure the minimum *propagation* delay of the technology. The origin of the digital clock rate limitation is then partially shifted to the accuracy limitations of the phase interpolator and the willingness to increase system complexity. Moreover, if there are many finely spaced clock edges, the design must ensure that the clocked digital circuit blocks are able to maintain their setup and hold timing budgets⁶.

To summarize, it is important for a wireline circuit designer to be aware of the limitations of data rate in terms of achievable spectral efficiency, PAM complexity and existing channel bandwidth to be able to get a feel of the available design space. In addition, a good understanding of the maximum digital clock frequency in the given transistor technology will indicate the challenges in architecture or circuit synthesis later on.

⁵The rise (fall) time is the time period it takes for a signal to increase (decrease) in voltage between the thresholds $0.1V_{DD}$ and $0.9V_{DD}$.

⁶The setup time is the minimum time that a clocked memory element requires the data to be present at its input before the rising edge of a clock. The hold time refers to the minimum time that the data should stay steady *after* the same clock edge to ensure successful sampling. Together, they limit the allowable speed of the clock and logic gates.

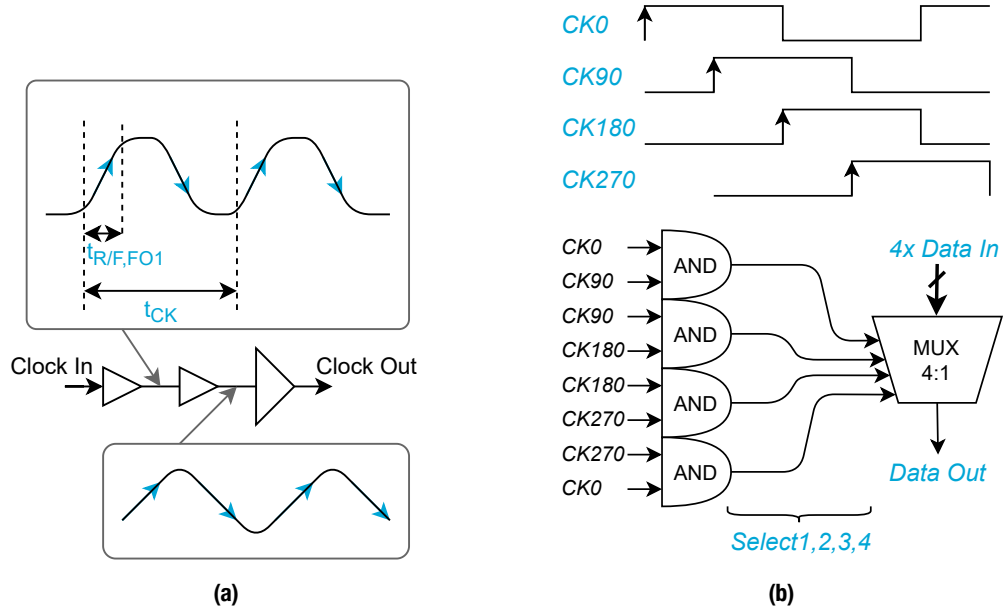


Figure 2.3: Digital clock considerations. (a) An increase of rise time due to a large load limits the maximum clock frequency. (b) Four-way multiphase MUX arrangement (“quarter-rate”).

2.2 PAM Spectral Design Space Chart

In the above discussion, we provided the necessary perspectives for a wireline circuit designer to start thinking about the effects of a chosen data rate. However, especially in low voltage supply designs, circuit-level constraints can additionally limit the data rate to a lower number than expected from the technology’s potential speed.

To the best knowledge of the author, there are no concise analyses available that show the interrelation of such limitations. Commercial wireline standard specifications are usually formed behind closed doors. Nevertheless, in the view of thesis project’s intent, there is a great need for the realization of a method that allows instinctive assessment of system design target feasibility to aid the circuit designer. It would be of great help to know the best possible design result.

Consequently, we develop a new wireline link system design theory based on a combination of the topics discussed in [5] and [6]. As a result, this thesis work also contributes to the future design of high-speed wireline links at ELCA, from a high-level perspective. We call this theory the PAM *spectral design space*.

2.2.1 Constellation and Errors

Recall that the BER can often easily be related to the *quality* of the eye diagram (e.g. in the case where jitter is not a dominant cause for errors). In this work, we quantify this quality with the metric Q_{EYE} , the “eye quality”.

We define Q_{EYE} as the ratio between the distance from a signal constellation point to the neighboring decision threshold, and the rms-noise at the detector:

$$Q_{EYE} = \frac{V_{PK,EYE}}{V_{N,rms}}. \quad (2.3)$$

For a given BER target and constellation diagram, the required Q_{EYE} can be derived. Fig. 2.4 reviews the relation of time domain waveform and eye diagram in this context for a PAM-4 signal. Fig. 2.5

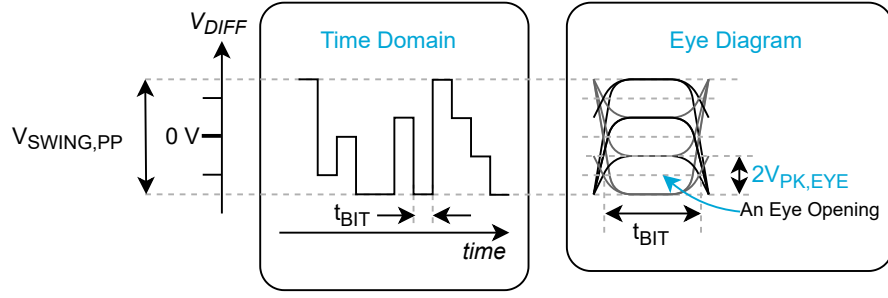


Figure 2.4: Review of the eye diagram for a PAM-4 constellation.




Figure 2.5: Eye quality definition. Here depicted for the lowest symbol of a PAM-4 signal. In our derivations, we take into account that the noise on inner symbols can cause a bit error in either direction.

visualizes the definition of $Q_{\text{EYE}} = V_{\text{PK,EYE}} / V_{\text{N,rms}}$. Looking at the eye diagram in wireline signaling, $V_{\text{PK,EYE}}$ should be interpreted as the peak voltage swing around the detection thresholds in the eye diagram. In other words, $0.5 \times$ the height of each eye opening.

In the case of a **bandwidth-limited** system—which is reasonable for wirelinks as a specific channel profile is often specified, increasing modulation complexity can result in a higher spectral efficiency η_{BW} . As discussed earlier, this can increase the system transfer rate if the BER remains sufficiently low.

For example, in (wireless) communications constellation sizes with 16 possible symbols are not unheard of. We can refer to such forms of modulation as “non-orthogonal” signaling, as opposed to orthogonal signaling, e.g. QAM vs. FSK modulation. Logically, for wireline links, it is possible to increase from PAM-2 to PAM-4, 8 etc.

However, an increase of the modulation complexity has generally other critical effects for non-orthogonal signaling, in addition to the improvement of η_{BW} :

- The average signal power drops dramatically at a fixed energy expense as Q_{EYE} decreases *exponentially* with the increasing number of bits.
- As the SNR decreases strongly, the signal swing must be increased heavily to compensate for this effect.

At the circuit level this swing is not without bound: an obvious reason is the supply voltage. Or, a constraint on slew rate might exist, imposing a limit on the product of swing and achieved bit rate.

Let us illustrate this trade-off in the following case study, where we consider the limited bandwidth vs. power trade-off for the virtually omnipresent used pulse amplitude modulation (PAM-) signaling.

2.2.2 Development and Demonstration: PAM Spectral Design Space Chart

Introduction Let us consider a wireline link with PAM-M signaling, i.e. M different symbols mapped to evenly spaced voltage levels where each symbol contains $k = \log_2(M)$ bits of information. We constrain ourselves to an analysis where M is an even number greater than or equal to 2.

The question is: *what is the maximum bit rate that we can achieve if M is our degree of freedom?* We perform a demonstration for the following constraints:

1. The target BER equals 1 out of one trillion (10^{-12}).
2. Differential swing limit of 1 V.
3. Limiting channel bandwidth of 5 GHz.
4. Differential rms-noise voltage of the system referred to the detector's input is $10 \text{ mV}_{\text{rms}}$; and
5. The receiver bandwidth is not perfectly matched to the achieved symbol rate, resulting in an overhead of one-third (33%). In this context, matched means the case of a matched filter in a receiver that would optimize the SNR.

Derivation of Theory From the last constraint (no matched conditions), we must recognize that for this case

$$\eta_{\text{BW}} = \frac{R_{\text{B}}}{BW_{\text{CHAN}}} = \frac{kR_{\text{SYM}}}{(R_{\text{SYM}}/2)(4/3)} = 1.5k = 1.5 \log_2 M \quad (2.4)$$

as a matched filter would have $BW_{\text{CHAN}} = R_{\text{SYM}}/2$ [6] and the overhead of 33% justifies the additional 4/3 factor. So, for a general description of the spectral efficiency, we can use

$$\begin{aligned} \eta_{\text{BW}} &= \frac{R_{\text{B}}}{BW_{\text{CHAN}}} = \alpha_{\text{MISMATCH}} \log_2(M), \\ \alpha_{\text{MISMATCH}} &= \frac{0.5R_{\text{SYM}}}{BW_{\text{CHAN}}} \end{aligned} \quad (2.5)$$

The spectral efficiency η_{BW} only increases slowly with M .

Let us start with derivations of waterfall (BER vs. SNR) plots. Traditional waterfall plots (vs. (E_b/N_0) as a function of bit energy E_b and noise density N_0) are usually leveraged by communication system designers. However, in the context of this thesis it is much more practical to derive as a function of SNR to hand the circuit designer a more intuitive knob to turn. The SNR can be derived from Q_{EYE} as shown later.

As we are dealing with two or more symbols, we can compute the symbol error rate (SER) by means of the function $f_{\text{SER},M}(Q_{\text{EYE}})$ as

$$\text{SER} = f_{\text{SER},M}(Q_{\text{EYE}}) = \left(1 + \frac{M-2}{M}\right) \int_{Q_{\text{EYE}}}^{\infty} \text{PDF}_{\text{gauss}}(x) dx \quad (2.6)$$

assuming that the noise's gaussian probability density function $\text{PDF}_{\text{gauss}}$ is equal regardless of symbol. The integral of $\text{PDF}_{\text{gauss}}$ from x_0 to ∞ can be computed via the *complementary error function*⁷ $\text{erfc}(x_0)$ [7].

⁷The evaluation [7] goes as $\int_{x_0}^{\infty} \text{PDF}_{\text{gauss}}(x) dx = 0.5\text{erfc}(x_0/\sqrt{2})$ where in exact terms, $\text{erfc}(z) = (2/\sqrt{\pi}) \int_z^{\infty} e^{-x^2} dx$. Regular calculation programs such as MATLAB have the complementary error function built in.




Figure 2.6: Effect of PAM complexity as a function of SNR. (a) Bit error rate, (b) minimum peak to peak swing required with $10 \text{ mV}_{\text{rms}}$ noise at the detector.

The factor $1 + (M - 2)/M$ before the integral can be intuitively understood from the fact that most constellation points have two neighboring symbols for a large M in PAM- M (except for the outer symbols). This factor matches for the specific PAM-4 BER example as discussed in [6], but we generalize for higher modulation complexities. Thus, compared to a PAM-2 case, the symbol error rate almost doubles at the same rms-noise voltage and eye quality as M grows.

For each symbol error, we can assume that only one bit is faulty while the others in the symbol are okay [7], so:

$$\text{BER} = \frac{1}{k} f_{\text{SER},M}(Q_{\text{EYE}}) = \text{SER}/k \quad (2.7)$$

For PAM- M , we now generalize the relation between SNR and Q_{EYE} as

$$Q_{\text{EYE}} = \sqrt{\text{SNR}/L_M} \quad (2.8)$$

with L_M a factor that we define to relate the signal power to the spacing between constellation points ($\equiv 2V_{\text{PK,EYE}}$). Omitting the exact derivation⁸ for the sake of brevity, we have

$$L_M = \frac{V_{\text{SIGNAL,rms}}^2}{V_{\text{PK,EYE}}^2} = \sum_{x=0}^{M/2-1} \frac{2(1+2x)^2}{M} \quad (2.9)$$

where it should be remembered that $V_{\text{SIGNAL,rms}}$ is the rms-voltage of the overall signal, while $V_{\text{PK,EYE}}$ is the peak voltage of each eye opening in the eye diagram.

So, we find a generalized description for PAM- M BER vs. SNR as

$$\text{BER}(M) = \frac{1}{k} \left(1 + \frac{M-2}{M} \right) \int_{Q_{\text{EYE}}}^{\infty} \text{PDF}_{\text{gauss}}(x) dx, \quad (2.10)$$

⁸ L_M is obtained by computing the average value ($1/M$) of the squared value of all possible symbol voltage levels normalized to $V_{\text{PK,EYE}}$: $(1+2x)^2$. Since we square the values and the eye diagram is symmetric around 0 V, the sum can count up to $M/2 - 1$ while a factor 2 compensates.

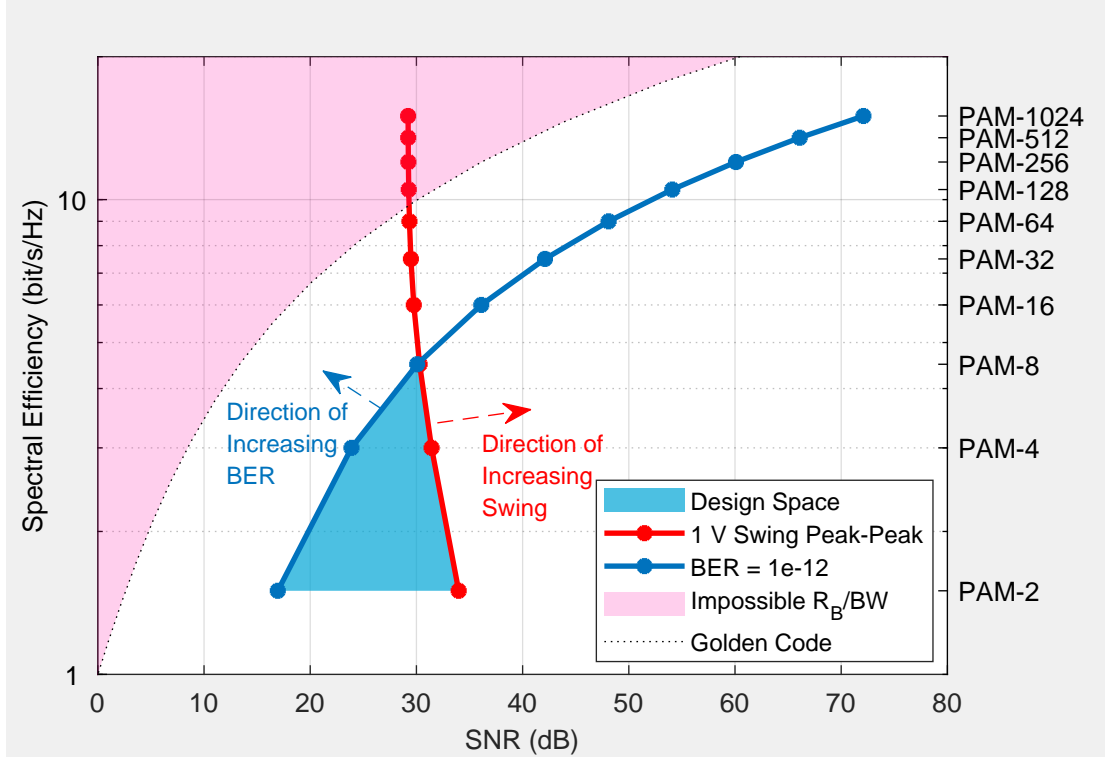


Figure 2.7: PAM spectral design space chart. The right y-axis shows the PAM modulation complexity that corresponds to the spectral efficiencies on the left for this specific demonstration.

$$Q_{\text{EYE}} = \sqrt{\frac{\text{SNR}}{\sum_{x=0}^{M/2-1} \frac{2(1+2x)^2}{M}}} \quad (2.11)$$

Fig. 2.6a shows the resulting waterfall plots up to $M = 64$. As expected, at a fixed BER the energy expense must increase in order to allow a denser signal constellation. This is the first influence of M that must be acknowledged.

Our derivation approach based on Q_{EYE} lends itself particularly to obtaining the required swing at the detector. That is, we can write

$$\begin{aligned} V_{\text{SWING}}^{\text{PEAK-PEAK}} &= (2M - 2)V_{\text{PK,EYE}} \\ &= (2M - 2)Q_{\text{EYE}} V_{\text{N,rms}} \end{aligned} \quad (2.12)$$

where the factor $(2M - 2)$ can be explained from the fact that a PAM-2 eye diagram has one eye opening, a PAM-4 eye diagram has *three* eye openings (i.e. $2M - 2 = 6$), PAM-8 has *seven* openings etc.

Substitution using Eq. 2.8 results in Fig. 2.6a, showing the swing requirements at the specified 10 mV_{rms} noise in this example. We can now see the second influence of M : at the same energy expense and rms-noise, the swing requirement increases drastically with M .

PAM Spectral Design Space Remember that the aim of this demonstration is maximization of the data rate and finding design requirements. With a given bandwidth, the key metric is allowable spectral efficiency.

Now, we come to the final product of the preceding derivations, by translating the information from Fig. 2.6a and Fig. 2.6b to the spectral efficiency vs. SNR plane. We form the *PAM spectral design space chart* as shown in Fig. 2.7.

The design-centric features of this chart are the modulation complexities drawn out on the right-hand y-axis and SNR on the x-axis. Table 2.1 summarizes the necessary equations in the spectral design space to perform the translation. Our generalized derivations with Q_{EYE} as a basis allow for easy incorporation of other types of constraints.

The dashed line black represents the maximum spectral efficiency that can be achieved under any SNR choice specification, using a theoretical golden code as discussed as the start of this chapter. So, pink area above this line is impossible to reach.

The blue line curving off to the right for increasing PAM complexity is obtained from the previous waterfall curves for the specified BER requirement while increasing M . That is, following a horizontal line in Fig. 2.6a. It is instructive to realize that this blue contour will move to the left when more BER is allowed.

Table 2.1: Overview of equations for generating a PAM spectral design space chart.

Equations	Description
<i>Main Equations</i>	
$\eta_{\text{BW}}(M) = \alpha_{\text{MISMATCH}} \log_2(M),$ $\alpha_{\text{MISMATCH}} = \frac{0.5R_{\text{SYM}}}{BW_{\text{CHAN}}} \quad (2.13)$	Spectral efficiency η_{BW} as function of M . A matched filter corresponds to the case where BW_{CHAN} matches half the symbol rate R_{SYM} , such that $\alpha_{\text{MISMATCH}} = 1$.
$Q_{\text{EYE}}(\text{BER}, M) = \frac{V_{\text{PK,EYE}}}{V_{\text{N,rms}}}$ $= \text{erfc}^{-1} \left(\frac{\text{BER} \log_2(M)}{1 + \frac{M-2}{M}} \right) \quad (2.14)$	Eye quality Q_{EYE} as function of BER (erfc^{-1} is the <i>inverse</i> complementary error function, see also discussion at Eq. 2.6)
$Q_{\text{EYE}}(V_{\text{SWING}}^{\text{PP}}, V_{\text{N,rms}}, M) = \frac{V_{\text{PK,EYE}}}{V_{\text{N,rms}}}$ $= \frac{V_{\text{SWING}}^{\text{PP}}}{(2M-2)V_{\text{N,rms}}} \quad (2.15)$	Eye quality Q_{EYE} as function of signal swing and expected noise level
$\text{SNR}(Q_{\text{EYE}}, M) = Q_{\text{EYE}}^2 L_M,$ $L_M = \frac{V_{\text{SIGNAL,rms}}^2}{V_{\text{PK,EYE}}^2} = \sum_{x=0}^{M/2-1} \frac{2(1+2x)^2}{M} \quad (2.16)$	SNR as function of eye quality Q_{EYE}
<i>Optional Equations</i>	
$C/BW_{\text{CHAN}} = \log_2(1 + \text{SNR}) \quad (2.17)$	Shannon-Hartley's golden code spectral efficiency C/BW_{CHAN}
$\frac{E_b}{N_0} = \frac{\text{SNR}}{\eta_{\text{BW}}} \quad (2.18)$	Conversion of SNR to E_b/N_0

If we also follow a horizontal line in Fig. 2.6b to represent the 1 V swing limit, we can plot the red line in the PAM spectral design chart that curves to the left for increasing PAM. At this specified swing limit and noise level, the SNR will not be higher than 35 dB for any PAM modulation complexity. A higher SNR could only be obtained by increasing the swing limit (or lowering expected noise level), thereby moving the red contour to the right.

In this demonstration, we cannot tolerate options on the left side of the constant BER curve or on the right-side of the constant swing curve. In other words, the chart shows that the design space is limited to the triangle-shaped area marked in cyan. For example, if employing PAM-2, the requirements can be met between roughly 17 dB and 34 dB of overall SNR over the link. The allowable SNR range narrows for PAM-4 and PAM-8.

So, under the conditions of this demonstration, we can conclude that PAM-8 would yield the ultimate spectral efficiency of 4.5 bit/s/GHz. With a 5 GHz channel bandwidth, this comes down to a data rate limit of

$$(4.5 \text{ bit/s/GHz}) \cdot (5 \text{ GHz}) = 22.5 \text{ Gbps.} \quad (2.19)$$

In this example, a choice for PAM-8 might not be feasible though, as the target SNR range has very little margin (the “top” of the cyan triangle). In addition, circuits operating at 22.5 Gbps may very well be impossible for the available transistor technology unless a multiphase clock architecture can be employed.

Spectral Design Space Chart Conclusion The derivations reveal that a combination of system constraints can heavily limit the design space in terms of signaling schemes. Our use of Q_{EVE} allows for the addition of low-level design constraints in an η_{BW} vs. SNR chart. Examples are BER, swing limits and slew rate limits: even effects such as jitter or ISI can be added to the analysis presented above to obtain more realistic curves in the spectral design space chart. Combined with a minimum channel bandwidth, this limited design space translates to a range of data rates.

The spectral design space chart is an effective, visual design tool. We can conclude that the wirelink transfer rate is not always limited by the device speed or a simple timing budget. A system designer must also consider limitations imposed on the signaling level of the design.

2.3 Thesis Design Targets

The insightful discussion on maximizing data rate and new development of an intuitive PAM design tool lead us to the creation of targets as shown in Table 2.2 for the remaining circuit design efforts in this work. We explain the choices in more detail below.

Goal and Constraints To review, the primary goal of this thesis is to design a wireline link with the potential to enable real-time RF-DAC testing. The following additional constraints are imposed:

Table 2.2: Summary of thesis wireline link design targets.

Link Characteristic	Target
Technology	40 nm CMOS @ 1.1 V supply
Modulation	PAM-2
BER (SNR)	$<10^{-8}$ average bit error/s (>15 dB SNR)
Speed	10 Gbps
Serialization factor	16×
Clock architecture	Forwarded or global clock
Power consumption	<25 pJ/bit (i.e. <250 mW @ 10 Gbps)

- The designs must be implemented in 40 nm CMOS, preferably with a nominal supply voltage of 1.1 V.
- The transmission channel must be differential to minimize emissions and allow future multi-lane operation.
- The transmission channel must be implementable on a regular FR-4 PCB with a channel length < 5 cm to emulate a realistic, worst case RF-DAC test scenario.
- Accounting for a tape-out lead time of two months, roughly 30 weeks are available for system design, circuit design and layout activities. Given that this project is a first development, there must be an emphasis on creation of insight and effective design.

Transmitter and Receiver Primary Functions The first iteration(s) of the ELCA wireline link will not employ clock recovery. Clock recovery is a complete topic on its own, and is not necessary for the envisioned use case. Therefore we target a receiver solution based on a forwarded clock or global clock architecture.

Testing RF-DACs requires the generation of in-phase and quadrature information. To comply with this need, we target for a (de)serialization factor of 16:1:16. This high serialization factor can support two 8-bit baseband streams.

Link BER This work targets a maximum bit error rate (BER) of 10^{-8} . It can be shown⁹ that this BER target corresponds with a reference application uncoded transmission of a QAM-1024 constellation at an error vector magnitude of -40 dB. For circuit design, the spectral design chart reveals that the overall corresponding link SNR must be greater than 15 dB ($Q_{\text{EVE}} = 5.6$).

Coincidentally, this BER requirement falls nicely between specifications of old, conservative commercial wireline links with $\text{BER} < 10^{-12}$ and modern, forward-error corrected links with $\text{BER} < 10^{-4}$.

Modulation and Speed Our PAM design space analysis reveals that PAM-8 is an efficient signaling scheme if a 1 V peak-to-peak swing limit is imposed on 5 GHz channel, that is, reaching 40 Gbps effectively.

In this work, we will target a PAM-2 link at 10 Gbps. This is a realistic target and keeps the design effort effective and challenging under the given constraints. Directly going for PAM-8 would prohibit the understanding of all important circuit blocks in the limited timeline available. A 1 V swing limit is generous with a 1.1 V supply and might be earlier limited by the speed of the process. Ultimately, it is important that the designs are flexible towards expansion to PAM-4 or beyond for future iterations of this work.

Power Consumption Power efficiency is not a priority for the use case of this work. However, to indicate a good overall design practice we specify a modest power efficiency target of < 25 pJ/bit¹⁰ for the first complete transceiver.

2.4 Chapter Conclusion

We discussed the theoretical limits of achievable data rate in wireline links. Our analysis of BER and spectral efficiency relations yields the development of a new design tool to optimize data rate

⁹This derivation goes as follows. The SNR is inversely proportional to error vector magnitude (EVM), that is, $\text{SNR} = 1/\text{EVM}$. The E_b/N_0 is equal to $\text{SNR}/\eta_{\text{BW}}$. Assuming a matched filter, $\eta_{\text{BW,QAM-1024}} = 10$. So, with $\text{EVM} = -40$ dB, we have $E_b/N_0 = 30$. MATLAB simulations show that this corresponds to $\text{BER} = 10^{-8}$ when employing QAM-1024.

¹⁰25 pJ/bit is equivalent to 25 mW/Gbps.

for PAM wireline link design, while considering practical constraints upfront.

As a result, this thesis work contributes to the future design of high-speed wireline links at ELCA, from a high-level perspective. We call this novel system analysis tool the PAM *spectral design space chart*.

The developed understanding in this work so far is the setup for the primary design target of this thesis: a differential PAM-2 wireline link at 10 Gbps and $\text{BER} < 10^{-8}$. It is worthwhile to remember the following observations and ideas from this chapter:

Data Rate

- Ultimately, the effective channel bandwidth limits the achievable data rate. In this case, spectral efficiency can be improved by a combination of extra SNR expenditure and multi-level signaling, such as PAM-4.
- A challenging channel response can be countered by the application of equalization techniques at the transmit or receive ends.
- If the channel bandwidth high, the circuit “bandwidth” or maximum speed of operation will limit the maximum data rate first. The minimum rail-to-rail rise times dictate the maximum clock frequency.
- Multi-phase clocking effectively increases the number of clock edges per clock period. So, the datarate can be higher than the maximum clock frequency in the system. In this approach, the minimum gate delay limits the speed of operation.

Spectral Efficiency

- By increasing the modulation complexity in PAM for a given SNR, the peak-to-peak swing must inevitably go up. Increasing the modulation complexity also requires a higher SNR for a given BER.
- Visually, the PAM spectral design space chart shows that in achieving a specific BER, the modulation complexity cannot be increased without boundary: the increase of required SNR can lead to a prohibitively large swing.
- The spectral design space chart is an excellent tool for testing the impact of system level design choices. It also shows a clear design space to know what the best realistic design target for data rate is.

3 10 Gbps Wireline Transmitter

Realizing a 10 Gbps (gigabit per second) wireline transmitter in 40 nm CMOS is not trivial. To reach such data rates or beyond, the high-speed digital circuits need careful design and non-conventional architectures. On the analog side, the required design procedures also deviate from the methods taught for traditional amplifier design. After all, we are dealing with information encoded into discrete amplitude levels.

In this chapter, we present the design and implementation results of a 10 Gbps wireline transmitter chip as a first major step towards the vision of a serial link for ELCA's digital RF front-ends. After giving an initial overview, we will cover the design of the core circuit blocks, specifically the digital 16:1 multiplexer and analog driver stage. The main focus points are the required careful thinking, achieving 10 Gbps operation while meeting the specified design targets. Moving on, we briefly discuss other auxiliary (but critical) subcircuits that completed this wireline transmitter chip. We close off with an overview of the layout, critical layout simulations.

We will briefly discuss initial, promising measurements of the fabricated chip. Full measurement results did not make it for this thesis. We did fully prepare for the measurements (e.g. PCBs are fully designed and fabricated), but the transmitter silicon returned late into the project when the receiver design was a priority.

3.1 Overview and Challenges

Fig. 3.1 breaks down the wireline transmitter chip into its constituent subcircuits. To review, the design goals of this wireline transmitter are a link bit error rate (BER) $< 10^{-8}$ at an output data rate of 10 Gbps using a 16:1 serializer. It is worthwhile to explore the overview in more detail and discuss the high-level design choices in view of these requirements.

3.1.1 Core Circuit Styles

Power Efficiency To build a wireline driver, one can choose between a current-mode (I-mode) or voltage-mode (V-mode) driver as first design consideration [8]. Fig. 3.2 simply illustrates these two options. The logical trade-off between these two circuit styles is power versus flexibility. V-mode drivers use switches to a fixed voltage reference to generate the pulse amplitude, while I-mode drivers consume a constant I (i.e. constant power) that is steered to one of two resistors to generate a voltage level. In effect, the V-mode driver is more power efficient judged by the final stage only.

However, in implementation we should also consider the *pre-drivers* to the switches of the V-mode drivers: they are very large due to the low-ohmic output resistance that the switches should equate to, in order to match the channel impedance. As such, looking at the overall picture it is hard to judge whether a V-mode transmitter is certainly more power efficient than an I-mode transmitter. The V-mode designer must be innovative (e.g. employ stacked drivers to recycle current) to compensate for their pre-driver power.

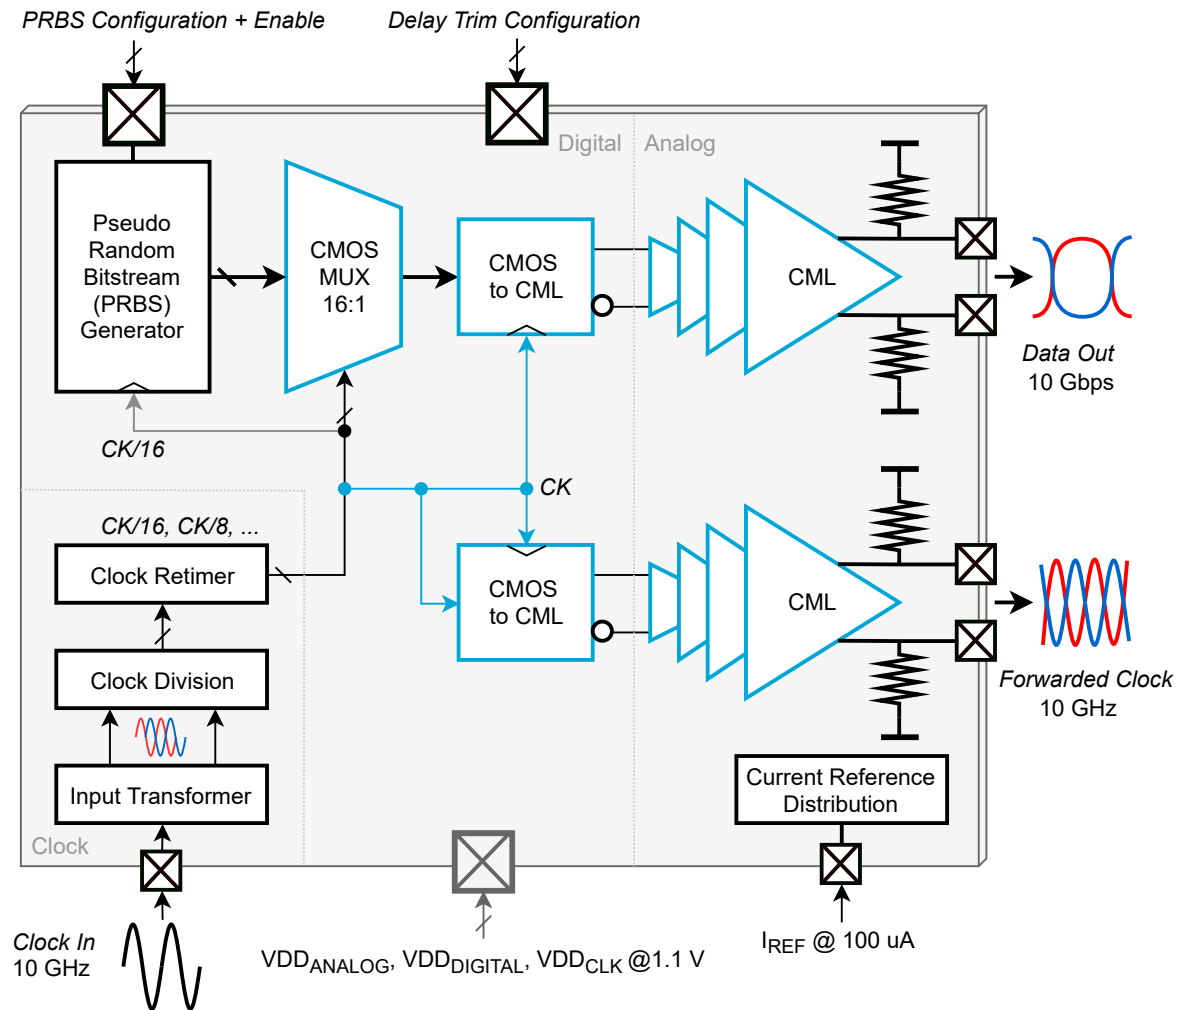


Figure 3.1: Overview of the wireline transmitter presented in this chapter. The transmitter core is highlighted in blue and the other blocks are regarded auxiliary circuits. CML stands for *current mode logic*.




Figure 3.2: General distinction of wireline transmitters. (a) Current mode, (b) voltage mode.

Driver Implementation From the perspective of signaling quality, the constant current draw is another appealing feature of the I-mode driver. There are no big transient switching currents from the supply, resulting in quieter eye diagrams and common-mode emissions for the same supply decoupling. Similarly, channel matching (termination) is not worrisome for I-mode drivers. In contrast to the need for linear (or linearized) $50\text{-}\Omega$ switches in V-mode drivers, the I-mode driver works with a high-impedance current source, so $50\text{-}\Omega$ resistors do suffice here.

Finally, it is easy to add up signals in the current domain. Therefore, I-mode drivers provide a rather easy pathway towards circuit expansion for equalization or pre-emphasis techniques. This fits our end vision best and motivates mainly the choice to go with an I-mode driver compared to a V-mode driver. The specific I-mode driver type that we propose to use is the *current mode logic* type (CML). As presented in Section 3.3, the design of efficient, fast CML driver chains requires a non-traditional but structured approach.

Multiplexing Speed Serialization or multiplexing is also a function that is easily performed with I-mode circuits. I-mode drivers should reach high-speeds without much effort due to the low voltage swings that are required. However, since we are aiming for a high multiplexing ratio, it is not smart to go with a pure I-mode design.

With high multiplexing ratios, an exponential number of 2:1 multiplexers (MUXs) are necessary to build a $2^M : 1$ MUX, yielding a high power consumption and tedious layout due to all the differential signals and scattered tail current banks. Therefore, we choose to investigate whether we can push *digital* CMOS MUX to 10 Gbps and bridge this CMOS MUX to the I-mode driver in the next sections.

3.1.2 Functionalities

It is clear that the multiplexer and data driver are the core design blocks of this wireline transmitter chip. However, building a framework around it to allow for operation and testing is separate question that our designed chip accounts for.

For one, to test the chip, the unserialized 16-bit words should be given to the transmitter. Some obvious options to consider are a large parallel digital interface at hundreds of Mbit per second per link or having an on-board SRAM memory block. The former solution is impractical in terms of pin numbers and both defeat the purpose of this project. Therefore, we decide to design a pseudo-random bit stream (PRBS) generator. The PRBS generator allows for the measurement of a realistic eye diagram at the output of the transmitter. In addition, a full transceiver can benefit from this on-board PRBS generator to use it for memory-less bit error rate testing.

Another concern for testing a wireline transmitter is the generation of an internal clock signal. In this work, we go for a distributed multiplexer in favor of pushing the speed of digital CMOS as motivated in the next section. To support this choice, our design features a clock divide-and-retchime subsystem operating from a single 10 GHz (nominal) sinusoidal clock. Four divided clocks ($1/2$ down to $1/16$) are required for such a distributed multiplexer. Some trimmable delay blocks placed in critical locations allow for extra flexibility in system operating clock frequencies.

The main 10 GHz clock is provided single-ended. Using an on-chip $50\text{-}\Omega$ transformer¹, we obtain a complementary clock as required for CML and our high-speed flip-flops. In order to maintain design flexibility for the receiver design, we also generate an equal-delay forwarded 10 GHz clock. In other words, the chip has another output path to transmit a clock that matches the data's phase. To this end, we feed the internal 10 GHz clock to an identical CMOS-to-CML and CML driver chain cascade at the point where the data is retimed for the last time.

¹Credits for the design of this transformer go to Mohammadreza Beikmirza.




Figure 3.3: Most basic 2:1 unit MUXs. (a) 3-NAND MUX, (b) TGATE MUX.

3.2 CMOS Multiplexer Design

Multiplexers (MUXs) are commonly encountered in analog-mixed signal circuits and systems, leading to a wide variety of implementations. In this section, we show how to design a fast CMOS MUX. The design removes the need for performing any MUX operation in our (analog) driver circuits. Let us first shortly review the basics of MUXs.

The 3-NAND MUX and MUXs based on variations of transmission gates (TGATES) are the most simple (“unit”) MUX circuits as shown in Fig. 3.3. Similar to circuit styles of driver circuits, the TGATE MUXs are easily adapted to current domain MUXs by inserting the TGATE in the tail of a differential pair.

Generally speaking, for a MUX serializing more than 2 input streams we require a specific clock arrangement to orchestrate the *SELECT* signals. Expanding on the TGATE MUX, we can easily come to the multiphase MUX architectures for more inputs. As shown in Fig. 3.4, this choice is attractive when a multiphase clock is readily available, e.g. in the presence of a quadrature clock divider, we only need to add some AND gates. However, with a high number of inputs such as for our 16:1 serialization goal, it becomes a difficult task to obtain this number of uniformly spaced clock phases.

A more structural approach for building multi-input MUXs is by cascading 2:1 unit MUXs, forming a “MUX tree”. As Fig. 3.5 shows, a MUX tree requires divide-by-2 clocks, as many as $\log_2(M)$ dividers for a $M : 1$ MUX for MUXs that provide a retimed output. If system latency is of concern, this architecture is not the optimal choice due to the M slices that the data must traverse to reach the serial output.

The shift register MUX is a special architecture that does not employ any unit MUX: it is very flexible just like the multiphase MUX while requiring only a simple clock arrangement as shown in Fig. 3.6. The major downside of this architecture is the heavy load of the clock CK running at the same rate as the data posed by the flip-flops, prohibiting reasonable power consumption of the transmitter’s clock buffers.

3.2.1 MUX Tree Speed Limitations

Based on the large serialization factor of 16, we decide to go forward with the design of a MUX tree for our transmitter. But, what makes a MUX tree fast?

To answer this question, we must start with the timing analysis of a simple tree consisting of 3-NAND MUXs. One unit 3-NAND MUX in 40 nm is quite fast² and can have a propagation delay down to ~ 40 ps. The necessary behavioral model of a 3-NAND MUX is shown in Fig. 3.7.

²A standard library NAND has FO4 delay of < 18 ps (fan-out is the relative capacitive load of a logic gate, FO4 denoting fan-out of four), while a standard library 2:1 MUX specifies a > 50 ps FO4 delay. It may be clear that the standard library is not fully optimized for high-speed applications.




Figure 3.4: Multiphase MUX architecture. D_{OUT} is the output data stream, while CK denotes a clock signal.

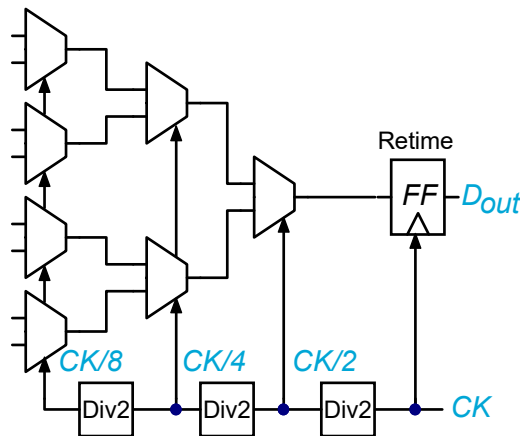


Figure 3.5: MUX tree architecture. The input clock is divided in frequency by two for each slice.

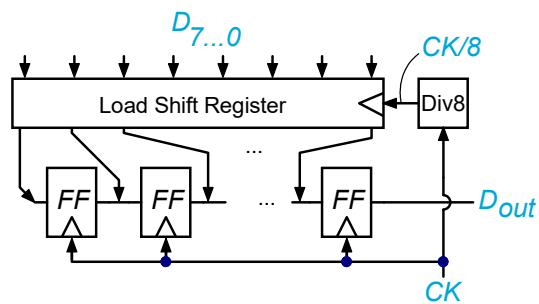


Figure 3.6: Shift register MUX architecture.

3-NAND MUX Tree Suppose that we build a 4:1 MUX tree using the 3-NAND MUX as shown in Fig. 3.8. This results in $\log_2(4) = 2$ slices (layers of MUXs). Now, looking at the propagation of data in time, each piece of data has different timing constraints corresponding to the order in which they appear at output. The most stringent constraint is placed on the first data that needs to appear, in this case $D1$ in pink. When all clocks rise, $D1$ must propagate through to *two* slices. Assuming that we have an input data register and an output retimer that are fast, $D1$ has at most t_{CK} seconds to




Figure 3.7: Behavior of a 3-NAND MUX over time.

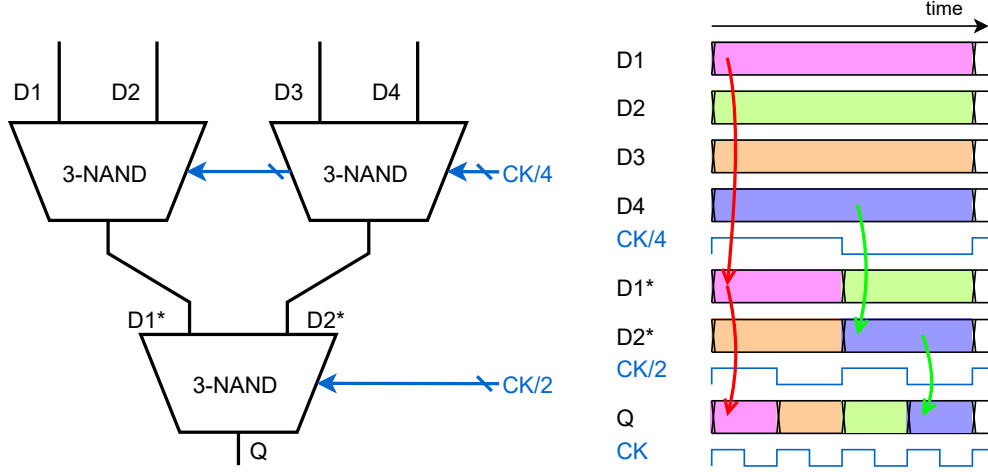


Figure 3.8: Data propagation in a 3-NAND MUX tree.

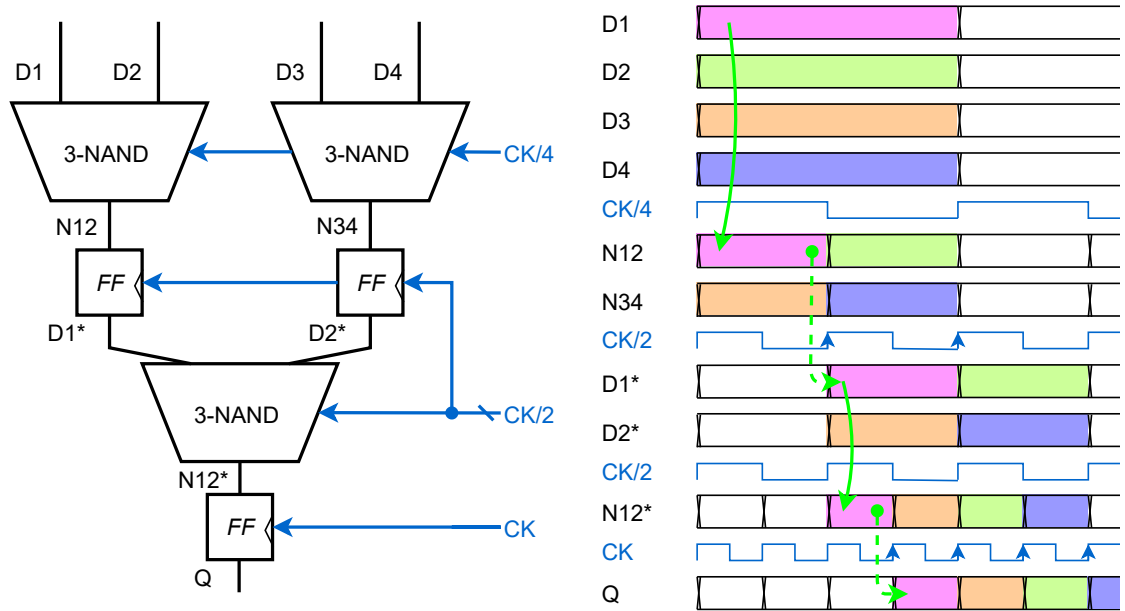


Figure 3.9: Data propagation in the proposed pipelined 3-NAND MUX tree architecture.

complete its traversal. As such, the output data rate $f_{\text{MUX,NAND}}$ of such an architecture for a $2^M : 1$ MUX is limited to

$$f_{\text{MUX,NAND}} \leq 1/t_{\text{CK,min}} = \frac{1}{\log_2(M) \cdot t_{\text{p3NAND,min}}}, \quad (3.1)$$

where $t_{p3NAND,min}$ is the minimum delay of the 3-NAND MUX. Thus, the 3-NAND MUX is a fast piece of logic, but on its own not suitable for building high-speed MUX trees with many slices. For example, for a 16:1 MUX this circuit can reach at most 6.25 Gbps. Namely, $1/(\log_2(16) \cdot 40 \text{ ps}) = 6.25 \text{ Gbps}$.

Pipelined MUX Tree We propose to improve the speed of this tree by addition of more flip-flops. Looking closer at the previous example, we can see that $D3$ enjoys a more comfortable propagation. Indeed, it has one t_{CK} to pass the first slice and needs to pass the second slice only when $CK/2$ falls low. We can force this relaxed data traversal for all data inputs with a pipelined data structure as shown in Fig. 3.9.

Looking again at $D1$ in pink in this pipelined MUX tree, we can see that the pink data is preserved by each subsequent flip-flop and has one t_{CK} available to make only one final slice instead of having to cross all slices at once. In general, the clock timing constraint for data transmission between flip-flops is [3]

$$t_{CK} > t_{CQ,FF} + t_{pLogic,min} + t_{SETUPFF} \quad (3.2)$$

Here, $t_{CQ,FF}$ is the clock edge to output delay of the flip-flop and $t_{SETUPFF}$ is the minimum (“setup”) time that the data must be present in advance at the next flip-flop in order to guarantee successful sampling. Thus, for this pipelined 3-NAND MUX tree, we can specify a maximum data rate of

$$f_{MUX,NAND+FF} \leq \frac{1}{t_{CQ,FF} + t_{p3NAND,min} + t_{SETUPFF}} \quad (3.3)$$

removing the dependency of M ! For large serialization factors (M), this flip-flop approach can pay off. Then, the speed of unit MUX and flip-flop limit the speed in exchange for an increase of latency and additional load for the low-frequency clocks.

To conclude for now, by inclusion of a pipelined data structure we make a sufficiently fast, large 3-NAND MUX tree. However, to push for operation at 10 Gbps or beyond, we will require fast registers with small propagation delays and setup/hold times as they add up to the data timing budget. The standard 40 nm CMOS library does not provide sufficient speed at all in this respect: its flip-flops have clock-to-Q delays of at least 100 ps, calling for a custom solution.

3.2.2 10 Gbps Inherently Pipelined MUX Tree

At this point, we are seeking for faster flip-flops and possibly even faster unit MUXs (compared to the 3-NAND MUX): the “Clocked CMOS” (C2MOS) digital circuit style can be exploited well for this purpose. In 2010, [9] showed a unit MUX based on a special arrangement of “clocked inverters”, without recognizing that they are in fact C2MOS latches. In the following, we show how to arrive at an optimum sizing arrangement of their C2MOS-based unit MUX and add our own MUX tree timing analysis. We intuitively explain the underlying mechanisms at play to motivate the final 16:1 MUX design. Our 10 Gbps wireline transmitter ultimately features an *inherently pipelined* C2MOS MUX tree.

C2MOS is a *dynamic* CMOS circuit style, meaning that logic functions are only performed at clock events (e.g. when the clock is high). As for most dynamic CMOS circuits, C2MOS achieves its dynamic characteristics by preserving the output of logic functions in the charge domain. For example, the C2MOS inverter as shown in Fig. 3.10 can act as a inverting latch, either transparent-high or transparent-low³.

³to prevent confusion we avoid the use of positive latch and negative latch in this text: *transparent-high* means that the latch can perform its logic when the clock = 1 and holds its output when the clock = 0.




Figure 3.10: A transparent-high C2MOS latch. The latch inverts when its high.

Sizing Consideration Sizing a C2MOS latch for speed depends on minimizing the RC delay between two latches. For this MUX application, it is desirable to have the clocked transistors near the output nodes to ensure minimal charge leakage and unwanted data feedthrough. Unwanted clock feedthrough is minimized by keeping the clocked transistors small⁴. The degree of freedom thus lies in the size of the outer transistors that act as the input node D of the latch. For a cascade connection of C2MOS latches, we find that a relative $3\times$ width of the outer transistors (NMOS and PMOS connected to ground and V_{DD}) compared to the width of the inner ones is a sweet spot for speed, minimizing the resulting clock-to-output delay $t_{CQ,C2MOS}$. As shown later, $t_{CQ,C2MOS}$ ultimately limits the speed of the this C2MOS tree MUX.

Inherent Pipelining Knowing the optimum performance of our C2MOS latch allows for evaluation of the speed of a MUX tree based on the unit MUX presented in [9]. We make the observation that this C2MOS unit MUX actually consists of two elements: a C2MOS master-slave flip-flop and a dynamic latch inverter. Fig. 3.11 shows this important model and its corresponding behavior. In our view, this C2MOS MUX has two major advantages:

- Compared to a standard 3-NAND MUX, the propagation delay of the data after clock switching is very low for this C2MOS MUX (14 ps compared to 40 ps). The dynamic nature allows for a simple take-over of the output node.
- Each unit MUX enjoys a built-in flip-flop that we can use to relieve the datapaths with the most tight timing constraints without needing explicitly placed pipelining flip-flops. Moreover, the C2MOS flip-flop operates fast as its latch is only a slightly degenerated inverter from CK 's perspective: the clock-to-Q delay is 20 ps at most.

As such, we can more intuitively arrive at the fast tree expansion in Fig. 3.12 as briefly presented in [9]. Compared to the timing diagram shown earlier for the simple 3-NAND MUX tree, we can see that the C2MOS flip-flop (FF) inputs save data such as $D1$ in pink that normally would have to traverse the tree at once. The tree is preceded by an extra slice of latches to ensure a well-retimed input word and additional inverters to compensate for unwanted inversions from input to output due to the inverting nature of the clocked inverter part of this unit MUX.

⁴Since C2MOS utilizes a rising and falling clock, the clock feedthrough is already partially canceled.

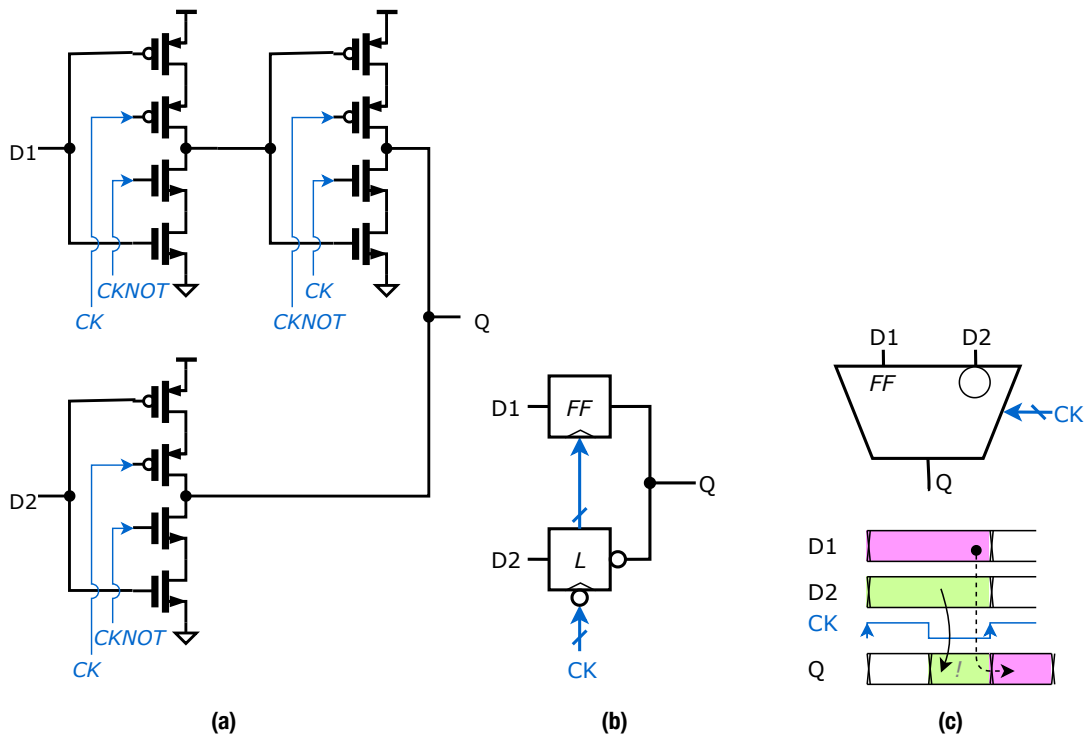


Figure 3.11: Fukuda's [9] C2MOS multiplexer. (a) Full circuit, (b) our functional recognition, (c) symbol and timing diagram. The exclamation mark in the timing diagram denotes an inversion.

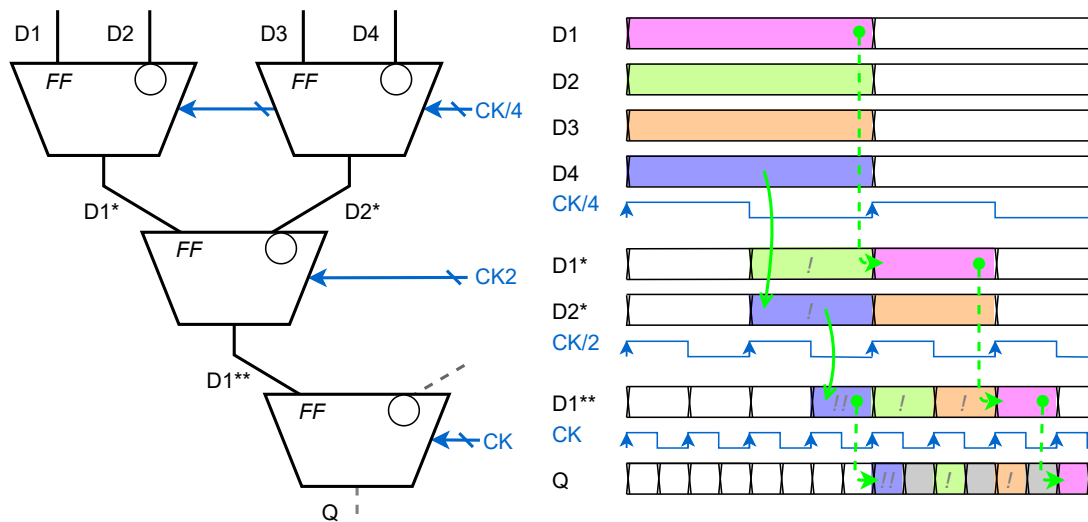


Figure 3.12: Data propagation in a C2MOS MUX tree. The inherent pipeline characteristics arise from the built-in flip-flops. Note that the applied clocks must be complementary and that the order in which the input data appears at the output is different from previous MUX analyses (reversed).

Speed Limitations We can quantitatively predict the maximum data rate of this C2MOS MUX tree. Shown in Fig. 3.13 are the possible combinations of data traversal situations. A C2MOS MUX can act as flip-flop or inverter, so there are four possible situations. The most stringent situation appears for $D3$ and $D4$, where $0.5t_{CK_{next}}$ is available to arrive at the inverting input terminal (INV, marked by the circle) of the next MUX.

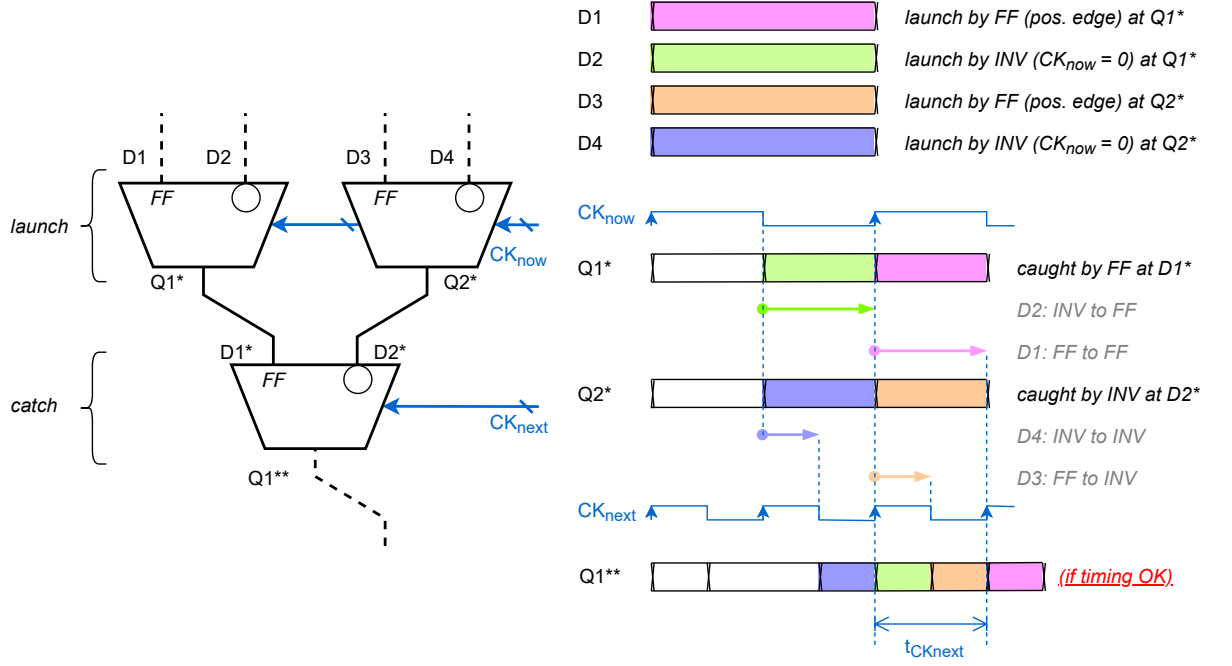


Figure 3.13: Timing constraints for all four combinations of inverter/flip-flop to inverter/flip-flop propagation in a C2MOS MUX tree. *D3* and *D4* have the most stringent constraints.

We can motivate that the C2MOS MUX is fast for two reasons. Similar to Eq. 3.2, this situation demands

$$0.5t_{CKnext} > t_{CQ,C2MOS} + t_{pLogic,min} + t_{SETUP,INV} \quad (3.4)$$

There are no discrete 3-NAND MUXs placed here in between the flip-flop elements in the tree (compare to Fig. 3.9). In fact, the C2MOS MUXs are directly connected by a metal connection, so that $t_{pLogic,min} \simeq 0$! Secondly, when acting as an inverter, $t_{SETUP,INV} = 0$ as the inversion operation continues after the clock edge. Therefore,

$$f_{CKnext} > \frac{1}{2t_{CQ,C2MOS}} \quad (3.5)$$

For a C2MOS multiplexer, the frequency of the last clock equals one half of the output bit rate $f_{MUX,C2MOS}$. So, we finally find that the C2MOS MUX tree is speed limited to $1/t_{CQ,C2MOS}$.

Conclusion and Final Design We conclude that the C2MOS MUX tree is the fastest MUX tree as summarized in Table 3.1

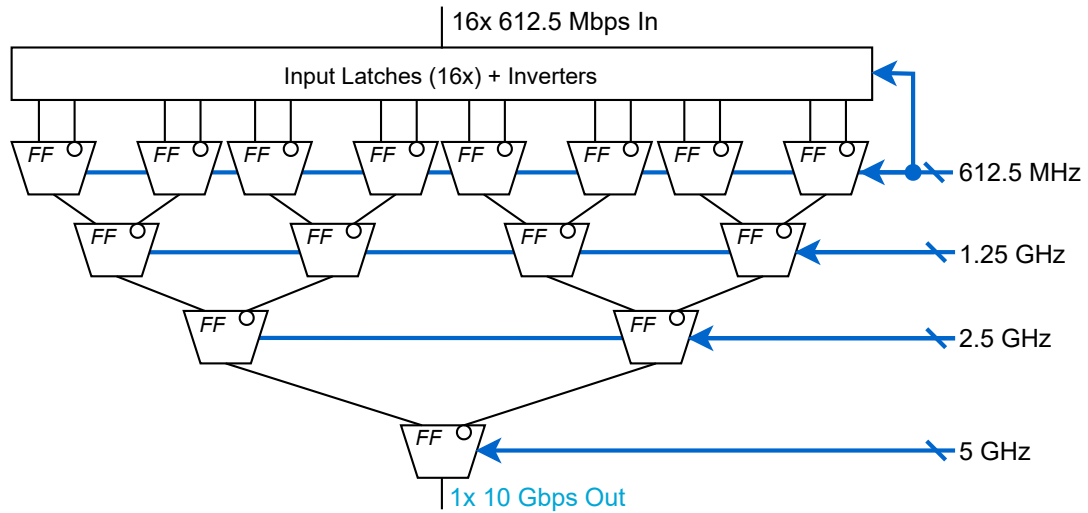
As such, we designed a full 16:1 MUX based on the C2MOS MUX as shown in Fig. 3.14. Post-layout simulations indicate that $t_{CQ,C2MOS} \simeq 25$ ps, so the realized tree can potentially run at 40 Gbps. However, as discussed in Chapter 2, this is a typical data rate at which a digital system is primarily limited in speed by its rise and fall times. For a realistic capacitive load, we measure an output rise/fall time of ~ 40 ps. Thus, our 16:1 40 nm CMOS MUX design is safe to run up to roughly $1/(2 \cdot 40 \text{ ps}) = 12.5$ Gbps.

3.3 Current-Mode Logic Driver Design

Our aim is to bring the output data to a 100-ohm differential channel in matched manner. Thus, the design must feature a relatively large driver (in terms of drive capability) to differentially bring out

Table 3.1: Summary of MUX tree speed limitations when rise and fall times are negligible.

Architecture	Maximum Data Rate R_B
3-NAND	$\frac{1}{\log_2(M) \cdot t_{p3NAND,min}}$
3-NAND and pipelined	$\frac{1}{t_{CQ,FF} + t_{p3NAND,min} + t_{SETUP,FF}}$
C2MOS MUX	$\frac{1}{t_{CQ,C2MOS}}$

**Figure 3.14:** Complete designed 16:1 inherently pipelined 10 Gbps+ MUX tree. The latches help to retime the incoming data while the inverters nullify the inversion in the MUX tree where necessary.

the data from the 10 Gbps CMOS MUX tree.

As discussed earlier, we will implement an I-mode driver. As shown in Fig. 3.15, I-mode drivers come in single-bridge or double-bridge mode, often referred to as “Current-Mode Logic” (CML) and “Low-Voltage Differential Signaling” (LVDS) transmitters, respectively. LVDS-type drivers exhibit superior supply noise and variation rejection, but require a significant headroom in order to be comfortably biased. Moreover, LVDS drivers need additional common-mode feedback.

In the view of our vision for a universal serial link at ELCA, we favor a highly integrable system and prefer the CML driver. This allows us to design for a complete wireline transmitter from a single low 1.1 V supply voltage.

In this section we propose a structural design approach that we developed to be able to size for a CML driver chain without a trial-and-error method. In fact, we are able to show the design space that we are working with. Finally, to convert CMOS to complementary differential signals, we designed an elegant circuit solution that achieves near-zero skew outputs.




Figure 3.15: Common current mode driver circuits. Left: CML, right: LVDS. An LVDS driver requires in addition a common-mode feedback loop.

3.3.1 Current Mode Logic Chain Issues

The CML unit cell is quite simple as it resembles a differential pair. For CML, the cell is not purposely performing amplification, so linearity and gain are not the prior design goals. Let us highlight the critical issues related to CML chain design that are not obvious from a typical analog design perspective.

Switching Efficiency The ideal differential output swing of a CML stage depends on the load resistance R_L and steered current I_{TAIL} . For example, $R_L = 100 \Omega$ and $I_{TAIL} = 1 \text{ mA}$ yields an ideal peak amplitude of $100 \text{ mV}_{PK,DIFF}$. One of the two output nodes stays at V_{DD} while the other node drops by $V_{A,PK,DIFF}$. Thus, the output common-mode voltage $V_{CM,OUT}$ depends on the swing as follows

$$V_{A,PK,DIFF} = R_L I_{TAIL}, \quad (3.6)$$

$$V_{CM,out} = V_{DD} - 0.5V_{A,PK,DIFF} = V_{DD} - 0.5R_L I_{TAIL} \quad (3.7)$$

If the input voltage swing is too small under incomplete switching conditions, it leads to a decreased output swing. We can define the switching efficiency in terms of drain currents of M1 and M2:

$$\eta_{SW} = 100\% \cdot \max(I_{D1} - I_{D2}) / I_{TAIL} \quad (3.8)$$

With perfect square law devices, η_{SW} reaches 100 % when the differential input voltage equals $\sqrt{2}V_{GT,EQ}$ [10]. $V_{GT,EQ}$ denotes the overdrive voltage $V_{GS} - V_T$ when M1 and M2 draw $I_{TAIL}/2$ (“equilibrium”).

For completely switched short-channel devices ($\eta_{SW} > 90 \%$), we observed a strong offset and much higher coefficient for the required input swing as a function of the equilibrium overdrive voltage. We notice that two effects are at mainly play that cause this deviation with respect to the typical $\sqrt{2}V_{GT,EQ}$ guideline:

- The drain voltage of the conducting transistor will fall. Channel length modulation and drain-induced barrier lowering will cause an effective increase of the threshold voltage of this transistor, reducing its expected steering ability. The converse happens for the other transistor that should turn off.
- In a CML chain, any increase of differential swing results in a reduction of common-mode voltage. Increasing the input swing to improve switching efficiency will lower the input common mode and the source node of M1 and M2 follows. The effective equilibrium overdrive increases as a consequence of the body effect (V_{SB} reduces).




Figure 3.16: Normalized required peak differential voltage $V_{\text{SWITCH},90\%}/V_{\text{GT},\text{EQ}}$ versus equilibrium overdrive voltage $V_{\text{GT},\text{EQ}}$ for 40 nm CMOS (simulated). Only for very strong saturation (moving to the right) the required switching voltage will reach the pure square law prediction.

Fig. 3.16 shows this trend: we will use this data in our design procedure. The data indicates that complete switching and high $V_{\text{GT},\text{EQ}}$ will not easily go hand in hand with a low supply voltage of 1.1 V.

Voltage Level Constraints A high switching efficiency yields a power-efficient design. But, there are more constraints: the swing and common mode at inputs and outputs cannot be of arbitrary value! To list, all important constraints in terms of the output amplitude of a driving stage in a CML chain are:

1. The conducting transistor of this differential pair must stay in saturation for high-speed operation. The gate will rise to V_{DD} for this transistor while its drain drops by the designed peak output amplitude. Therefore, its drain must not fall more than one threshold voltage below its gate⁵. In other words, $V_{\text{A},\text{PK},\text{DIFF},\text{OUT}} < V_{\text{T}}$.
2. The tail of the next driven CML stage must also stay in saturation to ensure the expected current. The input common-mode for the next stage should not be too low:

$$\begin{aligned} V_{\text{CM},\text{OUT}} &< V_{\text{GS},\text{EQ}} - V_{\text{X},\text{MIN}} \\ V_{\text{A},\text{PK},\text{DIFF},\text{OUT}} &= 2(V_{\text{DD}} - V_{\text{CM},\text{OUT}}) \\ V_{\text{A},\text{PK},\text{DIFF},\text{OUT}} &< 2(V_{\text{DD}} - V_{\text{GS},\text{EQ}} - V_{\text{X},\text{MIN}})_{i+1} \end{aligned} \quad (3.9)$$

where $V_{\text{X},\text{MIN}}$ denotes the minimum headroom required for the tail (of stage $i + 1$).

3. Given a minimum overall switching efficiency and equilibrium overdrive, the output amplitude must be greater than the prescribed voltage $V_{\text{A},\text{PK},\text{DIFF},\text{OUT}} > f(\eta_{\text{SW}}, V_{\text{GT},\text{EQ}})$ dictated by the previously shown trends.
4. A low swing is susceptible to noise. For a given noise margin (NM, referred to the input) and sizing, the input swing $V_{\text{A},\text{PK},\text{DIFF},\text{IN}}$ must be greater than $V_{\text{min},\text{NM}}$. In most cases, the threshold for switching efficiency exceeds $V_{\text{min},\text{NM}}$.

Fig. 3.17 shows an overview of these constraints. We use these observations to draw the CML design space after this subsection.

⁵if there is no strong inversion, this requirement becomes even more strict since V_{DSAT} does not perfectly follow V_{GT}

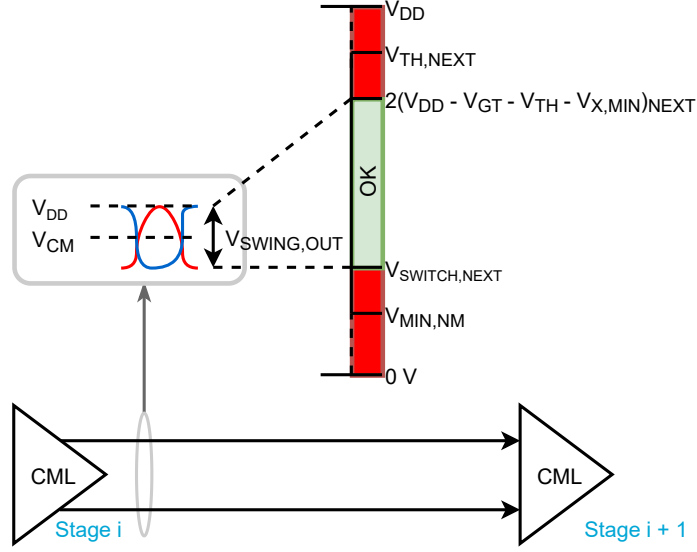


Figure 3.17: Operational constraints for voltage swing in a CML driver chain.

CML Tail The current tail of the CML stage directly dictates the achievable output swing accuracy and limits the output swing of a preceding CML stage. We should also keep in mind that node V_X preserves some voltage headroom for the tail but will move up and down. Due to mismatch of M1 and M2, imperfect complementary inputs and non-linear $C_{GS1,2}$, V_X can move tens of mV upon switching.

To preserve speed during switching, we must not break the headroom margin at V_X . A large absolute sizing in favor more quiet V_X voltage, larger headroom or good stage-to-stage matching leads to a large junction capacitance C_{TAIL} at V_X . But in sizing the tail, the parasitic node capacitance cannot be allowed to grow arbitrarily. A large C_{TAIL} can incur a significant *extra* current transient to ground:

$$I_{EXTRA} = C_{TAIL} \frac{dV_X}{dt} \quad (3.10)$$

If little supply decoupling is available, these transients can be a threat for the common-mode signal integrity.

For our driver design, we implement all tails using a default cascode structure. Non-cascoded tails are difficult to include for a structured design approach due to their weak current accuracy under unknown voltage at V_X .

Speed Considerations It would be practical if sizing guidelines from digital CMOS design can be carried over to CML design. For example, in digital CMOS, high speed is generally achieved with low propagation delays between registers, calling for design guidelines. In turn, overall digital propagation delay is minimized with a fixed fan-out ratio (e.g. FO4) and per-stage delay is minimized by minimizing the fan-out [3].

Unfortunately these mechanisms does not hold for CML design. There are a few distinct differences:

- The delay due to parasitic loading can be tuned by sizing the CML transistor pair. R_L is independent of the transistor. For CMOS, the intrinsic delay is fixed as here the drive strength is inversely proportional to its own parasitics.
- The drive strength of a CML stage represented by R_L can also be independently tuned.

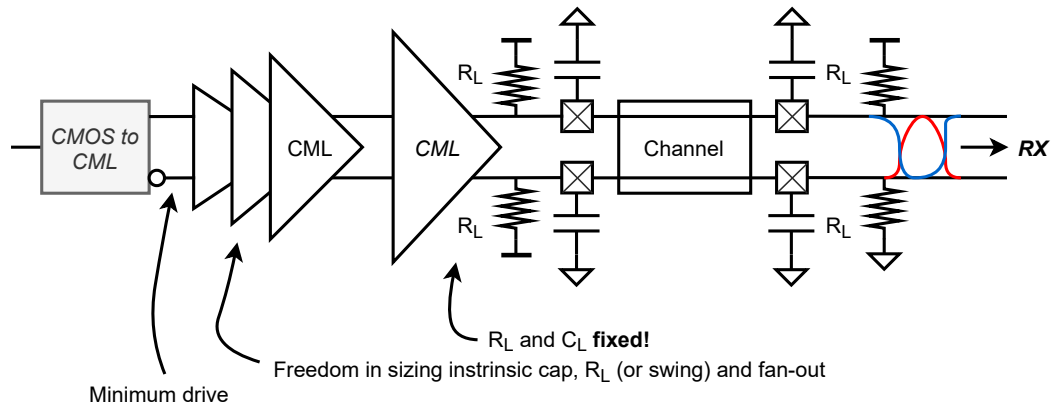


Figure 3.18: Variation of degrees of freedom along a CML driver chain.

- CML circuits are semi-analog circuits. Under the earlier listed voltage level constraints, we can also perform optimum sizing for speed at a chosen tail current (i.e. maximize f_T under given constraints).

3.3.2 Driver Design

To complete the above discussion on design for speed, let us have a look at the complete driver chain as shown in Fig. 3.18. On the left side of the chain, the CMOS circuitry has low drive strength and must be presented with a low capacitive load to preserve the overall system speed. On the right side of the chain, the CML is actually fixed. To match channel reflections $R_L = 50 \Omega$ and the capacitance of the pads or channel response will dominate the time constant seen at the receiver. We are only interested in scaling drive strength along the chain, not the swing.

Procedure Due to the minimum degrees of freedom at the last CML stage and based our previous design insights, we propose to employ the following design procedure for a CML driver chain design problem, summarized by depiction in Fig. 3.19.

1. We start at the last driver stage, the output driver. We identify the fixed load resistance. Then, we pick a safe output amplitude based on the BER requirement and derive the necessary tail current as $I_{\text{TAIL,OD}} = V_{\text{A,PK,DIFF,OD}}/R_L$. Note that R_L also includes the termination resistance at the receiver.

Generally, a CML TX is not very noisy or suffering from large mismatch due to the low-gain stages and large area for current sources and load resistors. For short-channel transistor designs, assuming $20 \text{ mV}_{\text{rms}}$ of pure noise is a safe start.

2. To size the transistor pair, we must make the design space visible. An example is shown in Fig. 3.20. Recall that having a narrow transistor pair will result in a high switching threshold, while limiting the swing of the previous stage to keep its tail in saturation. Drawing these two boundaries for different switching efficiencies, the design space becomes clear.

We must pick a desired switching efficiency curve that keeps the required driving signal below the maximum allowed swing. A trade-off between switching efficiency and high V_{GT} appears: power efficiency versus speed (presented load).

Generally, it is wise to have a switching efficiency $> 90 \%$ to minimize reduction of expected output swing and to keep the design procedure straight forward. Then, V_{GT} should be chosen

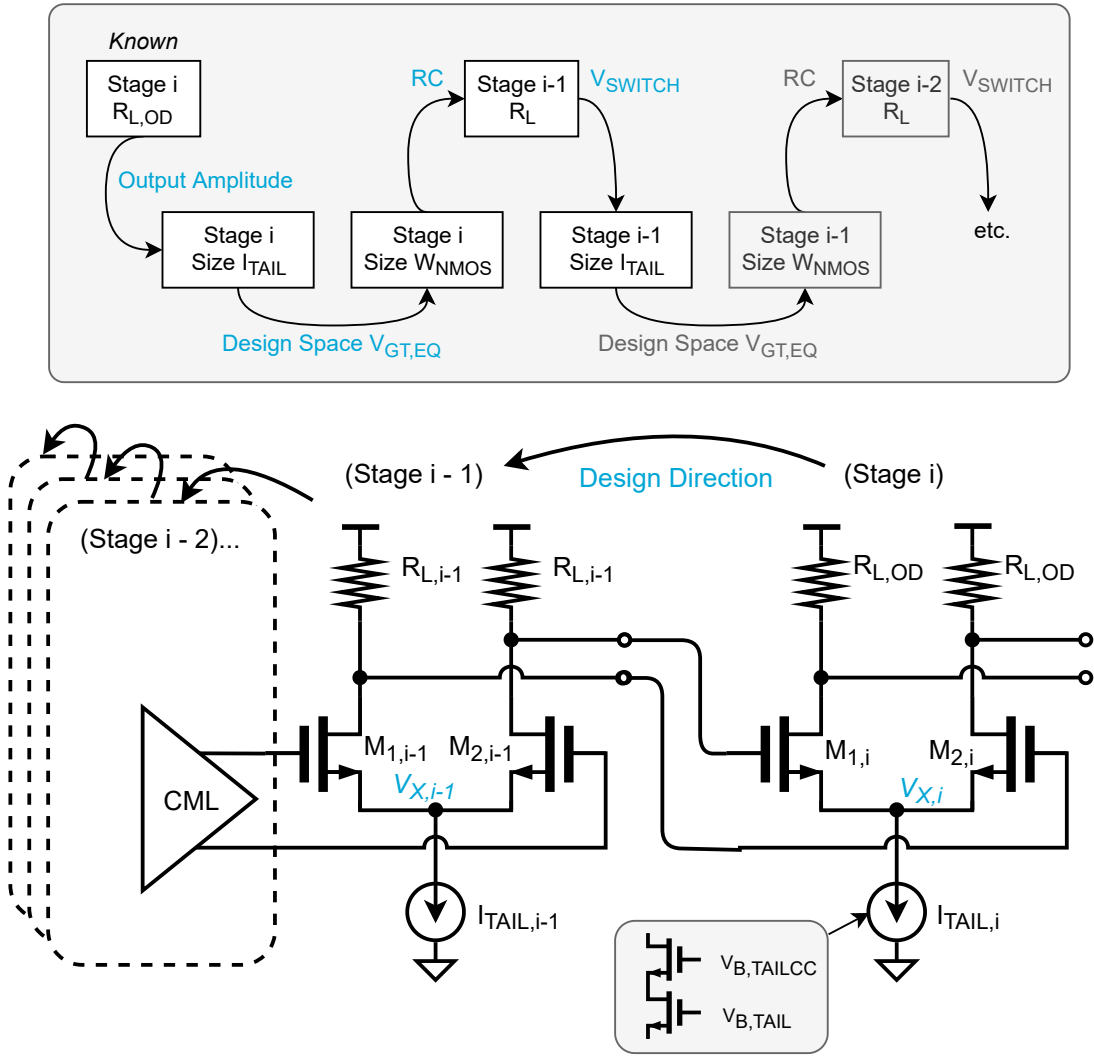


Figure 3.19: Proposed output to input CML driver chain design procedure.

maximum (narrow transistor) within the remaining design space to allow for quick tapering of the chain's sizing⁶

3. A quick simulation of the resulting output driver includes channel effects to measure the resulting rise and fall time at the receiver end. The corresponding RC-constant must be maintained along the other preceding CML stages. In other words, these stages must not be slower. An operating point simulation reveals the total gate capacitance $C_{gg,OD}$. Based on this value, and the previously obtained maximum RC-constant, pick R_L for stage $i - 1$.
4. The necessary switching voltage was obtained in step 2. The tail sizing of stage $(i - 1)$ $I_{tail,i-1}$ must be chosen such that this switching amplitude is achieved for the output driver.⁷ Stage $i - 1$ is now fully designed.
5. Sizing of stages $i - 2, i - 3$, etc. continues by maintaining the same RC-constant and I_D/W -ratio. That is, find C_{gg} and R_L for RC, then I_{TAIL} for switching, followed by sizing for the same I_D/W .

⁶The sizing can be settled on by applying a gm/I_D characterization, thereby obtaining a convenient number for I_D/W that corresponds to the chosen V_{GT} .

⁷Assuming the same switching efficiency for all stages, we can size the transistor pair of $i - 1$ for the same I_D/W -ratio as the output driver.




Figure 3.20: Design space of choice for a $V_{GT,EQ}$ and input amplitude combination marked in white, assuming a minimum tail headroom of 300 mV in 40 nm CMOS. A differential pair with high overdrive voltage cannot sustain a CML input amplitude higher than the blue curve. It is only possible to relax the design space by sacrificing the requirement on η_{SW} .

Table 3.2: Summary of simulated 4-stage CML driver chain after following our proposed procedure. The low stage delays allow for easy 10 Gbps data passing and the targeted 90% switching efficiency is met on average.

	Stage 1	Stage 2	Stage 3	Stage 4
Tail Current I_{TAIL}	100 μ A	330 μ A	1.25 mA	6 mA
Load Resistance R_L	2.5 k Ω	1.4 k Ω	310 Ω	50//50 Ω
NMOS Gate Width ⁹	800 nm	3 μ m	10 μ m	45 μ m
Output Swing (Diff.)	0.68 V_{PP}	0.82 V_{PP}	0.63 V_{PP}	0.3 V_{PP}
Delay	10 ps	13 ps	13 ps	1.2 ps
Switch Efficiency	97%	86%	92%	90%

The design loop stops when C_{gg} is small enough for the CMOS stages to drive at the intended operating speed.

Design Result Our design approach leads us to the configuration as shown in Table 3.2. Following Fig. 3.20, we maintain a headroom of 300 mV for the tail. With just 4 stages and over 90 % switching efficiency, we are able to scale up from minimum size digital circuits to driving a 25 Ω load by sizing for maximum allowable $V_{GT,EQ}$ ⁸. By simulation with accurate RF transistor models, we confirm that the chain passes 10 Gbps data with success. The simulation includes any non-ideal driving effects by inclusion of the CMOS-to-CML circuit presented in the next subsection. The last two drivers dominate the total supply current consumption of less than 8 mA at 1.1 V.

⁸We should not forget that there is another 50 Ω (single-ended) termination at the receiver side. Therefore, the driver must drive a 50//50 = 25 Ω load in total.

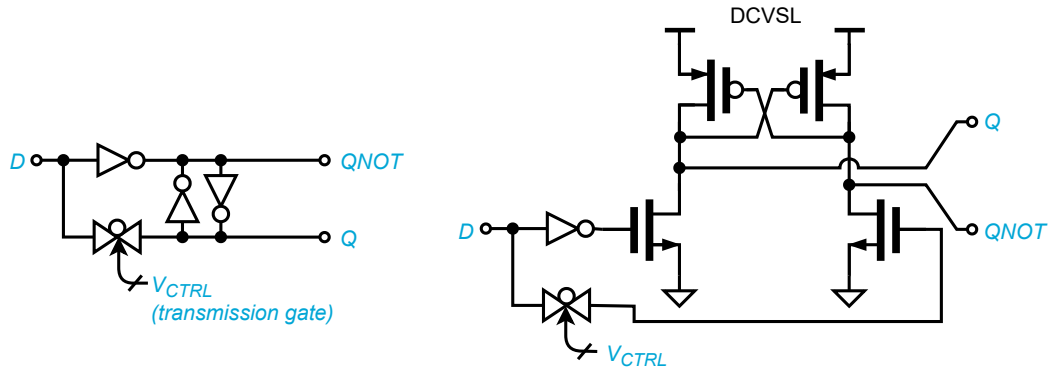


Figure 3.21: Basic signal complement generation options. Left: cross-coupled inverters, right: DCVSL.

3.3.3 CMOS-to-CML

The CML driver chain still requires a CMOS-to-CML chain to complete the driver and allow direct integration with the multiplexer. The two necessary functions are as follows:

1. Generate **complementary CMOS** signals, without significant skew or phase overlap.
2. **Level-shift** the CMOS voltages (rail-to-rail) to CML voltages (700 mV – 1.1 V), without slowing down the chain.

We precede the four-stage driver with a CMOS complementary signal generator and CML level-shifter.

Complementary Signal Generation The need for generation of complementary digital data signals is commonly addressed (for clock signals) by inverting the signal of interest and applying a compensating delay element to the original signal. The compensating element is often a (controlled) transmission gate.

At high data rates a significant phase overlap can still remain, which can be further reduced by addition of regenerative feedback. Basic examples as shown in Fig. 3.21 are cross coupled inverters, or cross-coupled PMOS loads as encountered in the DCVSL circuit style (Differentially Coupled Voltage Switching Logic). A different, common approach to this problem is to run an exact copy of the existing digital circuits which is fed the same data, but inverted. This approach is power hungry.

For this wireline transmitter, we propose a non-standard solution for this problem. The to-be-complemented data emerges from the MUX, so an in-phase clock can be used to simultaneously retiming the original and an inverted data stream on the falling edge as shown in Fig. 3.22. To simplify the solution even further, we can remove the slave latch and just use a transparent-high latch instead. By using a C2MOS latch and assuming that the latches can be placed nearby to enjoy identical clock signals, we are able to reduce the skew to less than 100 fs.

Thus, our latch-based solution performs better in terms of phase overlap (or skew) than the basic complementary signal generation circuits and does not require a full duplicate of the preceding logic, saving power.

Level-Shift to CML For a level-shifter to convert CMOS (0 – 1.1 V) to CML voltages (700 mV – 1.1 V), typical logic level-shifters cannot be easily used. Such traditional level-shifters only shift the logic-high voltage up or down.

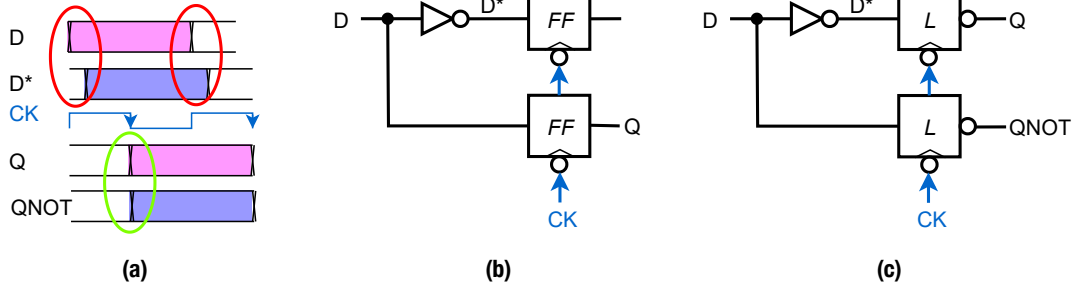


Figure 3.22: Proposed retiming solution to achieve skewless complement generation. The flip-flops can be replaced by latches. (a) Timing diagram, (b) flip-flop based implementation, (c) proposed latch solution.




Figure 3.23: Four-transistor level translator for CMOS to CML.

In this case, the logic-low voltage needs to shift up, requiring a circuit with a weak pull-down or an always-on pull-up. A simple option that we employ is to take a differential pair and take away its tail to allow operation with rail-to-rail input voltages (i.e. pseudo-differential logic). To arrive at an easy-to-layout solution, we replace the load resistors by active PMOS loads as shown in Fig. 3.23. This circuit's output swing is not very accurate, but that is of no concern as long it is able to fully switch the first CML stage.

3.4 Auxiliary Circuit Design

It requires a significant step to move on from core circuit designs to a layout, testable chip design. In this work, the most instrumental auxiliary subcircuit designs that complete the transmitter chip are the pseudo random bit stream (PRBS) generator and the aligned divided clock generator.

Here, we only provide a brief overview of these two auxiliary circuit designs. For the sake of reading agility, we omit a detailed explanation and description of the clock generation circuit. Instead, we provide the interested reader an elaborate text on the auxiliary circuits of this chip in Appendix B, including smaller design work such as current references.

Pseudo Random Data The inclusion of a PRBS generator avoids the need for on-chip storage of transmit data for this chip. A basic PRBS generator consists of a string of cascaded flip-flops with a XOR that feeds some of the intermediate nodes back to the input [6]. In our proposed transmitter design, we incorporate multiple selectable feedback options for a PRBS generator that can generate a random bit stream that is non-periodic up to the $2^{16} - 1$ th bit as shown in Fig. 3.24. As we need a 16-bit word for the multiplexer, the PRBS solution remains compact as we feed all nodes in the loop to the multiplexer as opposed to using only the output node (D15).

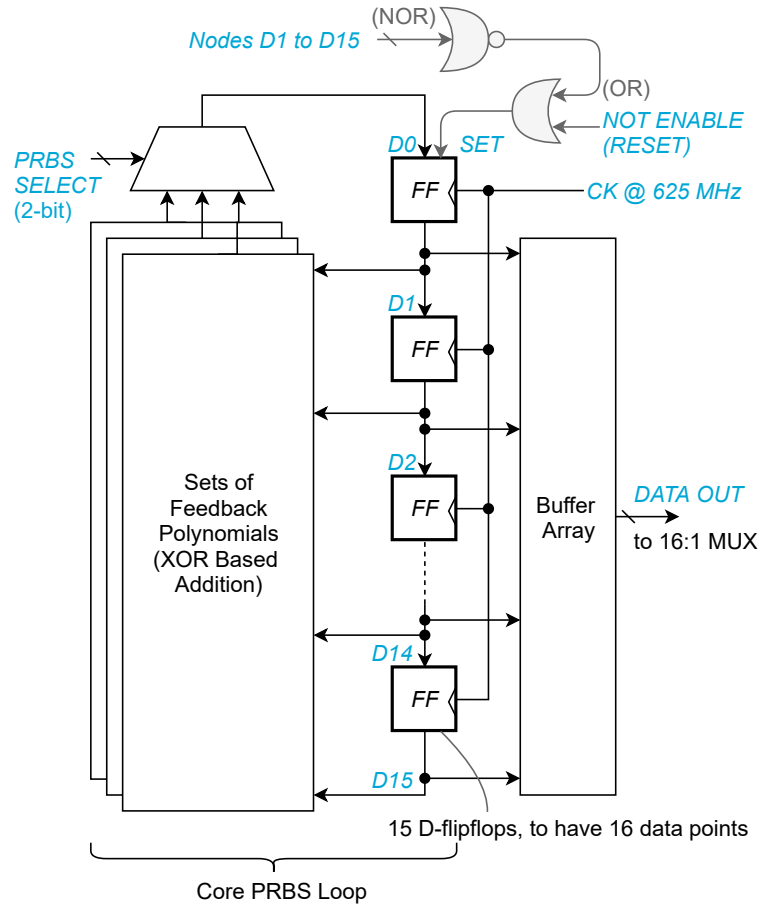


Figure 3.24: Overview of 16-bit PRBS generator in this work. The XOR feedback determines the periodicity of the sequence while the NOR feedback resets the loop if all D1 – 15 become zero in an unfortunate case (this would halt generation of new pseudo random bits).

Clocks Clock generation is critical in this wireline transmitter chip as the proposed tree multiplexer relies on the existence of four complementary, aligned clocks. Namely, the rising edges should arrive aligned at the multiplexer in order to be able to predict the order in which the PRBS data is serialized.

Shown in Fig. 3.25 is an overview of the designed clock generation circuitry. Our main design effort here is in the “aligned divided clock generator” where a special interclock aligner pre-aligns the divided clocks before a final retiming step. Additionally, we implement a variable delay element and clock inversion element to allow for possible transmitter operation at other clock frequencies than 10 GHz during measurements. These “delay trimmers” essentially extend the timing criteria of the clocked digital blocks.

3.5 Physical Implementation and Performance Verification

The overall chip overview shown at the start of this chapter (Fig. 3.1) should now be understood. The proposed wireline transmitter has been physically implemented for fabrication in a 6-week timespan, as a tape-out opportunity suddenly arose. In this section we highlight layout considerations and present core simulations results that verified post-layout functionality and performance.

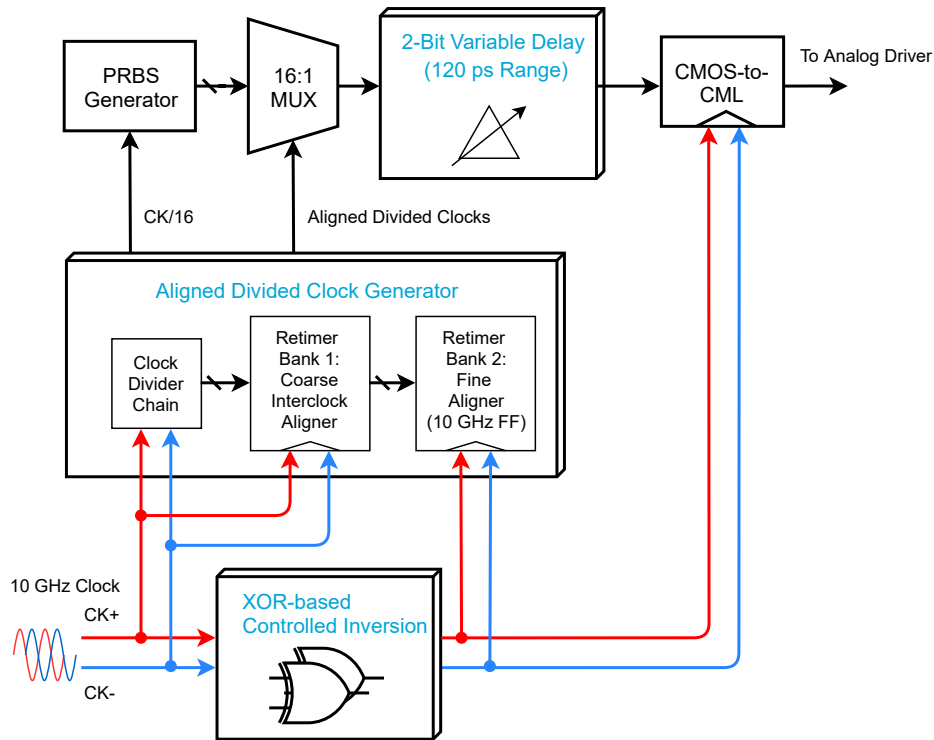


Figure 3.25: Overview of clock generation and delay trimmers in this work, highlighted in cyan. The 10 GHz complementary clock is obtained after passing an external 10 GHz reference through an on-chip transformer. The transformer design is kindly provided by PhD student Mohammadreza Beikmirza.

3.5.1 Chip Layout

Shown in Fig. 3.26 is the complete layout of the wireline transmitter chip. Including pads, the layout occupies 1.1 mm^2 . The core circuit area is equal to 0.093 mm^2 , to which the clock generation circuits contribute 0.023 mm^2 , the digital circuits 0.005 mm^2 and the CML drivers 0.065 mm^2 .

For the sake of brevity, we omit an extensive description of the layout implementation in this main text. For the interested reader, we provide more details on the pin planing in Appendix C. We only highlight the most important layout considerations below.

Supply Isolation The three supplies (analog, digital and clock) are separated using the NTN layer and double guard rings. The NTN layer is an extra resistive, undoped “substrate” layer. The NTN layer width of the “rotated T” separating the domains is roughly 30 um where possible. For hierarchy in layout, the core also has a small NTN ring around it. The width of this ring is not critical as it does not realize necessary isolation and the area between core and pads is filled with decoupling capacitors. The outer ring of NTN around top level is extra wide to provide necessary isolation.

Clock Distribution The clock generation circuits and clock distribution networks are laid out carefully. The placement of the subcircuits ensures that there is no need of clock line overlaps. All complementary clocks with equal frequency are closely routed together to minimize coupling to other lines. From the clock generation circuit to the multiplexer input, equal length routing is applied to uphold the timing budget.

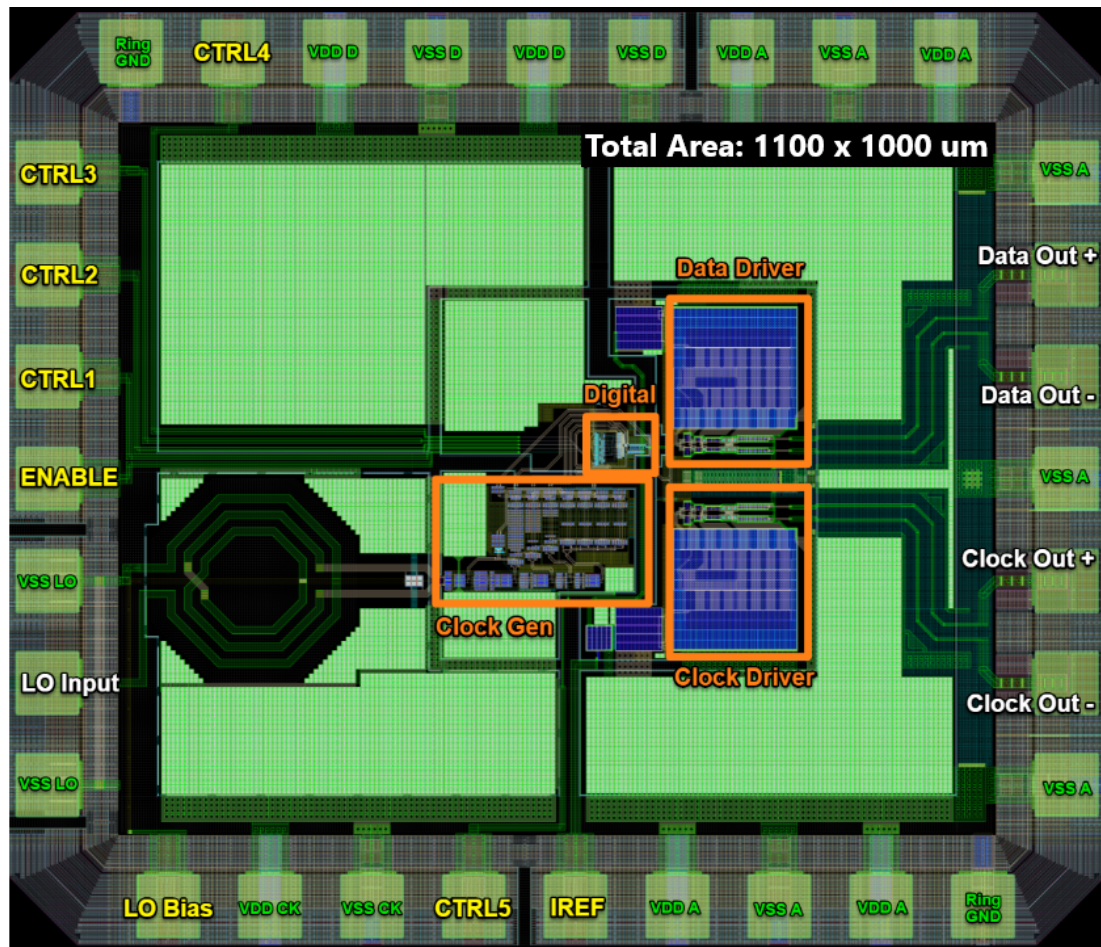


Figure 3.26: Full 10 Gbps wireline transmitter IC layout.

Matching The following choices were made in the process of layout to enhance matching:

- For the big driver tail current matrix, all tails are together in one rectangle, with a ring of dummy tails all around to improve tail-to-tail matching. For the smaller current mirrors there is no need for an all-around matching ring since we only make two copies. Here, dummy transistors provide equal environments above and below.
- The terminating $50\ \Omega$ CML load resistances are layed out by means of 13 parallel units each, in favor of reduced mismatch.
- For current reference distribution we chose to go for current-based copying instead of connecting gates between mirror cells across the chip so that we do not suffer from IR drops or noise in the ground return paths.
- The difference in delay of CML clock and data drivers are kept minimal by means of symmetry and maintaining an equal environment.

Driver Output Lines The CML output lines transition into $100\ \Omega$ differential edge-coupled coplanar waveguides to minimize reflections from the interface towards the PCB. As simulated $3\ \mu\text{m}$ width, $13\ \mu\text{m}$ spacing and $12\ \mu\text{m}$ clearance achieves a characteristic impedance of $104\ \Omega$ at $1\ \text{GHz}$, dropping of to $94\ \Omega$ at $10\ \text{GHz}$. The line is routed in ultra thick metal and presents less than $1\ \Omega$ of series wire resistance.

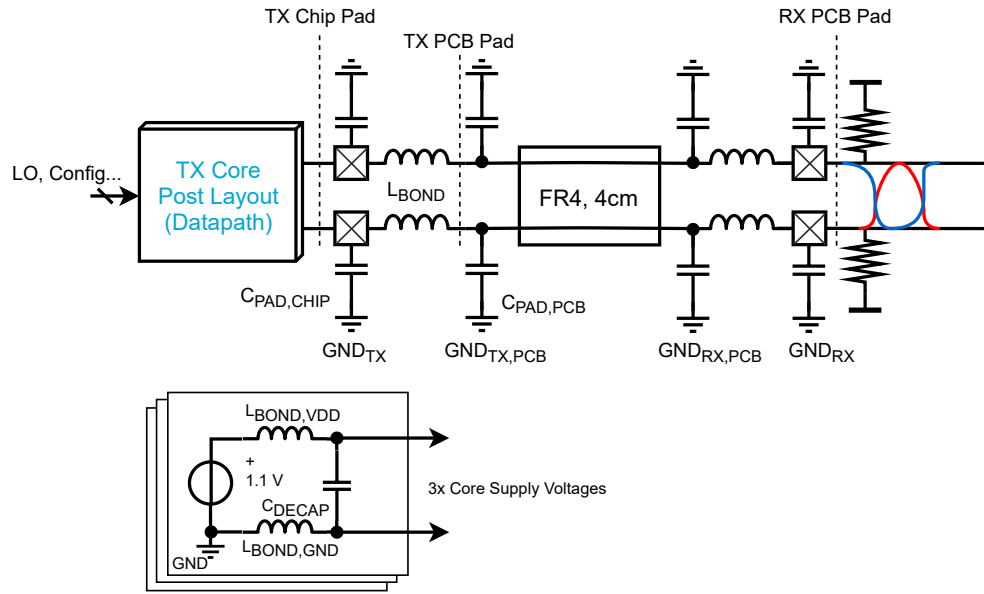


Figure 3.27: Wireline transmitter core simulation testbench, here shown only for one of the two differential outputs. The PCB and RX grounds are also connected to each other by means of parasitic inductances, but not shown.

Table 3.3: Summary of post-layout simulated wireline transmitter performance at 10 Gbps, measured over 4000 symbols.

Metric	Performance (10 Gbps, 60 °C)
Single Lane Driver Power (Efficiency)	9.9 mW (0.99 mW/Gbps)
Digital Power (Efficiency)	0.24 mW (0.024 mW/Gbps)
Clock Gen. Power (Efficiency)	20.5 mW (2.05 mW/Gbps)
Eye Height @ TX PCB Pad	>240 mV peak-to-peak
Output Risetime @ TX PCB Pad	42 ps on average to reach worst case eye height
Core Area	0.093 mm ²
Noise and Offset	1 mV _{rms} , 12 mV _{1σ}
Features	On-chip 16-bit PRBS, Forwarded Clock, Delay Trimmers

3.5.2 Core Simulations

Simulations of the core physical layout (complete layout excluding pads, decoupling capacitors and input transformer) confirm functionality and good performance.

Testbench In the corresponding testbench shown in Fig. 3.27, the transformer is modeled with an S-parameter model and its output is passed to the post-layout simulation model. 200 fF parasitic ESD capacitances model the high-speed output pads, while 100 fF is assumed for the PCB landing pad capacitances.

With respect to parasitic inductances, we model bondwires for supply lines (1 nH each), bondwires for high-speed pads (2 nH each and $k = 0.5$ mutual coupling) at TX and RX, and inductance between grounds of supplies on the PCB (0.5 nH each). These inductances are chosen pessimistic on purpose. Lastly the FR-4 PCB channel is modeled using rfTline library (by Cadence). We assume a realistic 4 cm long channel for all simulations here.

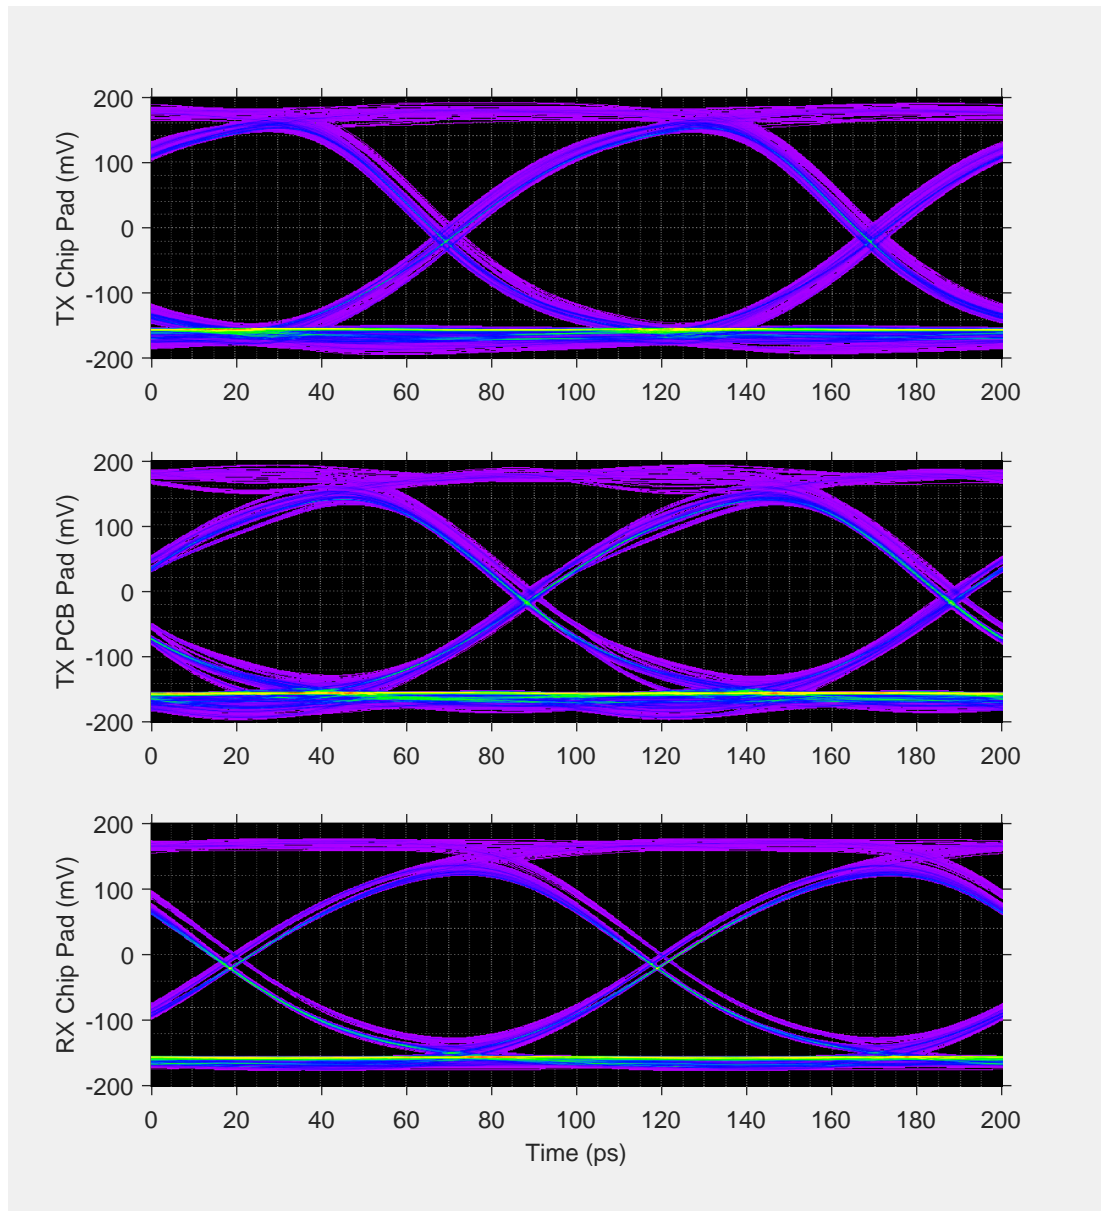


Figure 3.28: Eye diagram simulation results at various locations in the testbench.

Simulated Performance Table 3.3 summarizes the simulated performance. A transient simulation of 400 ns at 10 Gbps results in eye diagrams collected over 4000 symbols as shown in Fig. 3.28. We measure the eye at various locations in the testbench as marked in Fig. 3.27. We raise the temperature of operation to 60 °C for a more realistic die operating temperature.

The two CML drivers draw 18.7 mW in total, resulting in a 0.99 mW/Gbps driver efficiency. As expected from design, the digital circuitry draws negligible power, so the data path on its own will consume about 1 mW/Gbps. The clock generation circuitry has an efficiency of 2.05 mW/Gbps. The overall chip power efficiency is well below our target of 25 mW/Gbps. Simulations confirm that there is no significant ringing on the decoupled internal supply voltage despite the pessimistic bondwire parasitics in the testbench.

The eye height at the transmitter's PCB side is 300 mV peak-to-peak at average, with a worst-case of 240 mV due to ISI. The rise and fall times are about 42 ps, so we expect extra margin for

speed of operation.

The on-chip generated PRBs data sequences match the MATLAB predictions. The output forwarded clock and delay trimmers have been successfully been tested for functionality.

In a separate noise simulation¹⁰, we find that the driver generates 1 mV_{rms} of differential noise. The output referred offset has standard deviation of 12 mV¹¹ differential. We stay well within the SNR budget for BER < 10⁻⁸.

3.6 Initial Measurement Results

As mentioned in the introduction of this chapter, we will briefly discuss initial, promising measurements of the fabricated chip. The transmitter silicon returned late into the project when the receiver design had priority.

Measurement Setup Shown in Fig. 3.29 are the fabricated wireline transmitter die and PCB assembly. We provide more detailed documentation of the PCBs in Appendix D, for the interested reader. The used measurement setup is based around the two stacked boards as shown in Fig. 3.30, requiring an LO source, DC power supply and high-speed oscilloscope.

The used setup has two limitations. As a consequence of other higher-priority measurements in our group, these initial wireline transmitter measurements are carried out with a VNA acting as LO source (Keysight Streamline series). This LO source has a limited output power of ~ 7 dBm, thereby prohibiting successful on-chip clock generation faster than 8 GHz.

Secondly, the oscilloscope is DC-coupled when used in 50 Ω measurement mode, pulling down the transmitter output common-mode level by at least 500 mV. To prevent this overlooked aspect, DC-blocks are placed at the input of the oscilloscope. The wireline transmitter common-mode level is much better conditioned, but still too low by roughly 100 mV. In a future test setup, the receiving pull-up resistors should be placed on the measurement PCB while probing the channel with a high-speed probe.

First Results We are able to report that the chip is functional, indicating a successful tape-out. With the limited external LO power, the wireline transmitter operates up to 8 Gbps. The forwarded clock output continues to work up to 9 GHz. Fig. 3.31a and Fig. 3.31b show the corresponding signals measured. At higher frequencies, both signals vanish. The output data vanishes earlier because it requires the complete divider chain to work, which loads the (weak) on-chip LO.

Table 3.4 summarizes the first results in more detail. The nominal differential swing at maximum reference current¹² is over 250 mV peak-to-peak. It was observed that the DC-block and measurement cable alone already amount to just over 3 dB loss at 10 GHz. Therefore, we believe that the nominal swing is better in a realistic chip-to-chip case.

As it stands, the measured noise performance (after subtraction of oscilloscope noise) results in an SNR of over 26 dB, much better than the required 15 dB SNR for BER < 10⁻⁸. Still, it should be investigated why the noise performance is slightly worse than expected, possibly calling for a small increase of noise margin in the receiver design of future projects renditions.

The rise and fall time to reach the measured worst case eye height (~ 65 ps) are 35% more than simulated. The mentioned lower output common-mode level and additional channel losses contribute to this reduction of speed. Additionally, the parasitic capacitance involved in bonding (on

¹⁰Both AC noise and PSS noise simulations match.

¹¹We measure this by means of a transient Monte Carlo simulation and measuring the variation of output voltage level at a fixed moment in time.

¹²Maximum reference current is when the voltage on the reference pad is close to 1.3 V. Any higher applied voltage might damage the ESD diodes.




Figure 3.29: Fabricated die and PCB assembly. Fig. 3.26 shows the corresponding layout in greater detail.

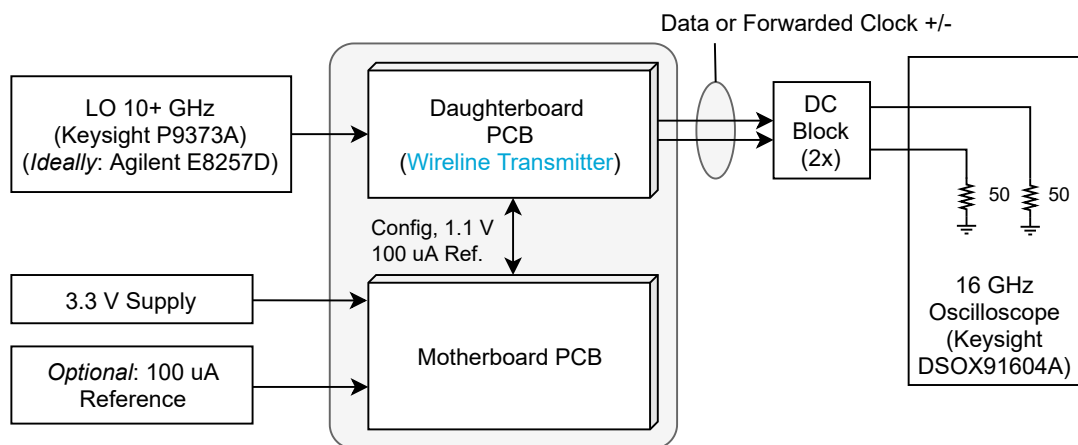


Figure 3.30: Overview of used measurement setup.




Figure 3.31: Initial wireline transmitter measurements with limited LO power. The signal quality starts to degrade significantly at 8 Gbps and 9 GHz for data and forwarded clock, respectively. (a) Highest speed eye diagrams, (b) highest speed forwarded clocks.

Table 3.4: Summary of first measurement results at 8 Gbps under the condition of limited available LO power.

Metric	Measured Performance (8 Gbps)
Nominal Output Swing	250 mV peak-to-peak
Noise (SNR)	6 mV _{rms} (26.3 dB)
Rise time to minimum eye height	65 ps
Analog Power Efficiency	0.92 mW/Gbps/lane
Clock Power Efficiency	1.57 mW/Gbps/lane
Digital Power Efficiency	0.024 mW/Gbps/lane
Total Power Efficiency	2.51 mW/Gbps/lane

the ESD pad) is about 100 fF higher than simulated for, according to our technician. Nonetheless, this speed reduction will not prohibit an open eye diagram at the targeted 10 Gbps.

Finally, the measured power efficiency is in the expected range. The output driver power consumption is slightly lower than expected, at 0.92 mW/Gbps per lane and is near constant over frequency. The complete chip is projected to consume only 2.51 mW/Gbps per lane when driven at 10 Gbps.

3.7 Chapter Conclusion

We presented the design and implementation of a complete, testable 10 Gbps wireline transmitter IC in 40 nm CMOS. The thorough design considerations, effective analyses and structural decision making are key in this achievement.

Based on an extended analysis of the C2MOS MUX, we were able to push a MUX tree to the speed limit of a digital system running on a single-phase clock. The proposed multiplexer can run up to 12.5 Gbps, much faster than a standard 3-NAND MUX tree could achieve in this process.

For the analog part of the wireline transmitter, we developed a new design procedure for CML driver chains. The method takes into account the many effects of voltage swing constraints and current steering efficiency. The proposed design procedure yields a four stage CML driver chain than successfully runs at 10 Gbps as seen from post-layout simulations. The driver is complemented by a novel skewless CMOS-to-CML bridge.

In addition to these core circuits, we designed a pseudo random bit stream generator and novel aligned divided clock generator, among others. These auxiliary circuits allowed for the tape-out of a complete, testable wireline transmitter chip.

After careful layout of the transmitter chip including the addition of a phase-matched forwarded clock path, post-layout simulations confirm successful operation at 10 Gbps with some room to spare while achieving a high SNR compared to the design target. The core circuits occupy less than 0.1 mm² and simulated overall power efficiency is 4 mW/Gbps to generate 10 Gbps data *and* forward a 10 GHz clock (two lanes).

We were able to perform initial measurements of the fabricated chip with customly designed motherboard and daughterboard wireline test PCBs. Due to the limited availability of LO sources, the chip has been characterized up to 8 Gbps. The measured eye diagram indicates that operation at 10 Gbps is certainly possible. The fabricated wireline transmitter consumes 2.51 mW/Gbps per lane, well within our set design targets.

4 Design of Comparators for Wireline

This chapter marks the beginning of the third part in this thesis. After spending a major effort to expose constraints in PAM wireline links and designing a complete wireline transmitter chip, it is sensible to also seek the design of the receiving end. Shown in Fig. 4.1 is early work at ELCA, where the design of a 1 Gbps wireline detector was attempted based on LVDS requirements. In the scope of this thesis, the vision is to push for a $10\times$ speed increase with simple circuits in the same 40 nm technology.

The comparator is the focus of this chapter and the most critical component in a wireline link as it takes the final decisions. The next chapter follows up by incorporation of the presented results here, to take on the design of a self-synchronizing wireline receiver.

We open this chapter with a wireline-focused examination of comparator design considerations. Among others, we develop an intuitive model for hysteresis effects in PAM receivers. Ensuing this is an instructive section reviewing the latest developments of comparators for analog-to-digital converters. Finally, we present the analysis and design of a two-tail comparator to match our 10 Gbps wireline transmitter. This work includes the critical simulation methods that enable fast, complete comparator design.

4.1 Considerations for Wireline Receivers

Fig. 4.2 shows a generalized model of a comparator with voltage-domain inputs and outputs. The primary goal of a comparator is to accurately extract the sign of a differential input voltage. The bit error rate (BER) measures the ability to do fulfill this function. Aside from this, discrete-time operation, pre-amplification, latch incorporation or performing domain conversion are all examples of optional design choices. Let us motivate and examine the correct approach for the third part of this thesis.

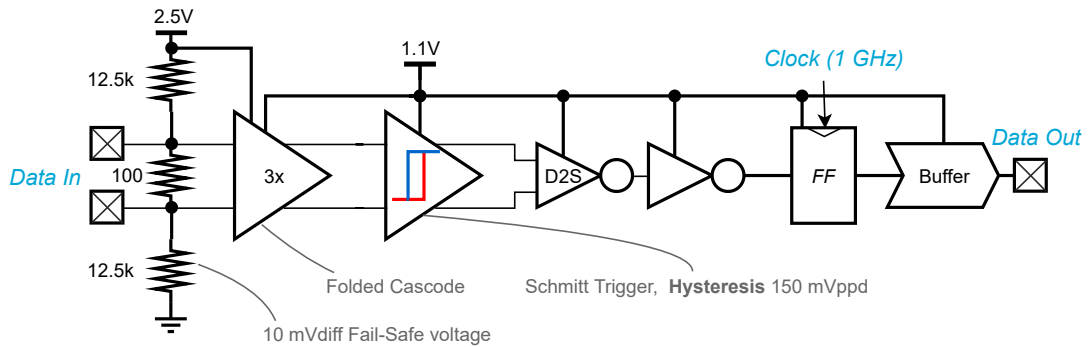


Figure 4.1: Overview of LVDS detector design effort from 2015 at ELCA. This work by PhD student Milad Mehrpoo in part motivates our intent to explore the design of a full wireline link.




Figure 4.2: Comparators exist with many variations of added functionality.

4.1.1 Clocked Operation

Comparators can be as simple as a single (limiting) amplifier operating in continuous time. Such comparators can be appropriate for slow single-shot ADCs, specific asynchronous conversions or other applications with a continuous feedback loop. Depending on application, local positive feedback in the form of a regenerative latch or hysteresis can lock-in the output decision if needed. In contrast to the continuous time applications, comparators operating with a *fixed* timebase are always desired in high-speed AD conversion, or rather expected for our case: wireline data detection.

Clocked comparators provide this discrete-time operation. Clocking increases complexity and clock driver power consumption, but allows for many choices that can lead to high-speed and/or low-power operation. Without exception, clocked comparators have a latching element controlled by a clock signal. The latch does not need to be regenerative: regeneration can provide a boost in operation speed or robustness of decision. In cases where this is not needed, it is also possible to decouple the decision and pass it to a flip-flop.

In ADCs it is common to have an additional clocked track-and-hold (T/H) operation preceding the latch. This makes sense: the comparator circuits need a finite amount of time to take a decision, while ADCs must resolve a random, continuously varying signal at fine resolutions. It is important to understand that for wireline applications, such T/H operation is not necessary. In a wireline link, the job of the receiver is to ensure that its detection process (i.e. the sampling clock) is in phase with the incoming data. Therefore, the *synchronized* comparator will not see much of a signal variation when enabled.

An optional preamplifier can enhance comparator performance. The amplification reduces the relative effect of internal noise and offset from the receiving end and provides isolation against unwanted feedback from the latch (kickback noise). Clocking a preamplifier further removes unnecessary power consumption. An ADC will surely incorporate a preamplifier to guarantee high-speed operation: the speed of a regenerative latch, and thus the overall ADC speed, drops exponentially for small input signals (and small internal signals for successive approximation ADCs).

Again, the design of comparators for wireline receivers is different in this respect. The expected input amplitudes under synchronous operation is constant due to the discrete nature of PAM signaling. Speed degradation due to metastability is not a great concern for wireline comparators, since under nominal conditions, the eye diagram is sampled at one of these discrete voltages. For an ADC, the sampled input voltage is fully random and can cause metastability under nominal conditions.

Lastly, clocked comparator usually have a reset phase where an additional circuit brings the comparator (e.g. amplifier output and latch) back into balance after making a decision. Thus, the speed limit of a resettable comparator is determined by both the comparison and reset times, whichever is greatest.

Thus, this work must aim for a clocked wireline comparator with the right combination of amplification and latch speed under the expected synchronous scenario. T/H is not necessary, metastability




Figure 4.3: Input-output relations for a comparator with hysteresis.

is of no concern and a reset phase can be used, but should not bottleneck the speed of operation.

4.1.2 Hysteresis

Resetting a comparator may not make immediate sense and should be motivated. A reset clears the previous decision results, removing any hysteresis (memory effect). In pure digital systems, it is a common understanding that hysteresis cleans up noise in data signals by insensitivity to signal bounce and chatter around the decision threshold as shown in Fig. 4.3. This hints at an improvement of BER in the presence of hysteresis.

On the other hand, for ADCs it is well known that a comparators hysteresis window limits the how fine its resolution can be. For example, for an ADC with two comparators having 10 mV of two-sided hysteresis cannot have corresponding reference voltages closer than 5 mV together. Though, it is reported that the special class of oversampling ADCs is rather insensitive to hysteresis or can even benefit from the effects.

Considering these contradicting aspects, we have developed a theoretical analysis in this wireline comparator work for general PAM signaling to show the under which circumstances the effects of hysteresis will be unrewarding, expressed in terms of BER.

PAM-2 Recall from Chapter 2 that the BER of PAM-2 is a function of the eye opening peak amplitude $V_{PK,EYE}$ to rms-noise $V_{N,rms}$ ratio, that is, the eye quality Q_{EYE} (Fig. 2.5). Moreover, For PAM-2 only, Q_{EYE} is equal to the SNR. Any bit error is caused by the probability of the gaussian noise tail (with normalize probability density function PDF_{gauss}) to spill over the decision threshold in the wrong direction. To review:

$$Q_{EYE,PAM-2} = \frac{V_{PK,EYE}}{V_{N,rms}} \quad (4.1)$$

$$BER_{PAM-2} = \int_{Q_{EYE}}^{\infty} PDF_{gauss}(x) dx \quad (4.2)$$

Fig. 4.4 shows this relation as a function of SNR. In addition, we plot an example result of a PAM-2 receiver MATLAB time-domain simulation with a 1 V input peak amplitude (differential) and 400 mV hysteresis around the 0 V decision threshold for a random data sequence. Clearly, the performance degrades due to an increase of the BER.

The key in finding the exact equation $BER_{PAM-2} = f(V_{HYST}, SNR)$ lies in recognizing that under hysteresis conditions, two cases can occur in the received data sequence:

- Case A: the comparator receives the *same* bit twice in a row $\Rightarrow Q_{EYE} = Q_{SAME}$.




Figure 4.4: PAM-2 BER as a function of SNR with non-zero hysteresis voltage from MATLAB time-domain simulation data (10^8 random samples per point). The data signal has a 1 V differential peak amplitude.

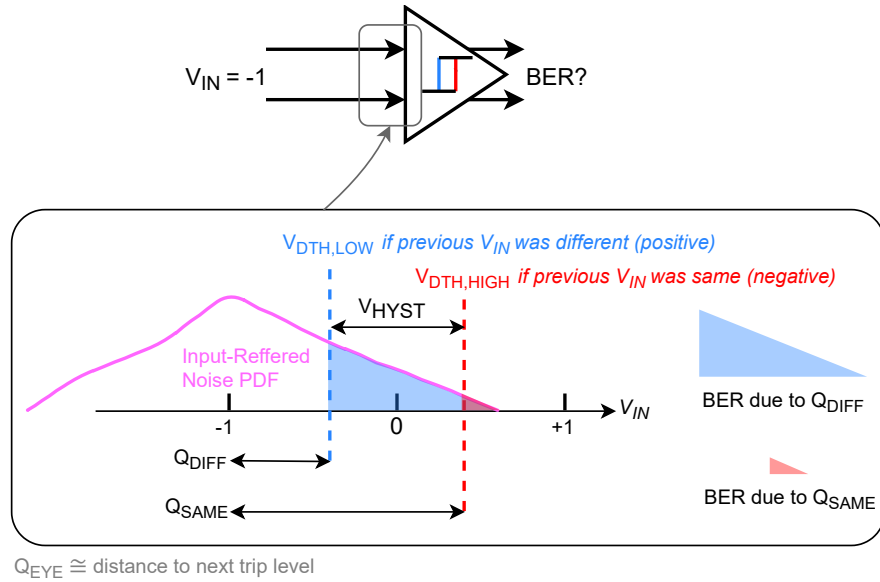


Figure 4.5: Probability of an error depends on the previous result. For example, if -1 is received it is favorable if the previous data bit was -1 too.

- Case B: the comparator receives a *different* bit compared to the previous one $\Rightarrow Q_{\text{EYE}} = Q_{\text{DIFF}}$.

In case A, the hysteresis will cause the decision threshold to have shifted to a beneficial position with respect to the “same” bit: the area of the noise PDF responsible for bit errors shrinks. In case B, the effect is opposite and the “different” has a higher change erroneous detection. The two cases are distinguished in Fig. 4.5.

To be exact, the eye quality increases for the “same”-bit case while it reduces for the “different”-bit case

$$Q_{\text{SAME}} = \frac{V_{\text{PK,EYE}} + 0.5V_{\text{HYST}}}{V_{\text{N,rms}}}, \quad Q_{\text{DIFF}} = \frac{V_{\text{PK,EYE}} - 0.5V_{\text{HYST}}}{V_{\text{N,rms}}} \quad (4.3)$$




Figure 4.6: PAM-2 BER as a function of SNR and hysteresis voltage, as predicted by our theoretical analysis, assuming fully random (fast) data and 1 V differential peak input amplitude.

So, the total BER amounts to

$$\text{BER}_{\text{PAM-2,HYST}} = P_{\text{SAME}} \int_{Q_{\text{SAME}}}^{\infty} \text{PDF}_{\text{gauss}}(x) dx + P_{\text{DIFF}} \int_{Q_{\text{DIFF}}}^{\infty} \text{PDF}_{\text{gauss}}(x) dx \quad (4.4)$$

where P_{SAME} and P_{DIFF} are the probabilities of case A and case B respectively. Of course, the sum of the two probabilities is equal to 1. P_{SAME} is responsible for an improvement of the BER while P_{DIFF} is responsible for a degradation of the BER.

For a purely random signal, P_{SAME} and P_{DIFF} are equal to 0.5. Plotting this case and verifying these equations with short-run time-domain simulations results in the trends seen in Fig. 4.6. Across the full SNR range, the BER is better when there is less hysteresis. In other words, the negative effect of Q_{DIFF} is stronger than the beneficial effect of Q_{SAME} . Fig. 4.7 shows intuitively why this is occurs: with Q_{DIFF} , the gain of responsible integrated noise PDF area is greater than the loss of the responsible area when Q_{SAME} is in action.

Thus, for high-speed (flash) ADCs, hysteresis will start to impact the BER before reaching the resolution limit. Similarly, the high-speed wireline comparator should have a reset phase to ensure (almost) zero hysteresis.

A subtle observation is that in the presence of higher SNR (low BER), the same hysteresis voltage has a stronger effect on the resulting BER. For a high SNR, the nominal Q_{EYE} is very far away from the decision threshold. The shape of a gaussian PDF makes that the loss of the responsible area when Q_{SAME} is very small. So, the effect of Q_{DIFF} area gain is much more amplified. The area under the tail of a gaussian PDF is increasingly sensitive to changes in detection threshold for lower nominal probabilities.

For a slowly varying input signal (e.g. data at 1 Gbps oversampled at 10 Gbps) or for data that is encoded such that there is a large probability of long sequences of ones or zeros, we have $P_{\text{SAME}} \gg P_{\text{DIFF}}$. It is consequently beneficial for the BER for these cases to have hysteresis for low SNR values (up to the point that Q_{DIFF} 's effect starts dominating again). An example is shown in Fig. 4.8 in the case where $P_{\text{SAME}} = 0.9$.

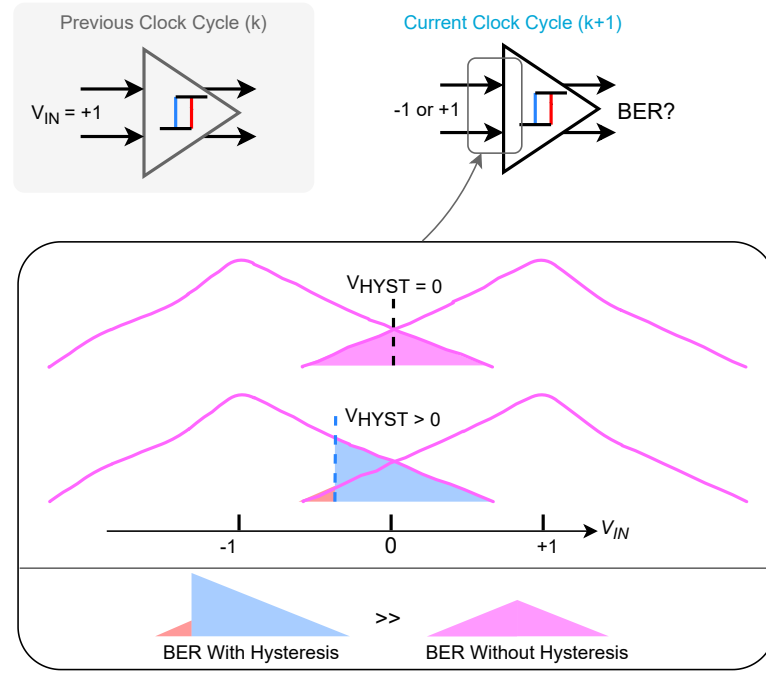


Figure 4.7: Total BER increases with hysteresis when P_{SAME} and P_{DIFF} are 0.5.




Figure 4.8: PAM-2 BER as a function of SNR and hysteresis voltage, as predicted by our theoretical analysis, assuming slow data (10× oversampled, i.e., $P_{\text{DIFF}} = 0.1$) and 1 V differential peak input amplitude.

Concluding Remarks on Hysteresis In the above analysis for PAM-2, the three main insights that contribute to our understanding are

1. The work in this thesis focuses equal-rate detection of random PAM-2 signals. We should minimize hysteresis to maximize BER.
2. To compute the effects of a varying detection threshold (hysteresis) on BER, it is required to know the original SNR and peak eye opening $V_{\text{PK,EYE}}$ to be able to evaluate the new Q_{EYE} .
3. At higher original Q_{EYE} , the relative impact of hysteresis on BER worsens.

As part of the work contributed by this thesis work, we expanded the above analysis of BER degradation due to hysteresis for the general PAM-M case in the form of a slide set. Appendix E presents this analysis.

A case that does not fully apply to the above analysis is the comparator with hysteresis in digital circuits often taught in electronics: the Schmitt trigger. To clear up this last contradiction, it should be recognized that a Schmitt trigger is a time-continuous comparator with hysteresis which only finds it use in (mechanical) switch debouncing and oscillator design. For switch debouncing, the Schmitt trigger is a logical choice, because the human input is expected to be very slow compared to the internal reactionary circuit. Finally, oscillator design does not fall in the category of data detection: BER does not play the same role.

4.1.3 Input Common-mode and Offset

For NMOS differential pair based input stages, input-referred offset, noise and speed grow with input common-mode voltage [11]. For ADCs, the fully random input voltage is problematic in two ways. The most evident concern is that the offset will cause a static input referred error, which limit the overall ADC accuracy in terms of e.g. conversion linearity. Another problem is that the offset of the ADC (and noise + speed) itself will vary if the fully random ADC input voltage directly sets the input stage's common mode voltage.

In the case of wireline comparators, the input common-mode fortunately does not vary much, but is high when using current-mode drivers on the transmit side. Consequently, the offset could be excessively high for simple links, limiting achievable BER. To circumvent this, an AC-coupled channel can be utilized, though only under the conditions that the data can be scrambled and the receiver or channel can set the far-end common-mode voltage.

A reset mechanism removes hysteresis does not necessarily nullify these (input-referred) offsets. The reset mechanism only ensures a default start scenario. Thus, the effect of offset on BER should be well understood in order to be able to assess design choices.

Fortunately, the above analysis for hysteresis effects easily adapts for offset too. For fast PAM-2 wireline data detection, a given input-referred offset V_{OS} will incur a net increase of BER due to 50 % of the bits enjoying a favorable detection threshold that is relatively far away from the expected bit value by V_{OS} , while the other 50 % suffers from a harmful detection threshold that shifted closer by V_{OS} . The former case corresponds to the change in eye quality for hysteresis with Q_{SAME} while the latter case corresponds to Q_{DIFF} (Eq. 4.3).

The main difference in thinking is that unlike hysteresis, it is not possible to appoint a fixed voltage to the offset. The offset is a random variable, varying from comparator to comparator and must be characterized by its voltage standard variation σ_{OS} ¹, and mean value of 0 V for differential comparators. Moreover, for differential wireline comparators we only care about the absolute value of the offset of each sample.

Only a requirement for yield can complete the design target for σ_{OS} . Here the yield is defined as the probability that the absolute offset of a sample stays below $z\sigma_{OS}$, i.e. $P(|V_{OS,sample}| < z\sigma_{OS})$. For example a value² of $z = 1$ corresponds to a yield of 68 %, $z = 2$ corresponds to a yield of 95 %.

So, the yield requirement demands that only for samples with $|V_{OS,sample}| \geq z\sigma_{OS}$ the BER may

¹In some circuit design papers, the offset is characterized by an rms-value. Technically, this is allowed as the mean offset is 0 V. In our opinion however, this notation misleadingly hints at the feasibility of calculating the BER by adding the offset rms directly to the noise rms. The fundamental mechanisms are completely different as discussed.

²In statistics referred to as absolute z-score or excess standard deviations. We are expecting a large number of samples, therefore we can assume a normal distribution.

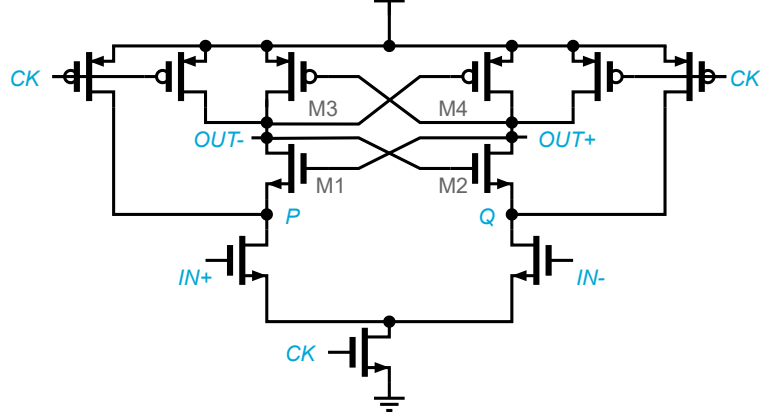


Figure 4.9: Classic StrongArm comparator.

exceed its design target. Consequently, the necessary BER calculation for offset starts with finding

$$Q_{OS,GOOD} = \frac{V_{pk} + z\sigma_{OS}}{V_{N,rms}}, Q_{OS,BAD} = \frac{V_{pk} - z\sigma_{OS}}{V_{N,rms}} \quad (4.5)$$

and the total offset degraded BER amounts to

$$BER_{PAM-2,OS} = 0.5 \int_{Q_{OS,GOOD}}^{\infty} PDF_{gauss}(x) dx + 0.5 \int_{Q_{OS,BAD}}^{\infty} PDF_{gauss}(x) dx \quad (4.6)$$

to find the design target for σ_{OS} . Also for this calculation, it is imperative to know the SNR and peak eye opening first: BER from a high initial SNR is affected more strongly by offset.

4.2 State-of-the-Art Comparators

The performance and design of comparators have mainly matured in view of advancements in the field of ADCs. Knowing the key considerations for a wireline comparators, it is now insightful from the perspective of this thesis' work to assess some of the most prominent comparator designs published for ADCs.

StrongArm Shown in Fig. 4.9 is the StrongArm comparator, considered to be a popular base design for clocked comparators. The StrongArm comparator operates as follows: a reset is facilitated by pulling up nodes $OUT+/-$, P and Q high when $CLK = 0$. Note that the tail is turned disabled. At the onset of $CLK = 1$, the tail turns on as the pull-up transistors turn off. A differential current proportional to the differential input voltage draws from P and Q while the output nodes are temporarily left floating at V_{DD} . $M1$ to $M4$ form a classic inverter-based latch and remains disabled as long as $M1$ or $M3$ do not conduct. As soon as P or Q drop by one threshold voltage, the latch turns on and makes a decision based on the pre-amplified voltage formed between P and Q .

The absence of an always-on tail for preamplification and of a path from V_{DD} to ground after decision making promotes energy efficiency. Moreover, the pre-amplification gained by sourcing current from the floating ("high impedance") P and Q nodes positions the StrongArm into the class of charge steering comparators. In the scope of this thesis, it is good to point out that the charge steering input stage manages high input common-mode voltages well without losing speed. P and Q start at V_{DD} , in contrast to equilibrium voltages of traditional differential pairs.

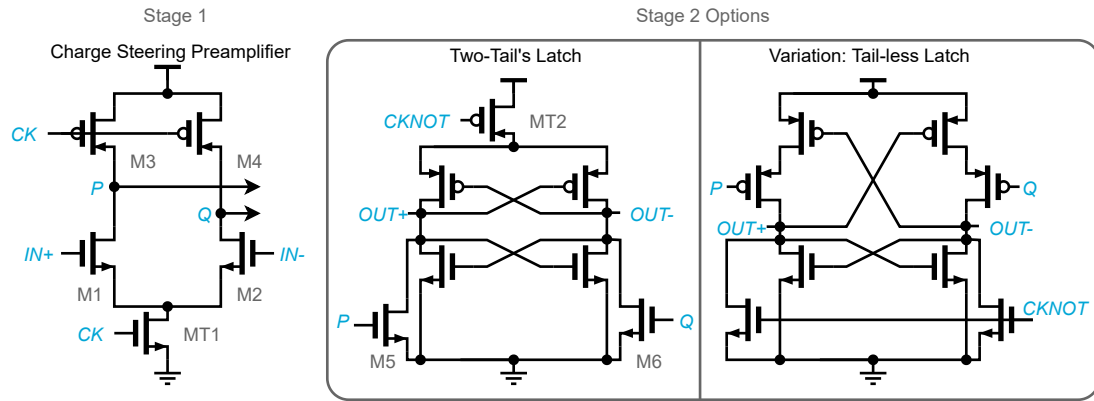


Figure 4.10: Two-tail comparator [12] and alternative design in [13].

Two-Stage The two-tail comparator in Fig. 4.10 is a transformation of the StrongArm comparator to create more design freedom [12]. The preamplification stage now has its own tail, and can be separately sized for its dominating effect on input-referred offset and noise without immediately dictating the bias condition and speed of the latch.

With two stages carrying out the precharge during reset, power-efficiency reduces. On the other hand, the two stage setup doubles the possibility of preamplification, as M5 and M6 also act as a charge steering pair. In an attempt to reduce the loss in power efficiency, [13] removes the tail in the second stage as shown on the right hand side: the main trade-off here is that M6 and M7 take some time to turn on completely resulting in loss of maximum operation speed. A follow-up on this two-stage single-tail approach was presented by [14] by modifying the first stage to prevent a full discharge of nodes P and Q, saving power. Speed of operation is further limited as the first tail's current drops over time.

Triple-Tail In successive approximation ADCs, metastability is a common problem due to the small voltages that might need to be discerned. The corresponding problem is that the conversion times of these ADC depends heavily on the input value. It is well known [15] that the speed of a latch is obtained as

$$t_{\text{latch}} = \tau_{\text{latch}} \ln \left(\frac{V_{\text{END}}}{V_{\text{INIT}}} \right) \quad (4.7)$$

where τ_{latch} is the latch's time constant³, and V_{END} and V_{INIT} are corresponding final and initial voltages at the output. The initial voltage is set by preamplifiers.

To tackle this issue, [16] presents a triple-tail comparator that relies on the availability of a lot of gain (three stages) to boost the speed of small input amplitude conversions as shown in Fig. 4.11. The same author presents an improved version in [17] that adds a feed-forward path that repairs the speed loss compared to two-tail comparators in the case of large input amplitude conversions. The triple-tail approach clearly increases power consumption, but shows well how high-speed operation is achieved. At small input amplitudes, multiple charge steering stages will benefit the speed of the latch.

³The time constant of a regenerative latch is equal to its load capacitance divided by its equivalent trans conductance, that is, $C_L/g_{m,\text{latch}}$ [15].

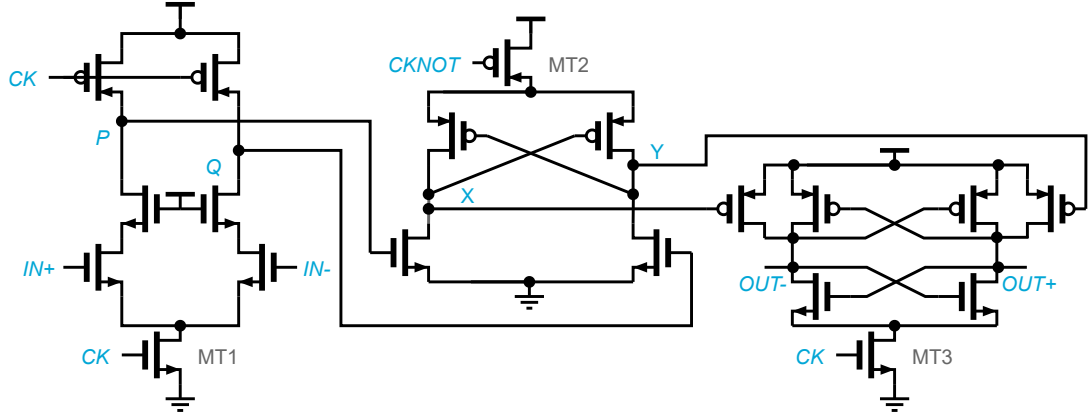


Figure 4.11: Triple-tail comparator by [16] boosting performance at small input voltages.

4.3 Two-tail Comparator Analysis and Design

For the wireline receiver work in this thesis, the basic two-tail comparator is an excellent candidate for clocked comparator data detection. Based on the preceding considerations and literature study, it is clear that input voltage-wise this work concerns a narrow working range. The two-tail comparator offers design flexibility for high speed and offset/noise tuning as opposed to the single-tail StrongArm comparator, without the extra overhead of a triple-tail setup.

4.3.1 Summary of Requirements

In this section, we design a two-tail comparator for a wireline receiver to complement the transmitter chip of the previous chapter. Post-layout simulations of the wireline transmitter suggested that a well synchronized receiver will sample the input data with a worst case (differential) peak amplitude of 120 mV.

To give the overall receiver some design margin and to allow for extra freedom in designing the complete synchronization solution (next chapter's topic), we target a minimum comparator input voltage amplitude specification of 60 mV_{PK}.

Moreover, under this condition the designed comparator should be at least able to operate at 10 Gbps while guaranteeing a BER of 10^{-8} at most for the total transceiver link, as specified in Section 2.3. Specifically, our proposed design choices for high-speed operation and a set of developed fast comparator simulations are the key enablers here for the presented synchronized receiver shown in the next chapter.

4.3.2 Analysis of Operation

The two-tail comparator exhibits more complexity than a StrongArm comparator. Shown in Fig. 4.12 is the full schematic.

Reset The reset starts when $CK = 0$, balances the comparator and is ready when both output nodes are pulled to ground. It prepares for the next compare phase, so it is good to know how the details of reset first. A two-step process takes place:

1. P and Q start fully discharged (MT1 was on in the previous compare phase). M3 and M4 now pull up P and Q to V_{DD} . The second stage is static until P and Q are high enough to enable M5




Figure 4.12: Two-tail comparator in this design and evolution of its main signals. A set-reset (SR) latch follows up to preserve the comparison result at $OUT+/-$ during the reset phase.




Figure 4.13: Two-tail timing details. (a) Comparison phase, (b) reset phase.

and M6. This first step requires a time duration t_{RESET1} of

$$t_{\text{RESET1}} = V_{T5,6} \frac{C_{PQ}}{I_{D3,4}} \quad (4.8)$$

where V_T is the threshold voltage of a transistor.

2. One output node is floating at V_{DD} and must be pulled down by M5 or M6, in turn driven by P or Q . The total reset time takes an additional time duration t_{RESET2} of

$$t_{\text{RESET2}} = V_{DD} \frac{C_{OUT+/-}}{I_{D5,6}} \quad (4.9)$$

The total reset time is the sum of t_{RESET1} and t_{RESET2} . The (dis)charge currents can be easily scaled for, as the M3,4 and M5 or M6 initially operate in saturation after enabling. Scaling is with limitations, due to contributions of their own parasitic drain-to-bulk capacitance. Fig. 4.13b illustrates the reset phase timing details.

Compare The comparison starts when $CK = 1$ also comes about in two steps:

1. P and Q start fully charged and floating, and $OUT+/-$ are both discharged. The latch formed by M7–10 does not enable its positive feedback action until either M7 or M8 has charged up by one $V_{T7,8}$. This pre-latch time t_{PRE} amounts to

$$t_{PRE} = V_{T7,8} \frac{C_{OUT+/-}}{0.5I_{T2}} \quad (4.10)$$

2. Meanwhile the differential pair draws a common-mode and differential current, creating differential gain by voltage integration ($Q_C = \int_0^t I_C(\tau) d\tau$) through pairs M1,2 and M5,6. At $t = t_{PRE}$ the latch makes a fast decision based on the built-up output voltage. This latch time t_{LATCH} amounts to

$$t_{LATCH} = \tau_{LATCH} \ln\left(\frac{V_{OUT,END}}{V_{OUT}(t = t_{LATCH})}\right) \quad (4.11)$$

where τ_{LATCH} is equal to $C_{OUT+/-}/g_{m,LATCH}$ and $V_{OUT,END}$ can be chosen as $V_{DD}/2$. This choice is justified as the comparator must be cascaded by an SR latch to preserve the decision while the reset ensues.

The latch speed is usually high since τ_{LATCH} approaches the transit frequency of the process. Thus, t_{PRE} dominates the comparison time and is most easily tuned by scaling of MT2. Fig. 4.13a illustrates the comparison phase timing details.

To preserve the comparison (latch) result as long as possible, another set-reset (SR) latch consisting of two NOR gates follows the two-tail comparator as shown in Fig. 4.12. A set-reset latch does not change its output Q when both inputs are 0 V. Otherwise Q follows $OUT+$.

Overall Timing To summarize for the time-domain relations, we now understand that

$$T_{CLKPERIOD}/2 > \max(t_{RESET}, t_{COMP}) \quad (4.12)$$

$$t_{RESET} = t_{RESET1} + t_{RESET1} \quad (4.13)$$

$$t_{COMP} = t_{PRE} + t_{LATCH} \quad (4.14)$$

with t_{RESET1} , t_{RESET1} , t_{PRE} and t_{LATCH} as identified in the previous paragraphs.

Gain, Offset and Noise The gain follows from two charge steering voltage integrations during t_{PRE} . First via integration of current generated by the product of input voltage and transconductance $g_{m1,2}$ of pair M1,2 onto C_{PQ} , accumulating charge Q_{PQ} . This charge results in an equivalent voltage $V_{PQ} = Q_{PQ}/C_{PQ}$. Simultaneously, this integration continues to V_{OUT} via the current generated by M5,6.

That is, the gain A_V grows over time up to t_{PRE} as

$$A_V(t) = \frac{V_{OUT}}{V_{IN}} = \frac{1}{C_{OUT}} \int_{t_2}^t g_{m5,6} \left(\frac{1}{C_{PQ}} \int_0^t g_{m1,2} dx \right) dy \quad (4.15)$$

such that

$$A_V = \left[\frac{V_{OUT}}{V_{IN}} \right]_{t=t_{PRE}} = \frac{g_{m1,2}g_{m5,6}}{C_{PQ}C_{OUT}} (t_{PRE}^2 - t_2^2) \quad (4.16)$$

where t_2 is the delay until M5,6 start to enter saturation (to draw significant current). The voltage gain clearly⁴ grows quadratically with t_{PRE} .

⁴Certainly, the transconductances in the equations are not exactly constant, but we should realize that M1,2 operates in saturation and M5,6's initial gain in linear mode can be disregarded via t_2 . So we can assume saturation and easily perform sizing by applying a g_m/I_D procedure to scale transconductance predictably. The g_m/I_D sizing procedure [18] is an empirical design method based on pre-characterization of transistor transconductance-to-current performance versus inversion level.




Figure 4.14: Simulated two-tail comparator characteristics versus differential input peak to peak swing, confirming our understanding.

Mismatch contributions of the first stage differential pair dominate the overall input-referred offset. With a fixed common-mode input voltage set by the transmit side, we must size up M1,2 and MT1 to reduce the input-referred offset without impacting of charge steering gain ($g_{m1,2}/C_{PQ}$ depends on inversion level). If this gain by M1,2 reduces, the contributions by the latch and M5,6 referred to the input will increase. Increasing C_{PQ} will additionally cause t_{RESET1} to raise, reducing speed of operation.

The noise calculation of a two-tail comparator is not straight-forward and an intuitive explanation is fitting here to reach the desired insight. The pairs M1,2 and M5,6 cascade each other while creating gain, so the noise contributed by M1,2 will also dominate the input-referred noise. The noise current from M1,2 is integrated over C_{PQ} , resulting in a linear growth of its variance (i.e. mean noise power) over time. M5,6 have the same effect while integrating the resultant noise voltage once more. Thus, the output noise power grows quadratically⁵ with t_{PRE} . Consequently, the output rms noise voltage must grow linearly with t_{PRE} . Dividing by the voltage gain shows that the input-referred rms noise voltage is inversely proportional to t_{PRE} .

To conclude, this leads to the belief that for a given sizing arrangement, the two-tail comparator's input-referred noise is almost constant over input amplitude, exactly like the behavior of t_{PRE} . The input-referred SNR therefore drops predictably in a linear fashion with respect to input amplitude, even if the comparator starts to slow down.

Simulations during design confirm our preceding analysis as shown in Fig. 4.14.

⁵To be exact, the delay t_2 until M5,6 start to act results in a small additive term. This term is insignificant for the purpose of this explanation.

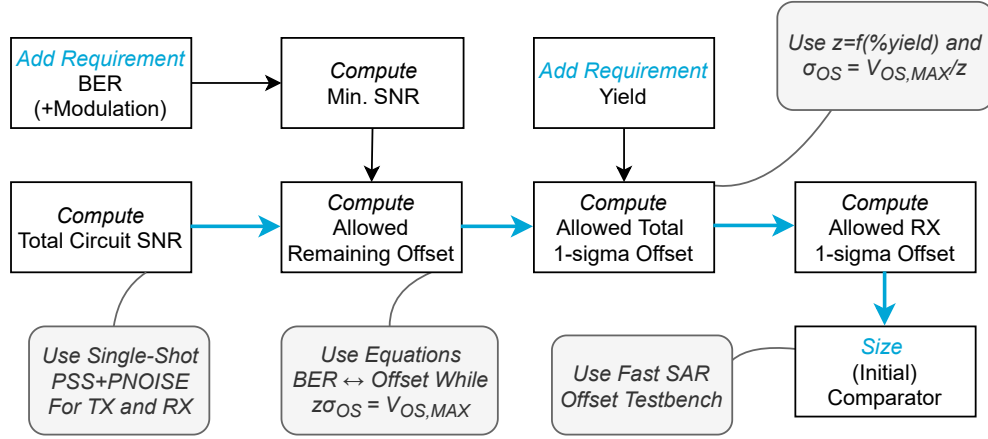


Figure 4.15: Two-tail design procedure developed during this work. The simulation methods are explained in the next subsection.

4.3.3 Design Choices and Results

The thorough understanding created allows for our design choices that lead to a satisfactory 10 Gbps two-tail comparator.

10 Gbps Speed The speed of operation is limited by the greatest of $t_{\text{COMP}} = t_{\text{PRE}} + t_{\text{LATCH}}$ and $t_{\text{RESET}} = t_{\text{RESET1}} + t_{\text{RESET2}}$. For 10 Gbps operation, both must be less than 50 ps. In this design, we exploit our process technology's freedom to specifically place high- or low- V_T transistors and break paradoxical sizing limitations.

t_{PRE} dominates the comparison time and depends on the charge speed of $C_{\text{OUT+/-}}$ by MT2. To push t_{PRE} to less than 50 ps, simply sizing up MT2 has its limits. MT2's increasing drain capacitance and clock feedthrough will start to hamper significant speed up. In this design, MT2 is implemented as a high- V_T PMOS transistor to overcome diminishing returns from sizing. The choice allows MT2 to stay operating in strong saturation while its drain voltage follows the latch output common-mode rise.

For the reset phase, t_{RESET1} depends on the delay until M5,6 start conducting. M3,4 cannot be sized much larger as soon as their drain capacitance start dominating C_{PQ} . Similarly, to reduce t_{RESET} M5,6 cannot be sized too large in order to keep contributions to C_{PQ} and $C_{\text{OUT+/-}}$ moderate. Thus, for this high-speed design, M5,6 are implemented as low- V_T NMOS transistors. Compared to normal- V_T devices with same sizing, t_{RESET1} reduces as M5,6 turn on earlier while t_{RESET2} reduces as a higher overdrive is enjoyed.

99% BER Yield Application of the above techniques allows for a speed faster than 10 Gbps with a 60 mV peak amplitude input signal. The circuit still must be scaled for sufficient BER ($< 10^{-8}$) with this worst case peak input amplitude.

The correct procedure at this point of design is shown in Fig. 4.15. Recall from our previously developed derivations in Section 4.1.3 that BER calculation based on offset requires knowledge of the SNR. Simulations confirm a constant input-referred noise level of 2 mV_{rms} for the current-size comparator, regardless of input amplitude. Including TX noise (referred to the comparator) for Eq. 4.6 shows that an absolute comparator input-referred total offset of 46.5 mV can be tolerated to meet our BER requirement.

Our fast offset simulations show that the comparator currently contributes $\sigma_{\text{OS,RX}} = 25$ mV. However, to obtain a yield of 99% ($\equiv 2.5\sigma$), the total offset $\sigma_{\text{OS,IN}}$ should be limited to $46.5/2.5 = 18.6$

Table 4.1: Summary of achieved comparator performance at 10 Gbps, 60 mV input peak amplitude and 1 V common-mode input voltage.

Metric	Pre-layout RF Transistor	Post-layout Normal Transistors
Power Efficiency	0.15 mW/Gbps	0.07 mW/Gbps
Comparison Time	40 ps	25 ps
Reset Time	35 ps	25 ps
Input Referred Noise	1 mV _{rms}	<i>Not Measured</i>
Offset	12 mV _{1σ}	<i>Not Measured</i>
Yield at BER 10 ⁻⁸	99%	<i>Not Measured</i>

mV. The TX will be signaling directly to the comparator, so the TX output-referred offset contributes directly to the comparator's input-referred offset. So, as presented in the last chapter, the TX will contribute $\sigma_{OS,TX} = 12$ mV to this 18.6 mV budget.

Subtracting the TX offset contribution⁶ reveals that the comparator should be sized up for an input-referred offset of at most 14 mV. A full size-up of 4 \times achieves this results.

Final Design Table 4.1 summarizes our achieved comparator design performance. We should explain first that we co-designed the receiver in the next chapter together with the comparator, at schematic level. For those pre-layout designs, we used RF transistor models to ensure that we would not be disappointed later on after layout.

The pre-layout design is significantly slower than the post-layout design, but safely runs at 10 Gbps ($t_{COMP} < 50$ ps). For the post-layout comparator, we used non-RF transistors (but same gate length and widths), which can be layed out very compact, thereby explaining the higher speed. We surmise that the RF transistor models overestimate the parasitics compared to our compact layout.

The pessimistic pre-layout RF two-tail comparator already passes all the requirements stated at the start of this design subsection. For both implementations, the power efficiency is excellent. Fig. 4.16 shows the layout of the faster two-tail comparator.

4.3.4 Fast Simulation Methods

A large challenge in this comparator design effort was the delivery of a satisfying design in a short time span. Comparators are time-varying and non-linear, requiring specific simulations. Traditional, statistical noise and offset simulation methods for comparators proved to be very slow, prohibiting completion of this design challenge.

As such, another significant contribution in this thesis work is our investigation for a set of fast *and* accurate simulation methods for comparator design that can be used in future iterations of this project. This successful investigation leads to simulation methods comprising a single-shot deterministic noise simulation and fast offset search algorithm.

We realize that the reader may not be familiar with Cadence, so we present these fast simulation methods in Appendix F instead for those interested.

4.4 Chapter Conclusion

This chapter constitutes the effort to gain understanding, background knowledge, design intuition and fast simulations for wireline comparator design.

⁶In the variance domain, i.e. $\sigma_{OS,RX} = \sqrt{\sigma_{OS}^2 - \sigma_{OS,TX}^2}$

We contributed to the goal of this thesis with an analysis of BER degradation due to hysteresis and offset. We found that this computation requires knowledge of the SNR and either absolute signal or noise level as derived. When sizing for offset, an additional yield requirement is necessary.

More importantly, we successfully designed and implemented a fast, low-power two-tail comparator for our application. The pre-layout design safely achieves a speed of 10 Gbps while ensuring that we can meet the set wireline link design targets in terms of bit error rate at 99% confidence. By virtue of a compact layout strategy, we were able to improve the potential speed of the post-layout two-tail comparator to almost 20 Gbps.

Finally, we recognized that comparator design can be sped up significantly by employing PNOISE simulations and SAR algorithms for noise and offset simulations, respectively. This developed set of simulation methods are of great use for future (wireline) comparator design efforts.




Figure 4.16: Compact, fast two-tail comparator layout. The data enters from below on M6 and leaves on the top. The clock signals are provided from the left. The occupied area amounts to 64 um^2 .

5 Self-Synchronizing Wireline Receiver

We already have the means to do high-speed data detection, but are missing the critical ability to perform this detection (sampling) at the right moment in time. After all, the previous comparator design was performed with a static input voltage or a clock conveniently provided by the simulation testbench, that is, assuming good clock to data alignment.

This thesis work ends with the development and design of a non-traditional wireline receiver as presented in this chapter. We start with an overview of the current challenge and corresponding common wireline receiver design solutions. However, the time remaining for the project is limited: to achieve autonomous synchronization nevertheless, we go ahead with the exciting design of a synchronization loop based on comparator metastability that does not require fine phase control of the clock. It results in a simple and elegant solution for this critical, missing function.

5.1 Overview and Key Challenges

As explained in the introductory chapter, a complete wireline receiver performs the following critical functions:

- conditioning or generation of a clock, to have a sampling clock
- data phase detection, to know the data's position in time with respect to the sampling clock
- data detection, to resolve the received information
- data deserialization (demultiplexing), to undo the serialization at the transmitter side

Clocked data detection is a key challenge, but has been covered in the preceding work so far. Moreover, the design of a demultiplexer (DEMUX) of 10 Gbps data should not be a challenge anymore: the task for the first 1:2 DEMUX is only to hold the 100 ps bits for one period longer (or more). The fast C2MOS flip-flop featured in the wireline transmitter chapter can definitely perform this task, down to the data rates at which standard library components take over. Generating the necessary divided clocks for such a DEMUX tree poses no difficult timing problems because the clock division and data demultiplexing occur in the same direction of propagation.

5.1.1 Considerations for Sampling Clock Alignment

The remaining key challenge is to choose and implement an architecture to close the phase detection and sampling clock conditioning loop. As Fig. 5.1 shows, many options are possible. The clock needs to be sourced from somewhere and phase-corrected based on feedback after data phase comparison.

If transmission of a clock is not feasible or the number of pins is highly restricted, clock recovery is a necessity. As such, a clock should be generated on-chip with the same frequency as the data using a clock and data recovery loop (CDR). When a clock is forwarded, it may still be a good idea to filter the clock jitter in the phase domain by duplicating it with a 1:1 phase-locked loop (PLL)

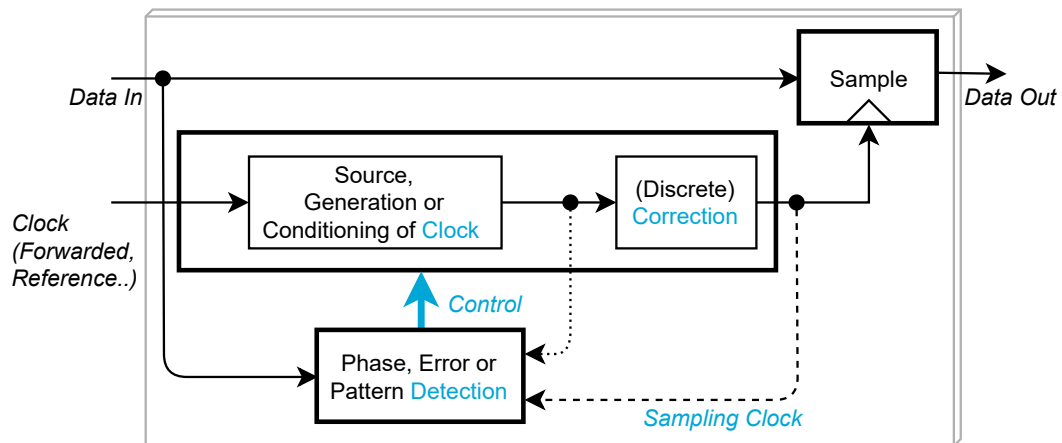


Figure 5.1: Options in sampling clock alignment architectures.

before aligning to the data. Especially for tough eye diagrams, the BER can be very sensitive to clock jitter. A CDR can provide a direct means to generate multiple clocks that individually align to multiple input data streams.

If a CDR is available, the phase correction can be applied inside the loop of clock generation. On the other hand, if the clock is readily available or when multiple clocks should be aligned to multiple data streams, it is more logical to implement phase correction using a controlled delay line. Power-wise it may not be economical to implement a large phase correction range with fine resolution delay steps. Therefore, a delay-locked loop may be used to enable a coarse-fine correction arrangement.

As discussed later, data phase detection is actually not trivial: data is not guaranteed to be periodic, requiring a special phase detector that can discern the phase from the available transitions in the random data for the CDR. As Fig. 5.1 illustrates: for correction, the extracted phase from the input data can be compared to the nominal sampling clock (dotted) or the corrected sampling clock (dashed).

In avoiding a special phase detector, it may be possible to precalibrate the correction settings of the receiver by sending a periodic pilot signal to the receiver's data terminal. Though, this requires external control of the calibration and normal states of the receiver, complicating the whole system operation. Moreover, this method might need to be coupled with a fine background feedback loop under normal conditions in the view of, e.g., temperature variations.

To fully avoid the need of phase detection, external control must be fully utilized. By addition of an additional human or machine in the loop, the sampling phase could be adjusted based on the reception of a predefined data pattern. Of course, this method is cumbersome and not as accurate. It should only be applied if the system operation can be disrupted.

5.2 Synchronization by Metastability

It is our wish to design the core of a wireline receiver that can complete the link from the presented transmitter chip in Chapter 3: a data rate 10 Gbps at BER < 10^{-8} for the full link.

The receiver conceptual architecture, based on metastability, is shown in Fig. 5.2 and explained further in this section. Given the short time available, we omit the design of the demultiplexer and CML-to-CMOS clock buffers. Instead, we focus on sampling clock alignment. In the vision of this project, this is the remaining prime challenge that needs progress.

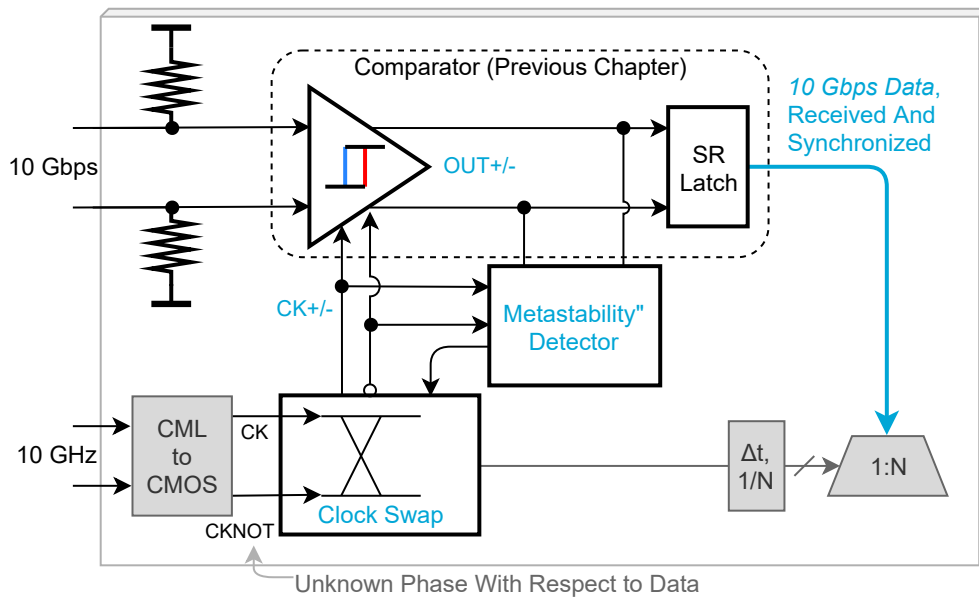


Figure 5.2: Overview of the target design in this chapter. A self-synchronizing receiver will enable multi-lane links with a single global clock.

5.2.1 Motivation of Architecture

Based on the general considerations in the previous section and limited design time, we made the following additional observations and choices for the target design:

- The design should not feature a CDR loop to generate the clock or filter the clock. The presented transmitter design sends data and clock, with only 1 mV_{rms} of noise, so the jitter on the forwarded clock should be acceptable in conjunction with the comparator performance. This makes the design a *synchronizer* as opposed to a CDR circuit.
- The synchronizer should preferably leverage the comparator design of the previous chapter to assist development speed.
- The phase detection and correction should be simple and built around the comparator since there will be no oscillator loop. A non-traditional architecture may be used to exploit the comparator's great performance.
- The synchronizer should operate autonomously (no externally initiated calibration or pattern detection) and continuously correct the sampling clock.

Edge Swap Phase Correction The idea embedded in the target design (Fig. 5.2) can be explained as follows. The forwarded clock inherently has two edges (rising and falling) per data period. Assuming a favorable random clock phase offset with respect to the data, the detector will have the rising edge in the middle of the eye: the falling edge should not be used to sample data. Of course, this phase relation cannot be guaranteed. Another case is when the edges are at 90° and 270° away from the eye crossing: in this case an ideal sampler would see the same input voltage regardless of the edge used to sample as shown in Fig. 5.3. For all other cases, the sampled voltage can be higher if the best edge is picked for sampling.

If a comparator has overhead in designed performance (SNR or BER), we can surmise that there is always one clock edge that can be used to sample the data with satisfactory BER. In effect, the




Figure 5.3: An equal rate clock has two edges that can be used for sampling. Depending on the clock to data skew, one of the edges has a superior SNR (eye height). Therefore the $90^\circ/270^\circ$ situation is a worst case if there is possibility to use either of the two edges.

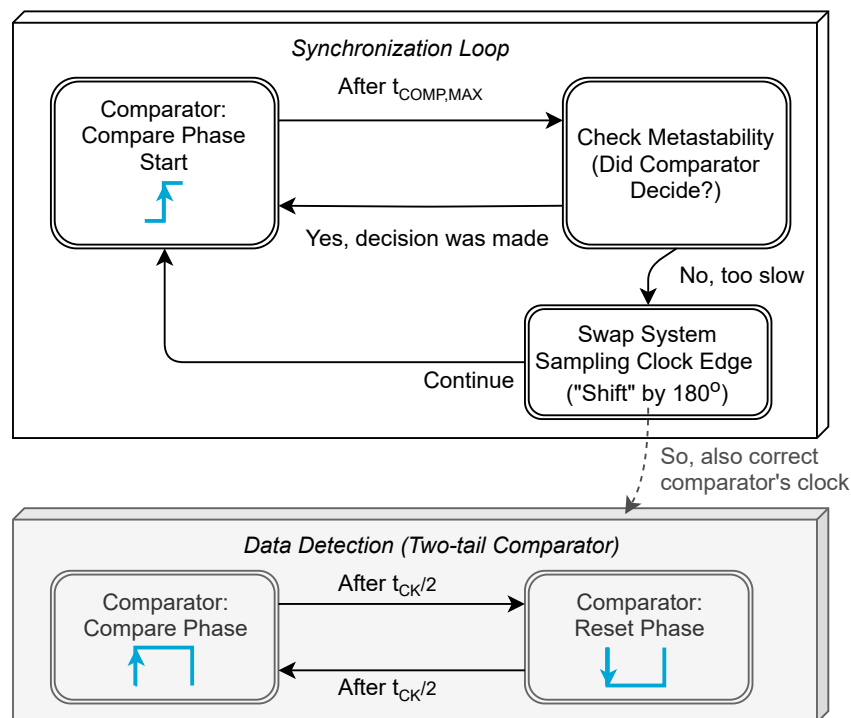


Figure 5.4: Initial state diagram of proposed receiver with synchronization based on metastability detection and correction by swapping the sampling edge.

clock phase correction can be as coarse as possible. To build a full synchronizer from this idea, several possibilities are available.

- A double-edge triggered comparator could sample twice per data period. From its 20 Gbps output stream, the correct output of the 1:2 DEMUX (determined by some phase detection feedback) could be picked and further deserialized.
- It could be possible to place two comparators in parallel, with one of the two driven by an inverted version of the forwarded clock. The data from the optimum clock edge is then further deserialized.

While this interleaving approach is feasible, this work proposes a clock swapper circuit instead, in order to save on power and area for the clock buffers. Fig. 5.4 overviews our proposed architecture in the form of a state diagram.

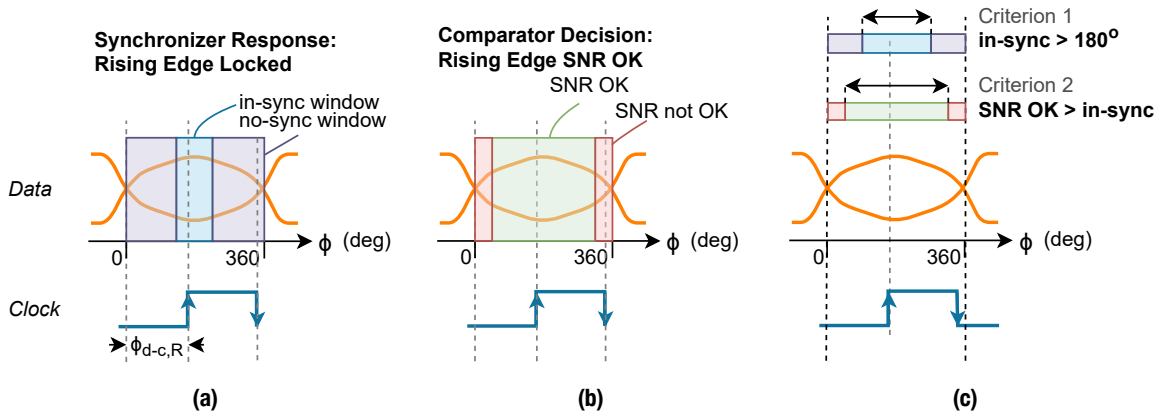


Figure 5.5: Requirements for the synchronizer feedback characteristics and comparator are defined in terms of (a) in-sync and (b) SNR OK windows, respectively. (c) In order to come to a successful receiver solution with edge swapping, the in-sync window must be wider than 180° and the comparator must provide sufficient SNR in that same window.

Optimum Edge Detection The challenge is data phase detection to pick the optimum clock edge. In essence, the figure of interest is the phase error. Instead of adding a separate phase detector at the input of the proposed receiver, this design exploits the characteristics of the comparator.

If the comparator samples far away from the middle of the eye diagram, the reduced input amplitude will cause a slower decision time by the comparator as the two-tail comparator features a regenerative latch. This *metastable* behaviour can be detected and indicate the need to swap the clock! A “metastability detector” can check whether one of the comparator outputs decided to rise after a fixed time delay. We expect that this delay (e.g. 50 ps) correlates to the actual SNR through the detected input amplitude.

5.2.2 Concept Soundness

Before moving on with design, we should prove that the two-tail comparator of last chapter is a proper candidate in terms of performance characteristics for the proposed edge swap based sampling phase correction. In other words: does the comparator slow down enough (become highly delayed due to metastability) when the BER rises as a consequence of incorrect sampling clock phase?

Window Criterion 1: Sync with Coarse Correction Two criteria must be met to confirm this suitability. First, looking at the eye diagram drawn in Fig. 5.5, there is a rising clock edge slightly off-center from the middle. This edge can be comfortably used for detection, and is in the “in-sync window”. With synchronization by metastability, use of the falling edge in the depicted scenario would cause a large metastable decision time and the feedback loop would command a clock swap. That is, the falling edge here is in the “no-sync window”. The clock swap concept embodies a 180° phase correction: to ensure that the complementary clock edge is located in an in-sync window, this in-sync window must be wider than $T_{CK}/2$!

Window Criterion 2: BER of Comparator The second criterion is related to the SNR specification that corresponds to the targeted BER. As shown in Fig. 5.5, the eye diagram can be linked to an “SNR OK” and “SNR not-OK” window based on the comparator performance. Samples initiated with a clock edge within the SNR OK window will result in a sufficiently low BER. Thus, to ensure that the synchronization yields a low, the second criterion is that the SNR OK window should be wider




Figure 5.6: In-sync window and SNR OK window simulations for the 10 Gbps two-tail comparator. For example, with a metastability detector comparison time threshold of 50 ps, the corresponding in-sync window lies between $t_{DtoC} = 83$ ps and $t_{DtoC} = 161$ ps. When the threshold lowers to 45 ps, the in-sync window shrinks, to lie in between 95 ps and 158 ps.

than (encapsulate) the in-sync window.

The criteria for the comparator are shown together in Fig. 5.5. It should be recognized that the width of the in-sync window depends on the metastability detector and comparator while the width of the SNR OK window depends entirely on the comparator design.

Two-Tail Criteria Testing The fast PNOISE comparator simulation methods from the previous chapter allow us to check whether the designed two-tail comparator passes the two in-sync and SNR OK window criteria. In short, the comparator is driven by a periodic clock and a period input signal resembling the expected eye diagram at the receiver input. The phase offset of the clock is swept to vary the sampling moment in the “eye”. For each phase offset, the output SNR and comparator comparison time is measured.

The simulation results for a 10 Gbps case are shown in Fig. 5.6. t_{DtoC} denotes the time between the rising edge of the data and rising edge of the clock. The additional metastability delay around the two crossings of the eye diagram is clearly visible. The maximum comparison times occur around $t_{DtoC} = 75$ and 175 ps instead of expected at 100 and 200 ps when the input signal crosses 0 V. This unexpected comparator characteristic can be explained by reminding ourselves that the two-tail comparator performs an integration over time after a rising edge, that is, it does not have a sample and held input. This means that the two-tail comparator makes its decision based on the moving average after the rising edge. Metastability therefore occurs somewhat earlier than the zero crossings of the input signal (“eye”).

Nonetheless, the criteria can still be tested against. It is revealed that the width of the in-sync window is 78 ps wide when the metastability detector would have a comparison time threshold of

50 ps. For a lower threshold of 45 ps, the in-sync window shrinks to 63 ps. Both are greater than $T_{CK}/2 = 50$ ps, so the first criterion is met.

A practical approach for checking the SNR-OK window is to measure whether the SNR at the edge of the widest in-sync window is sufficient (i.e. worst case allowable sampling moment). The widest in-sync window corresponds to a comparison time threshold of 50 ps. The synchronizer will never employ a comparison threshold higher than 50 ps, because at 10 Gbps the reset phase also claims 50 ps.

Applying this approach, it follows that the comparator will never make decision with an output SNR less than 60 V/V (e.g. at $t_{DIOC} = 85$ ps). A calculation using the known input-referred noise level, and equations from last chapter reveals that the BER remains below 10^{-8} including contributions of the transmitter, at 99% yield.

Thus, the designed two-tail comparator passes both criteria as illustrated underneath the simulation results and can be used to build the proposed synchronizer.

5.3 Self-Synchronizing Receiver Design

The proposed self-synchronizing receiver (Fig. 5.2) has three core functions:

- To detect data with a fast comparator.
- To detect a slow decision from the comparator (due to metastability).
- To swap the current clock configuration, i.e. make the comparator sample on the rising edge of the current complementary clock. At circuit level, this means that a *Swap Select* signal should be inverted.

This section covers the circuits designed for the metastability detector and clock swapper of this self-synchronizing receiver. In addition, the resulting system performance is presented.

5.3.1 Metastability Detector

Metastability of the two-tail comparator causes the neither of its output nodes to rise within a specified time after the start of the comparison phase. In terms of digital signal values: both output stay 0. A fast 3-input NOR gate¹ could check after the falling edge of the comparator clock whether metastability occurs (input = 000).

A simple but naive initial metastability detector is shown in Fig. 5.7. In the event of metastability a *FLAG* signal should rise to indicate the clock swapper to load the inverse of current *Swap Select*. The clock delay sets the maximum comparison time threshold $t_{COMP,MAX}$ and correspondingly the width of the in-sync window.

Key Issues There are two key issues to be solved with the initial detector. For one, a simple NOR gate is not sufficient to capture metastability at a specified moment in time. For example, if $t_{COMP,MAX} = 40$ ps, *FLAG* might be high at 40 ps, but might fall again at $t = 45$ ps because the comparator eventually makes a decision. What is more, the reset phase of the comparator will draw both output nodes low again. As a consequence, *FLAG* is only valid for a short amount of time if it is not supposed to be invoked.

An evident solution to this first issue is to make the NOR gate latching. To be exact, the inputs to the NOR gate should be frozen when the *FLAG* signal is to be evaluated. A track and hold (T&H)

¹Remember that a 3-input NOR gate gives a 1 output only if all input are 0.

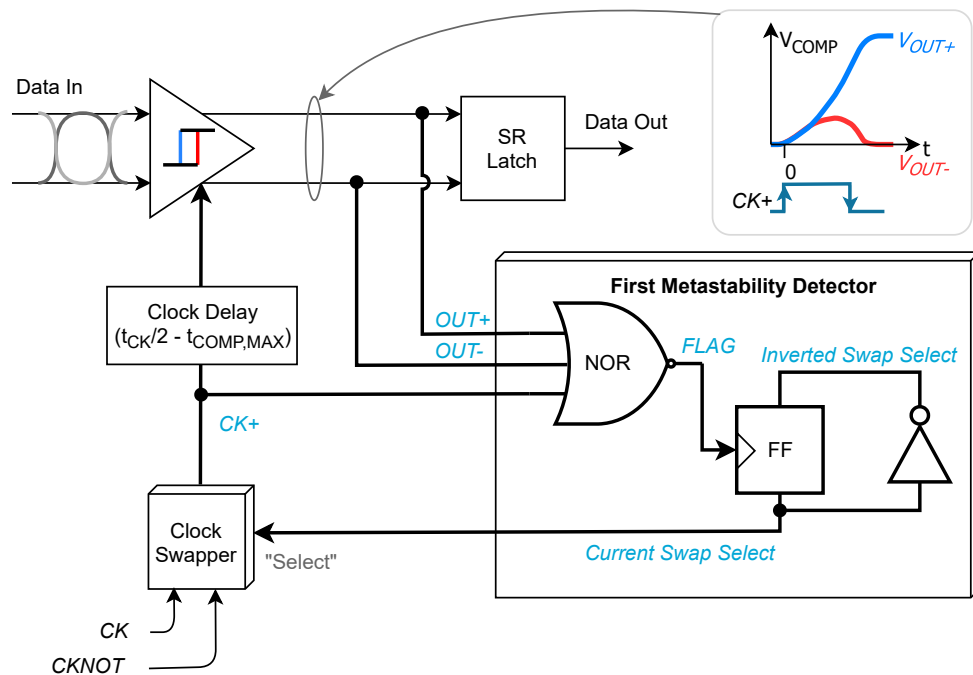


Figure 5.7: First proposed metastability detector loop utilizing a 3-input NOR and inverter based *Swap Select* computation.

transmission gate can fulfill this function. An SR latch would not be suitable because it would only be a remedy for the problem caused by the reset phase of the comparator.

The second issue involves the response of the clock swapper and the random nature of the data. The clock swapper will not respond directly to the command of the metastability detector, that is, there is a significant delay until the complete receiver has successfully swapped. This delay stems from the clock buffers chains etc. that need to have the new *Swap Select*'s effect fully propagated through. In an unfortunate case that the swap is in progress, but the detector once again detects metastability (for example 100 ps later), it will wrongly command the swapper to swap while it was still busy. In short, the latency of the clock swap operation might be too high compared to the input data rate.

This second issue can be solved by obtaining the latest *Swap Select* directly from the swapper (a delay element from the point of view of the metastability detector) instead of remembering it locally in the detector loop.

Fig. 5.8 shows this better implementation of the T&H circuit, NOR gate and feedback of the latest ("current") *Swap Select* from the clock swapper. When $CK+$ is high, the comparator is making its comparison. When $CK+$ drops low, the NOR's input is frozen and *FLAG* becomes valid. If *FLAG* = 1, then the computed *Swap Select* will be the inverse of the current *Swap Select*. At the following rising edge of $CK+$, a flip-flop consisting of two C2MOS latches will load in the computed *Swap Select* while another "compare" (and track) phase is initiated.

Final Asynchronous Loop However, now a third, new issue will pop up: when the clock swapper is swapping, the comparator might encounter a long stream of 1's or 0's. Since there are no data transitions (a "long run") the comparator will make fast decisions and the loop may wrongly command swapper to not swap anymore while it was correctly doing so based on a data transition a few bits ago. That is, the sampling metastability of the comparator will only appear when the data transitions from 0 to 1, or vice versa.

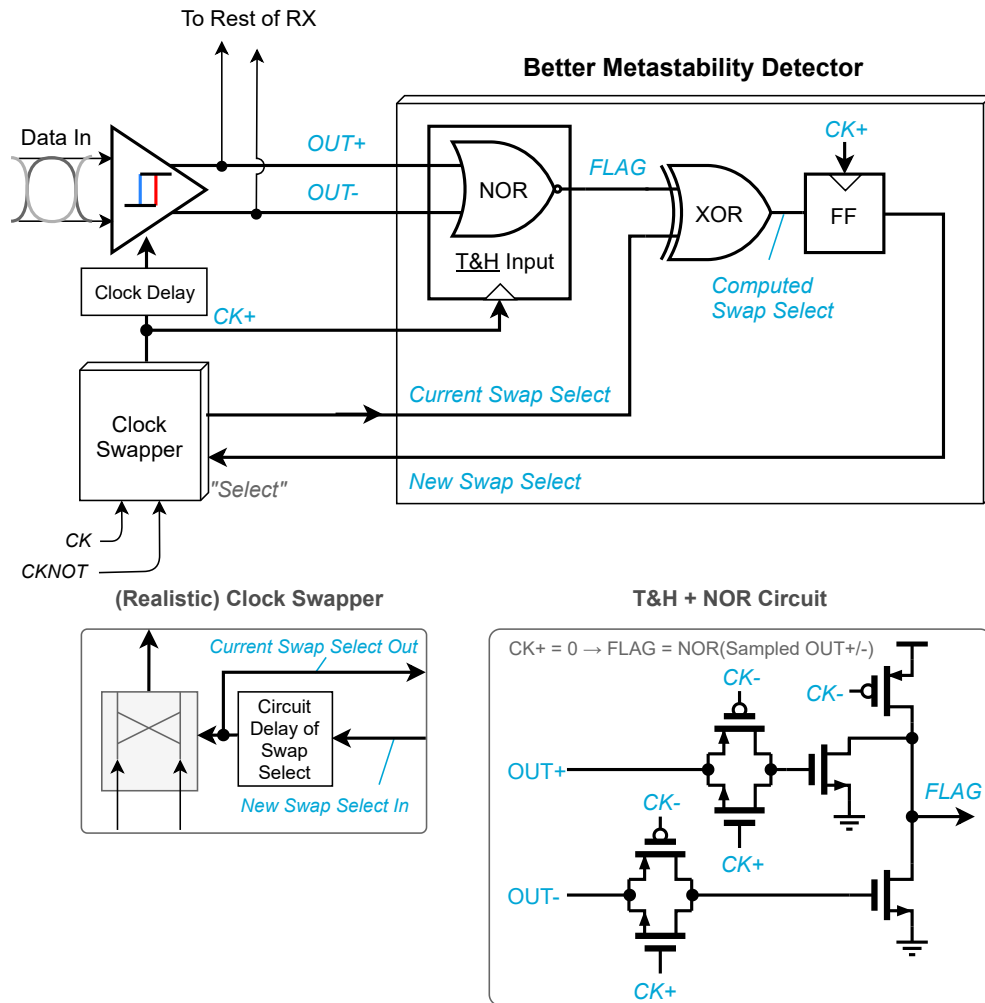


Figure 5.8: Second proposed metastability detector loop utilizing a track and hold before a pseudo-NMOS NOR gate. A realistic clock swapper will incur some delay before the swap is actually executed. The current clock swap select is not stored locally anymore.

To solve this final race condition, the total response time of *FLAG*, the clock swapper and the receiver's clock buffers should be less than 100 ps. This is an impossible target in 40 nm CMOS, even more so when avoiding current-mode circuits. Instead, we solve the problem with an *asynchronous* feedback mechanism: at the onset of a clock swap, the detection loop is disabled until the clock swap has completed.

The expansion to the final proposed asynchronous metastability detector loop state transitions and circuit are shown in Fig. 5.9 and Fig. 5.10. Its goal is to disable itself while the clock swapper is in progress of swapping. Asynchronous systems always need a *start* and *completion* signal [3]. We generate these signals as follows:

- **Start:** the *BUSY* signal acts as the start signal and is a flip-flopped version of *FLAG*. Once *BUSY* is 1, it will enable a freeze circuit to prevent the flip-flops from loading new values. Secondly, the current *Swap Select* before swapping is remembered by a latch.
- **Completion (Discharge Activation):** After successful swapping, the discharge block will stop the asynchronous freeze by discharging *BUSY* to 0. The discharge block is initiated when the stored old *Swap Select* is different to the updated current *Swap Select* (hence the XOR). In

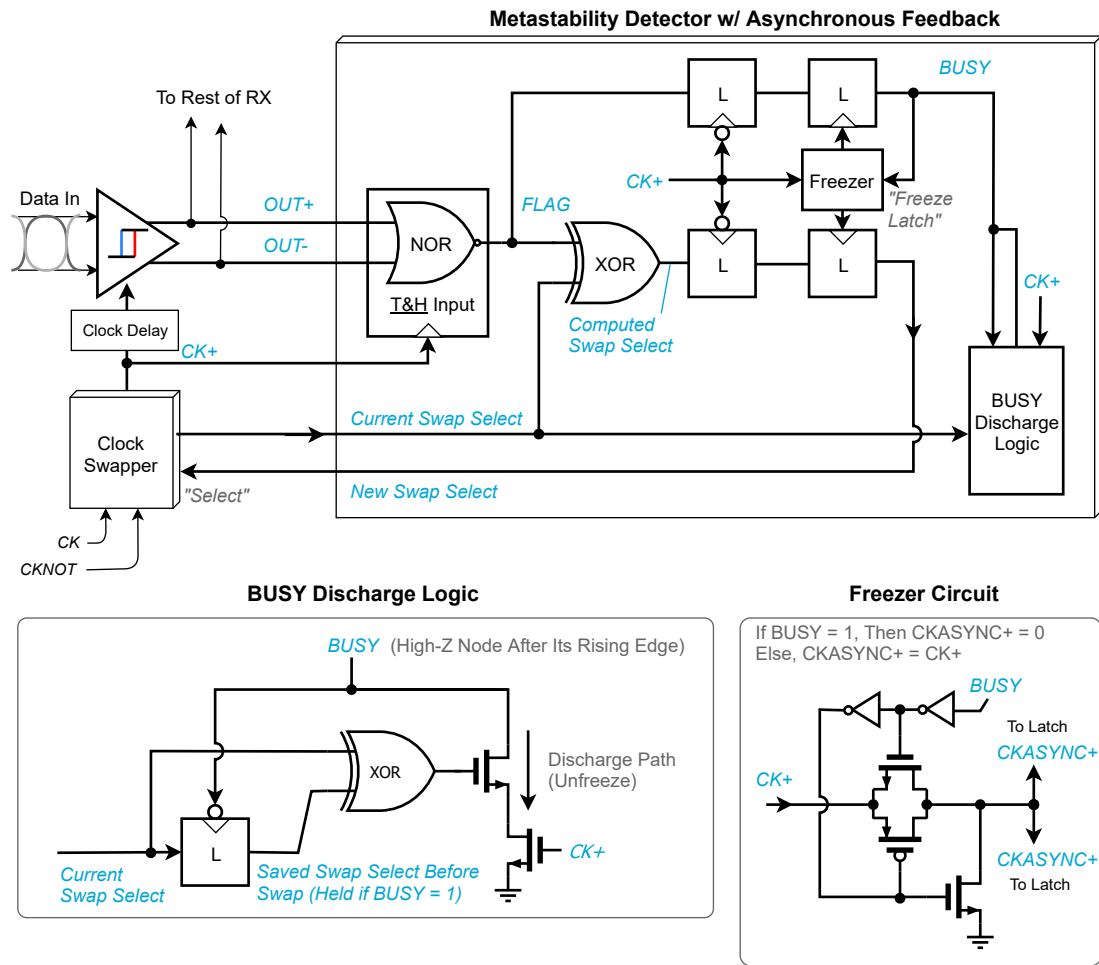


Figure 5.9: Final proposed metastability detector loop utilizing an asynchronous loop to prevent race conditions while swapping for synchronization. *BUSY* implements the start signal and freezes the latches, while the NMOSes in the discharge logic carry out the asynchronous completion signaling.

addition, the post-swap (!) *CK+* should be high to enable the discharge block, so that the unfreezing happens suitably at the start of a new comparison.

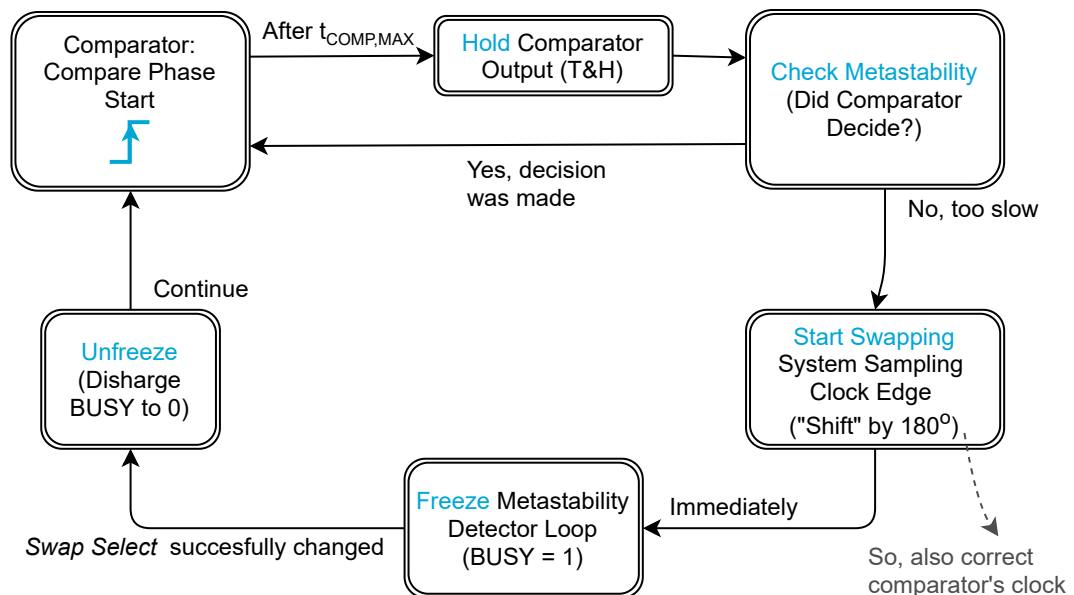


Figure 5.10: State diagram of final proposed metastability detector loop to solve the identified issues.

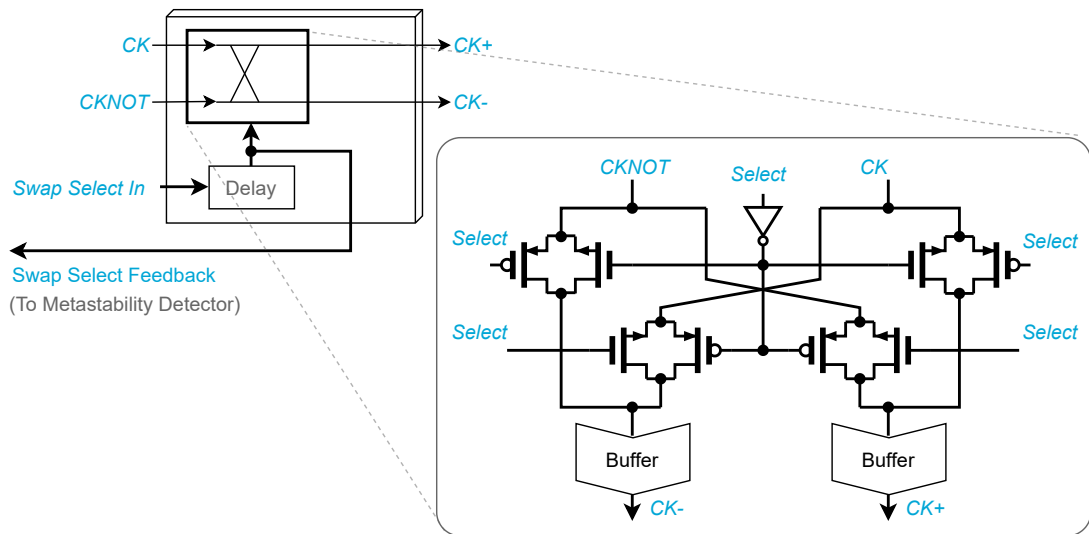


Figure 5.11: Transmission gate implementation of the clock swapping mechanism in this work. The transmission gates are large to prevent clock speed degradation, thereby causing a delay of the response to *Swap Select*.

5.3.2 Clock Swapper

By virtue of the flexible asynchronous metastability detector design, the clock swapper does not need to adhere to strict requirements. Essentially the clock swapper has two states of operation, picking one of the two forwarded clocks to be $CK+$ and the other to be $CK-$. In fact, any complementary equal-rate clock can be used for this receiver.

Since the forwarded clock is complemented, the clock swap could be implemented by a controlled inversion, that is, using a XOR gate for each clock path. Or, two NAND-based multiplexers could select which of the clocks is passed on. The minor downside of these two solutions is that they have some asymmetric delay resulting in a clock swap that is not exactly 180° . The in-sync window as explained in Section 5.2.2 should then be slightly wider than $T_{CK}/2$ which would have been a problem for a slower comparator.

Shown in Fig. 5.11 is the clock swapper designed for this work, instead built from four transmission gates. This solution is symmetric, achieving a perfect 180° swap. The transmission gates are sized wide to be able to pass the fast clock signals. The resulting delay to drive these transmission gates is not a problem—the metastability loop is asynchronous.

5.3.3 Simulation Results

The challenge of this receiver design lied (unexpectedly) in the complex creation of an innovative metastability detector: we were able to verify the performance of the self-synchronized receiver on schematic level (pre-layout), but did not manage to fully layout the complete self-synchronized receiver in the available time of the thesis project.

Given that the performance of this receiver depends only on the logical *reasoning* behind the metastability detection loop and the speed/BER performance of the two-tail comparator, there is luckily no major problem. Also, as discussed in the last chapter, we did manage to observe excellent performance for a compact, post-layout two-tail comparator.

The intended behavior of the detection loop stands firm in circuit simulations, as shown in Fig. 5.12 for a 10 Gbps pre-layout simulation. The top graph measures the comparison time of the




Figure 5.12: Pre-layout simulation of the receiver. The received data is purposely made out-of-sync after a short delay.

comparator over time. On purpose, we provide an out-of-sync data signal to the receiver after a short delay into the simulation. As soon as the comparison time exceeds ~ 46 ps, the computed *Swap Select* signal changes its value and the loop becomes frozen ($BUSY = 1$). At 500 ps into the simulation, the clock *CK+* successfully swaps and after a short delay, the discharge event takes place. We observe that the comparison time of the receiver's comparator stays constant after swapping for the correct sampling edge.

The simulation further indicate a total power efficiency of 0.28 mW/Gbps for the clock swapper and metastability detection loop together². The detection loop only draws 0.02 mW/Gbps. Thus, the proposed digital-heavy detection loop implementation pays off well in power efficiency.

Naturally, it is our intent to update this thesis with the final performance metrics once a layout is completed. We expect the clock swapper to occupy 400 μm^2 and the detection loop to occupy 100 μm^2 at most.

²the comparator draws 0.15 mW/Gbps as discussed in the previous chapter

5.4 Chapter Conclusions

This chapter discussed the general considerations for clock alignment in wireline receivers to come to the decision of designing a new approach for synchronization. The main priority of this design was to be simple and make use of the forwarded clock of the transmitter and the comparator design of last chapter.

We contributed to the goal of this thesis with the design of an innovative synchronization concept based on metastability detection and clock swapping. A practical design of a metastability detection loop proved to be extremely challenging, so we proposed an elegant asynchronous feedback mechanism inside the detection loop. This additional asynchronous feedback ensures proper data synchronization under all conditions of random data detection.

Simulations confirm that the proposed self-synchronizing wireline receiver loop can successfully operate at 10 Gbps and $\text{BER} < 10^{-8}$ while employing our fast two-tail comparator. The expected power draw and area of this proposed receiver are much smaller compared to those of the wireline transmitter. Thus, in the scope of this thesis, this chapter overcomes the final, critical challenge of sampling clock alignment, in agreement with the set wireline link design target.

6 Conclusion and Future Work

6.1 Conclusion

This master thesis presented the development of a wide variety of circuits and systems for fast wireline transceiver links to initiate this research topic at ELCA.

We developed a new perspective on the analysis of bit error rates in wireline links based on the eye quality. We proposed the use of the PAM spectral design space chart: a novel, visual system design analysis tool for PAM wireline circuit designers. At a glance, we can show designers what their highest achievable spectral efficiency is, what modulation complexity should be chosen and what link budget should be maintained.

We contributed to the development of wireline transmitter circuits by designing and taping out a < 3 mW/Gbps per lane, 10 Gbps wireline transmitter in 40 nm CMOS. The proposed inherently pipelined 16:1 multiplexer and current mode logic driver design procedure are the key enablers for this speed performance at low energy consumption. Moreover, our design of a flexible pseudo random bit generator removes the need for any on-chip test signal storage.

Finally, for the development of a 10 Gbps wireline receiver, we proposed a novel self-synchronized wireline receiver architecture. This contribution removes the need of a classical clock and data recovery loop. At its core, the proposed receiver comprises the design of a low power, high-speed two-tail comparator and a creative asynchronous metastability detection loop. In our journey towards these goals, we closely examined comparator circuit design for wireline applications. We contributed to our comparator design goals by a new analysis of bit error rate degradation due to hysteresis and a set of fast and accurate simulation methods.

For the sake of comparison, we can combine our measured transmitter performance and expected receiver performance to plot the fictive transceiver performance achieved in this work versus recent wireline transceiver publications, in terms of power efficiency and speed. Fig. 6.1 shows this plot and the best performance seen in 40 nm designs is marked by the dashed line. The design effort in this work does a good job, especially considering that the y-axis is logarithmic and we are taking up this research topic for the first time.

We conclude that we successfully opened the door towards highly integrable RF digital-to-analog converters at ELCA. In a short amount of time, we managed to tape-out a full chip and contributed to various new insights on high-speed wireline system, transmitter and receiver design.

6.2 Future Work

There is only a limited time available for the execution of a master thesis project. It is the intent of the author to further improve and develop the proposed receiver while publishing some of the contributions described above.

Nevertheless, to continue the advance at the topic of high-speed wireline links at ELCA, it would




Figure 6.1: Power efficiency versus data rate of recent wireline transceiver publications obtained via [19] (here limited to 40 Gbps). The pink diamonds mark the 40 nm designs. It is meritable to move towards the dashed line, that is to the right (high speed) and down (less pJ/bit).

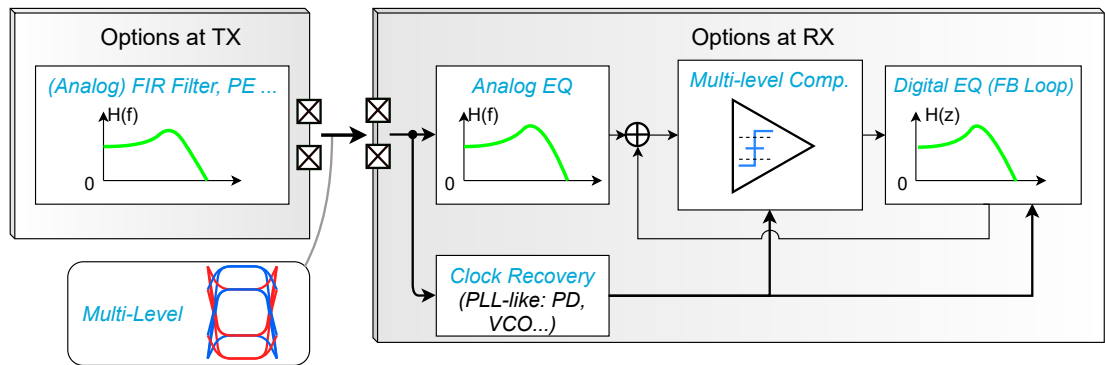


Figure 6.2: Overview of future possible wireline link research topics.

be interesting to either attempt to boost the data rate beyond 10 Gbps or introduce more complexity in functionality to overcome more challenging data transmission channels.

Spectral Efficiency To boost the data rate with the current available bandwidth demands the development of multi-level PAM signaling. The current-mode driver in this work can be expanded PAM-4 or -8 data transmission by connecting multiple drivers in parallel. The main challenge in going towards multi-level signaling is overcoming the tough SNR constraints at low supply voltages. So, novel data detection solutions at the receiver side should be developed.

Channel Equalization Another option to increase the the data rate involves the addition of channel equalization of bandwidth expansion techniques [2]. A classic, but area-hungry option would be the addition of inductive loads to cancel out the channel roll-off. More geared towards the advance of all-digital circuits and systems is the development of finite impulse response pre-emphasis at the transmit side or a digital (high-pass) equalization feedback loop at the receiver side. A challenge for this channel equalization is the need for implementing multi-phase clock systems, as the presented 10 Gbps designs are starting to see the limits of rise and fall time in 40 nm CMOS.

Clock and Data Recovery At some point in the future it might be necessary to expand our area of application beyond the interface to RF-DACs. It might become unpractical to work with a global (or forward) clock that is synchronized. Therefore, it is worthwhile to investigate the development of novel clock and date recovery loops.

Fig. 6.2 illustrates our future outlook. It is clear that the contributing efforts of this thesis project will be a great setup for interesting, follow-up wireline link projects at the ELCA research group.

Bibliography

- [1] M. Mehrpoo, M. Hashemi, Y. Shen, L. C. De Vreede, and M. S. Alavi, "A Wideband Linear-Interleaving DDRM," *IEEE Journal of Solid-State Circuits*, vol. 53, no. 5, pp. 1361–1373, 5 2018.
- [2] B. Razavi, *Design of Integrated Circuits for Optical Communications*, 2nd ed. Hoboken, NJ: John Wiley & Sons, 2012.
- [3] J. M. Rabaey, A. Chandrakasan, and B. Nikolic, *Digital Integrated Circuits: A Design Perspective*, 2nd ed. Upper Saddle River, NJ: Prentice Hall, 2003.
- [4] W. J. Dally and J. W. Poulton, *Digital Systems Engineering*. Cambridge: Cambridge University Press, 1998.
- [5] B. Sklar, *Digital Communications: Fundamentals and Applications*, 2nd ed. Upper Saddle River, NJ: Prentice Hall, 2001.
- [6] E. Säckinger, *Analysis and Design of Transimpedance Amplifiers for Optical Receivers*. Hoboken, NJ: John Wiley & Sons, 2018.
- [7] L. W. Couch, *Digital and Analog Communication Systems*, eighth ed. Upper Saddle River, NJ: Prentice Hall, 2013.
- [8] S. Palermo, "High-Speed Serial I/O Design for Channel-Limited and Power-Constrained Systems," in *CMOS Nanoelectronics Analog and RF VLSI Circuits*, 2011, ch. 9.
- [9] K. Fukuda, H. Yamashita, G. Ono, R. Nemoto, E. Suzuki, N. Masuda, T. Takemoto, F. Yuki, and T. Saito, "A 12.3-mW 12.5-Gb/s Complete Transceiver in 65-nm CMOS Process," in *IEEE Journal of Solid-State Circuits*, vol. 45, no. 12, 12 2010, pp. 2838–2849.
- [10] B. Razavi, *Design of Analog CMOS Integrated Circuits*, 2nd ed. Upper Saddle River, NJ: Prentice Hall, 2017.
- [11] M. Pelgrom, *Analog-to-Digital Conversion*, 3rd ed. Cham: Springer, 2017.
- [12] D. Schinkel, E. Mensink, E. Klumperink, E. Van Tuijl, and B. Nauta, "A double-tail latch-type voltage sense amplifier with 18ps setup+hold time," in *Digest of Technical Papers - IEEE International Solid-State Circuits Conference*, 2007.
- [13] M. Van Elzakker, E. Van Tuijl, P. Geraedts, D. Schinkel, E. A. Klumperink, and B. Nauta, "A 10-bit Charge-Redistribution ADC Consuming 1.9 μ w at 1 MS/s," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 5, pp. 1007–1015, 5 2010.
- [14] H. S. Bindra, C. E. Lokin, D. Schinkel, A. J. Annema, and B. Nauta, "A 1.2-V Dynamic Bias Latch-Type Comparator in 65-nm CMOS With 0.4-mV Input Noise," *IEEE Journal of Solid-State Circuits*, vol. 53, no. 7, pp. 1902–1912, 7 2018.

- [15] A. M. A. Ali, *High Speed Data Converters*. Institution of Engineering and Technology, 2016.
- [16] A. T. Ramkaj, M. Strackx, M. S. Steyaert, and F. Tavernier, "A 1.25-GS/s 7-b SAR ADC With 36.4-dB SNDR at 5 GHz Using Switch-Bootstrapping, USPC DAC and Triple-Tail Comparator in 28-nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 53, no. 7, pp. 1889–1901, 7 2018.
- [17] A. T. Ramkaj, M. S. Steyaert, and F. Tavernier, "A 13.5-Gb/s 5-mV-Sensitivity 26.8-ps-CLK-OUT Delay Triple-Latch Feedforward Dynamic Comparator in 28-nm CMOS," *IEEE Solid-State Circuits Letters*, vol. 2, no. 9, pp. 167–170, 9 2019.
- [18] P. Jespers, *The gm/ID Methodology, A Sizing Tool for Low-voltage Analog CMOS Circuits*. Springer US, 2010.
- [19] T. Anand, "Wireline Link Performance Survey." [Online]. Available: <https://web.engr.oregonstate.edu/~anandt/linksurvey>
- [20] J. Lee, P. C. Chiang, P. J. Peng, L. Y. Chen, and C. C. Weng, "Design of 56 Gb/s NRZ and PAM4 SerDes Transceivers in CMOS Technologies," pp. 2061–2073, 9 2015.
- [21] P. J. Peng, J. F. Li, L. Y. Chen, and J. Lee, "A 56Gb/s PAM-4/NRZ Transceiver in 40nm CMOS," in *Digest of Technical Papers - IEEE International Solid-State Circuits Conference*, vol. 60. Institute of Electrical and Electronics Engineers Inc., 3 2017, pp. 110–111.
- [22] J. Han, Y. Lu, N. Sutardja, and E. Alon, "A 60Gb/s 288mW NRZ Transceiver With Adaptive Equalization and Baud-Rate Clock and Data Recovery in 65nm CMOS Technology," in *Digest of Technical Papers - IEEE International Solid-State Circuits Conference*, vol. 60. Institute of Electrical and Electronics Engineers Inc., 3 2017, pp. 112–113.
- [23] K. Huang, D. Luo, Z. Wang, X. Zheng, F. Li, C. Zhang, and Z. Wang, "A 190mW 40Gbps SerDes Transmitter and Receiver Chipset in 65nm CMOS Technology," in *Proceedings of the Custom Integrated Circuits Conference*, vol. 2015-November. Institute of Electrical and Electronics Engineers Inc., 11 2015.
- [24] X. Zheng, C. Zhang, F. Lv, F. Zhao, S. Yue, Z. Wang, F. Li, H. Jiang, and Z. Wang, "A 4-40 Gb/s PAM4 Transmitter With Output Linearity Optimization in 65 nm CMOS," in *Proceedings of the Custom Integrated Circuits Conference*, vol. 2017-April. Institute of Electrical and Electronics Engineers Inc., 7 2017.
- [25] B. Razavi, "Historical Trends in Wireline Communications: 60CE Improvement in Speed in 20 Years," *IEEE Solid-State Circuits Magazine*, vol. 7, no. 4, pp. 42–46, 9 2015.
- [26] F. O'Mahony, J. E. Jaussi, J. Kennedy, G. Balamurugan, M. Mansuri, C. Roberts, S. Shekhar, R. Mooney, and B. Casper, "A 47 CE 10 Gb/s 1.4 mW/Gb/s Parallel Interface in 45 nm CMOS," in *IEEE Journal of Solid-State Circuits*, vol. 45, no. 12, 12 2010, pp. 2828–2837.
- [27] T. Kim, S. Jang, S. Kim, S. H. Chu, J. Park, and D. K. Jeong, "A Four-Channel 32-Gb/s Transceiver with Current-Recycling Output Driver and On-Chip AC Coupling in 65-nm CMOS Process," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 61, no. 5, pp. 304–308, 2014.
- [28] K. Gharibdoust, A. Tajalli, and Y. Leblebici, "Hybrid NRZ/Multi-Tone Serial Data Transceiver for Multi-Drop Memory Interfaces," *IEEE Journal of Solid-State Circuits*, vol. 50, no. 12, pp. 3133–3144, 12 2015.

- [29] W. S. Choi, G. Shu, M. Talegaonkar, Y. Liu, D. Wei, L. Benini, and P. K. Hanumolu, "A 0.45-to-0.7V 1-to-6Gb/s 0.29-to-0.58pJ/b Source-Synchronous Transceiver Using Automatic Phase Calibration in 65nm CMOS," in *Digest of Technical Papers - IEEE International Solid-State Circuits Conference*, vol. 58. Institute of Electrical and Electronics Engineers Inc., 3 2015, pp. 66–67.
- [30] R. Inti, S. Shekhar, G. Balamurugan, J. Jaussi, C. Roberts, T. C. Hsueh, and B. Casper, "A 0.5-to-0.75V, 3-to-8 Gbps/Lane, 385-to-790 fJ/b, Bi-Directional, Quad-Lane Forwarded-Clock Transceiver in 22nm CMOS," in *IEEE Symposium on VLSI Circuits, Digest of Technical Papers*, vol. 2015-Augus. Institute of Electrical and Electronics Engineers Inc., 8 2015, pp. C346–C347.
- [31] B. Razavi, "The StrongARM latch [A Circuit for All Seasons]," *IEEE Solid-State Circuits Magazine*, vol. 7, no. 2, pp. 12–17, 3 2015.
- [32] T. Rabuske and J. Fernandes, "Noise-aware simulation-based sizing and optimization of clocked comparators," *Analog Integrated Circuits and Signal Processing*, vol. 81, no. 3, pp. 723–728, 12 2014. [Online]. Available: <https://link.springer.com/article/10.1007/s10470-014-0428-4>
- [33] H. Omran, "Fast and accurate technique for comparator offset voltage simulation," *Microelectronics Journal*, vol. 89, pp. 91–97, 7 2019.

Appendices

A Wireline Transceiver Literature Review

We should preface the results of this review by saying that the observed developments are not representative for the starting point of this thesis project. In this thesis project, we are cultivating design experience by building a wireline link from the ground up in a short time span. The below observations are most useful for future development of this topic at ELCA.

A.1 Modern Wireline Transceivers

The increase of wirelink data rates whilst maintaining an economical power efficiency is partially made possible by active research efforts on wirelink transceiver ICs. Indeed, solid-state circuit conferences even have wireline transceivers as distinct topics.

This subsection reviews state-of-the-art transceivers (or relevant subsystems) published in literature. Most of the papers were picked based on their exceptional speed or power efficiency and originate from the wirelink survey available at [19]. This survey collects wirelink transceiver publications presented at prominent IEEE conferences. We provide reorganized tables after this section.




Figure A.1: Reported wireline TRX data rate versus technology. Data points via [19].

A.1.1 General Trends

Fig. A.1 shows the data rate per lane as a function of CMOS technologies. A gap exists between the 40 Gbps and 60 Gbps groups: this can be explained by the fact that almost all high-speed wireline




Figure A.2: Reported wireline TRX power efficiency versus technology. Data points via [19].

interfacing standards are inspired by the OIF-CEI standard. Currently, the highest lane data rates specified by OIF-CEI are 28 Gbps and 56 Gbps.

As expected, smaller technology nodes are able to achieved higher transfer rates as their FO4 delays¹ shorten, allowing faster clock networks. Publications using 80 nm technology or larger hardly report a data rate much larger that 20 Gbps. We can therefore assume that the 40+ Gbps rates only started to become feasible at the 65 nm node and below. Moreover, we can also expect that the smallest technology nodes can achieve data rates much higher than 60 Gbps. Extrapolating the fastest data points, we should expect that 7 nm designs at 100 Gbps should be possible. Of course this boldly assumes that suitable transmission channels are available.

Fig. A.2 shows the the overall link power efficiency versus technology used. The power efficiency has clearly improved by moving to smaller technologies and efficiencies around 1 pJ/bit have been reported now. This can be partially explained by the common use of current-mode implementations for the high-speed output driver. These rely on an always-on current source, so an increase of the data rate immediately leads to a better driver power efficiency in terms of pJ/bit. Moreover, the (limited) lowering of the supply voltage also helps to reduce power, assuming that the swing stayed constant.

Around the 40 nm era, the efficiency improvements seem to decline. We can possibly explain this by the increasing channel losses that are encountered by operating PAM-2 at high data rates. More power is then expended for equalization techniques.

Unfortunately, a large share of the publications do no report the channel characteristic. In spite of this, we can view Fig. A.3. Here, we see that obtaining excellent power efficiencies is currently not possible for 40+ Gbps transceiver links. So, it is a possibility that heavy channel losses account for a penalty or limit in power efficiency.

¹FO4 delay is the input to output delay of an inverter that drives another inverter that is 4 times as big, that is, with a fan-out of four. It is a common process technology speed performance metric (for digital circuits).




Figure A.3: Reported wireline power efficiency versus data rate for 65 nm or smaller. Data points via [19].

A.1.2 High Data Rates at 65 and 40 nm

The emergence of the wireline TRX at rates > 20 Gbps is seen starting at the 65 nm node. In this subsection, we highlight core techniques employed in TRX publications that successfully approach the speed limits of the technology while maintaining moderate system power efficiency. To this end, we constrain ourselves to TRXs operating at 40(+) Gbps with 65, 45 or 40 nm CMOS technology.

[20] presents a 65-nm PAM-2 TRX and a 40-nm PAM-4 TRX both operating at 56 Gbps. Two years later, the follow-up paper [21] managed to make the 40-nm design flexible to be able to run in PAM-2 or PAM-4 mode. The power efficiency was improved to 10.4 pJ/bit. The key enabling techniques used by the authors are the combination of quarter-rate clock, equalization (EQ) at the TX and RX, a fast phase interpolator at the RX (for demultiplexing) and smart layout for large distance clock distribution. The NRZ variant required very fine control of the final MUX at the transmitter side.

[22] manages to realize a 60 Gbps TRX at 4 pJ/bit in 65 nm. The highlights of the system are the extensive use of equalization at the receiver (both analog and digital EQ), the low power clock and data recovery system and large number of offset calibration loops to meet the BER target.

[23] achieves 40 Gbps at 4.75 pJ/bit in 65 nm. The authors circumvent speed limitations by means of multiphase clocking at the receiver and the ability to calibrate the timing of the transmitter. Of course, equalization was also implemented to combat the channel loss. It should be noted that the TX did not have MUX stage, so the real-life power efficiency should be worse.

It must be noted that in Fig. A.3, it seems like an even more efficient 65 nm transceiver has been published at 40 Gbps, i.e., [24]. However, this paper only reports the design and measurement of a transmitter², so it cannot be compared to power efficiencies of whole transceivers.

A.1.3 Low-Power Design

A wireline communication survey published in 2015 notes that a 30% improvement in power efficiency can be expected for the full link [25]. Based on the latest entry in that survey (6 pJ/bit in 2014), we would expect that the TRX power efficiencies published in 2019 float around 1 pJ/bit.

²The publication seems to be mistakenly included.

This prediction has not exactly been fulfilled: the best power efficiencies encountered in preparation of this thesis drop below 1 pJ/bit and were published in 2017 or earlier. In this subsection, we focus on the core techniques that such ultra power efficient employ.

[26] presents a SerDes link in 40 nm achieving a power efficiency of 1.4 pJ/bit and 10 Gbps per lane. Because their application allows for interconnect with many wires, they choose to share the forwarded clock among a group of lanes (the aggregated data rate was 470 Gbps). Therefore, power of the clock generation is divided across the lanes as a result of this system level choice. The reported TRX has a very low swing of 0.15 V differential also contributing to the good efficiency. The BER target was maintained by building in calibration for the detector to have the effective sampling thresholds at the ideal points.

In the category of reusing, an interesting concept was presented in 65 nm by [27]. They exploit the fact that the supply voltage greater than the differential swing by stacking voltage mode output drivers. Intuitively, the stacked drivers can be seen as a resistor ladder. The benefit is that the current than returns from the channel is recycled! However, the increased circuit complexity to ensure that the output resistance continues to be matched results partially in a moderate efficiency of 2 pJ/bit at 8 Gbps per lane.

The best power efficiencies realized with the 40 nm node have been reported by [28], touching 1 pJ/bit at 9 Gbps. They exploit the characteristic notch seen when looking single-endedly into a differential microstrip line. The second “bulge” in the channel profile is used to transmit some more data on a carrier. This is a smart use of concepts in wireless communication but adds complexity on the

Let us finally review the differential wireline TRXs with sub 1 pJ/bit performance. [9] is one of the first to report such performance: 12.5 Gbps per lane at 0.98 pJ/bit in 65 nm. The design aggressively employs CMOS-like structures to avoid idle (bias) currents. In addition, the paper does a great effort in reducing the power generated in clock distribution and clock-data recovery.

[29] achieves in 65 nm at 6 Gbps a power efficiency of only 0.58 pJ/bit. The key design choice at the system level in this design is the low voltage supply operation of down to only 0.45 V. In addition, eight clock phases are generated to allow for low-clock rate system operation. Phase calibration is implemented to evenly space the clock phases. The low supply voltage demands a voltage mode transmitter. Compared to [9] though, there is no CDR system as the clock is simply forwarded.

[30] reports another low-voltage TRX design. At 0.5 V supply voltage, they achieve 0.385 pJ/bit at 3 Gbps per lane in 22 nm. Two key techniques can be noted aside from going for a low supply voltage. At the circuit level implementation, the output driver is a new type of shunt terminated voltage mode driver, removing the need for a voltage regulator. For the clocking circuitry, the designers have managed to create quadrature phases of the clock for the receiver based on a all-digital delay line.

A.2 Original Data

The following pages contain the publications gathered in [19]. The data has been reordered and colorized for the comfort of the reader.

Year	Publication	Title	First Author	Process	[Gbps]	[pJ/bit]	[mW]	Loss (dB)	Loss Frequency (GHz)
2019	ISSCC	A 180mW 56Gb/s DSP-Based Transceiver	Tamer Ali	7	56	3.21	180	35	14
2019	ISSCC	A Sub-250mW 1-to-56Gb/s Continuous-Rate	Matteo Pisati	7	56	4.17	234	42.5	14
2019	ISSCC	A 60Gb/s PAM-4 ADC-DSP Transceiver in 16nm	Marc-Andre LaCroix	7	60	6.9	414	32	15
2018	VLSI	A 0.5-28Gb/s Wireline Transceiver with 15-Tap	Jay Im	7	28	10.11	283	30	
2016	ISSCC	A 32Gb/s Bidirectional 4-Channel 4pJ/b Cap	Chintan Thakkar	14	8	4	32		
2018	ISSCC	A 4-Lane 1.25-to-28.05Gb/s Multi-Standard	hammad Sadegh Jafar	14	28.05	6.06	170	40	14.025
2018	ISSCC	A 64Gb/s PAM-4 Transceiver Utilizing an AI	Luke Wang	16	64.375	5.8	373.6	29.5	
2018	ISSCC	A Fully Adaptive 19-to-56Gb/s PAM-4 Wirel	Parag Upadhyaya	16	56	9.7	545	32	14
2016	VLSI	A 56Gb/s PAM4 Wireline Transceiver using	Yohan Frans	16	56	9.82	550	25	
2016	VLSI	A Fully-Adaptive Wideband 0.5-32.75Gb/s	P. Upadhyaya	16	32.75	17.6	577	30	
2014	VLSI	A 36 Gb/s 16.9mW transceiver in 20nm CM	T. Hashida	20	36	16.94	609.9	20	
2015	ISSCC	A 0.5-to-32.75Gb/s Flexible-reach wireline t	P. Upadhaya	20	32.75	23.97	785	10.4	
2017	ISSCC	A transmitter and receiver for 100Gb/s cohe	Jun Cao	20	64	24.53	1570		
2015	VLSI	A 0.5-to-0.75V, 3-to-8 Gbps/lane, 385-to-79	Rajesh Inti	22	8	0.79	6.32	12	
2015	VLSI	A 1.2-2.5Gb/s 1.4-2pJ/b Serial link in 22nm	Sudip Shekhar	22	5	2	10		
2014	ISSCC	a 205mW 32Gb/s 3-tap FFE/6-tap DFE bidi	J. Jaussi	22	32	6.41	205	26	
2016	ISSCC	Pin-Efficient 20.83Gb/s/wire 0.94pJ/bit Forw	Amin Shokrollahi	28	20.83	0.94	19.58	3	
2016	ISSCC	A 38mW 40Gb/s 4-Lane Tri-Band PAM-4 /	Wei-Han Cho	28	10	0.95	9.5		
2018	ISSCC	A 126mW 56Gb/s NRZ Wireline Transceiver	Marc Erett	28	56	2.25	126	8	
2013	VLSI	A 2.8 mW/Gb/s quad-channel 8.5?11.4 Gb/s	A. Nazemei	28	11.4	2.82	32.2	21	
2019	ISSCC	A 32Gb/s 2.9pJ/b Transceiver for Sequence	Aurangozeb	28	28	2.9	93	30	8
2017	ISSCC	12Gb/s over four balanced lines utilizing NR	Yeonho Lee	28	12	3.08	36.94		
2019	ISSCC	A 25.6Gb/s Uplink-Downlink Interface Empl	Takashi Toi	28	25.6	3.69	94.4	1.84	6.4
2015	VLSI	A 3.8 mW/Gbps Quad-Channel 8.5-13 Gbps	Tamer Ali	28	13	3.8	49	35	
2016	VLSI	A 125 mW 8.5- 11.5 Gb/s Serial Link Transc	Bharath Raghavan	28	11.5	4	46	38	
2016	ISSCC	A 56Gb/s NRZ-Electrical 247mW/lane Serial	Takayuki Shibasaki	28	56.2	4.39	247	18.4	
2018	ISSCC	A 4.9pJ/b 16-to-64Gb/s PAM-4 VSR Transc	Emanuele Depaoli	28	64	4.9	180	16.8	16
2014	ISSCC	a 130nm 20Gb/s half-duplex serial link in 2	V. Balan	28	20	6.5	130	20	
2016	VLSI	A 28.3 Gb/s 7.3 pJ/bit 35 dB Backplane Tra	Hiroki Miyaoka	28	28.3	7.3	206.59	35	
2013	ISSCC	A 6.4Gb/s near-ground single-ended transce	K. Kaviani	28	6.4	9.1	58.24		
2015	ISSCC	a 28gb/s multi-standard serial link transceiv	B. Zang	28	28	10.54	295	40	
2016	VLSI	A 125 mW 8.5- 11.5 Gb/s Serial Link Transc	Bharath Raghavan	28	10.3125	12.12	125		
2016	ISSCC	A 25Gb/s Multistandard Serial Link Transce	akayasu Norimatsu	28	25	16.12	403	50	
2014	ISSCC	28Gb/s 560mW Multi-Standard SerDes with	H. Kimura	28	28	20	560	34	
2016	ISSCC	A 40/50/100Gb/s PAM-4 Ethernet Transceiver	arthik Gopalakrishna	28	28	21.4	600		
2012	CICC	Design of high-speed wireline transceivers f	J. Savoj	28	13.1	21.53	282	31	
2014	VLSI	A 40Gb/s serial link transceiver in 28nm CM	E-H. Chen	28	40	23.18	927	20	
2015	ISSCC	Multi-Standard 185fsrms 0.3-to-28Gb/s 40d	T. Kawamoto	28	28	25.1	702.8	40	
2012	VLSI	A wide common-mode fully-adaptive multi-s	J. Savoj	28	12.5	26.8	335	33	
2015	CICC	A 1.8-pJ/bit 16x16-Gb/s Source Synchroniz	Timothy O. Dickson	32	16	1.9	30.4	10	
2013	ISSCC	A scalable 0.128-to-1Tb/s 0.8-to-2.6pJ/b 64	M. Mansuri	32	16	2.58	41.28	18	
2009	VLSI	A 12-Gb/s transceiver in 32-nm bulk CMOS	Sopan Joshi	32	12	3.15	37.8	10	
2010	ISSCC	A 78mW 11.8Gb/s serial link transceiver wit	F. Spagna	32	11.8	6.61	78	25	
2011	ASSCC	Extending HyperTransport? technology to 8	B. A. doyle	32	8	11.81	94.44		
2012	ISSCC	A 28Gb/s 4-tap FFE/15-tap DFE serial link t	J. Bulzacchelli	32	28	14.04	393	29	
2017	VLSI	A 28.05Gb/s Transceiver using Quarter-Rat	Gautam Gangasan	32	28.05	17.25	484	48	14.03
2015	VLSI	A 4x9 Gb/s 1 pJ/b NRZ/Multi-Tone Serial-Di	arash Gharibdoust	40	9	1	9		
2015	CICC	A 5.4-mW 4-Gb/s 5-Band QPSK Transceiver	Wei-Han Cho	40	4	1.35	5.4		
2010	ISSCC	A 47 #x00D7 10Gb/s 1.4mW/(Gb/s) parallel	F. O'Mahony	40	10	1.4	14	8	
2011	VLSI	A 5.6Gb/s 2.4mW/Gb/s bidirectional link wit	J. Zerbe	40	5.6	2.4	13.4		
2009	VLSI	A 4.3GB/s mobile memory interface with po	R. Palmer	40	4.3	3.3	14.19	10	
2015	VLSI	A 100-GbE Reverse Gearbox IC in 40nm C	Taehun Yoon	40	25	3.92	98		
2010	VLSI	A 40nm 7Gb/s/pin single-ended transceiver	Seung-Jun Bae	40	7	3.94	27.58		
2015	CICC	A power-and-area efficient 10 x 10 Gb/s bo	Joon-Yeong Lee	40	10	4.16	41.6		
2014	ISSCC	A pin-and power-efficient low-latency 8-to-1	A. Singh	40	12	4.29	51.5	15	
2018	ISSCC	A 20Gb/s Transceiver with Framed-Pulsewi	Sejun Jeon	40	20	4.53	90.6	12	7.5
2011	CICC	A 19 mW/lane Serdes transceiver for SFI-5.	S. Fallahi	40	3.125	6.08	19		
2010	ISSCC	A 1.296-to-5.184Gb/s Transceiver with 2.4n	K. Maruko	40	5.184	6.21	32.2		
2011	ISSCC	An 8.4mW/Gb/s 4-lane 48Gb/s multi-standa	M. Ramezani	40	8	9.13	73	22	
2017	ISSCC	A 56Gb/s PAM-4/NRZ transceiver in 40nm	Pen-Jui Peng	40	56	10.39	582	24	14
2015	VLSI	56Gb/s PAM4 and NRZ SerDes Transceiver	Jri-Lee	40	56	11.96	670		
2015	VLSI	56Gb/s PAM4 and NRZ SerDes Transceiver	Jri-Lee	40	56	12.6	710		
2009	ASSCC	A 5Gb/s low-power PCI express/USB3.0 ready	PHY in 40nm C	40	5	25	125		

Year	Publication	Title	First Author	Process	[Gbps]	[pJ/bit]	[mW]	Loss (dB)	Loss Frequency (GHz)
2014	ISSCC	A 780mW 4x28Gb/s transceiver for 100GbE	U. Singh	40	28	27.86	780	20	
2010	CICC	A 2.5-8Gb/s transceiver with 5-tap DFE and	Wei-Cheh Chen	40	8	27.91	223.25		
2013	ISSCC	A sub-2W 39.8-to-44.6Gb/s transmitter and	B. Raghavan	40	44.6	43.05	1920	20.4	
2010	ISSCC	A 5-to-25Gb/s 1.6-to-3.8mW/(Gb/s) reconfig	G. Balamurugan	45	10	3.8	38		
2011	VLSI	An 8x10-Gb/s Source-Synchronous I/O Sys	Timothy O. Dickson	45	9	6.1	54.9	16.3	
2011	CICC	A 16-Gb/s backplane transceiver with 12-tap	G. R. Gangasani	45	16	24.06	385	32	
2009	VLSI	Self-calibrating transceiver for source synch	Young-Chan Jang	60	1.6	10.99	17.5875		
2017	ISSCC	A time-based receiver with 2-tap DFE for a	Il-Min Yi	65	12	0.46	5.52	13	6
2015	ISSCC	A 0.45-to-0.7V 1-to-6Gb/s 0.29-to-0.58pJ/b	Woo-Seok Choi	65	6	0.58	3.48		
2010	ISSCC	A 12.3mW 12.5Gb/s complete transceiver in	K. Fukuda	65	12.5	0.98	12.3	12	
2016	VLSI	A 32 Gb/s Rx Only Equalization Transceiver	Sewook Hwang	65	32	1.44	46.1	22	
2018	ISSCC	A 0.5-to-0.9V, 3-to-16Gb/s, 1.6-to-3.1pJ/b	Vshwin Ramachandra	65	10	1.8	18	27	5
2009	ISSCC	A 10Gb/s compact low-power serial I/O with	Liu Yong	65	8.9	1.91	17	19	
2018	ISSCC	A 7.8Gb/s/pin 1.96pJ/b compact single-end	Soeun Lee	65	7.8	1.961	15.3	4.98	
2011	ISSCC	An 8.4Gb/s 2.5pJ/b mobile memory I/O inte	Byun, Gyung-Su	65	8.4	2.5	21		
2017	CICC	A 4-40 Gb/s PAM4 Transmitter with Output	Xuqiang Zheng	65	40	2.55	102	6	
2016	VLSI	A Low-EMI Four-Bit Four-Wire Single-End	Il-Min Yi	65	6.4	2.67	17.088		
2015	VLSI	A 2.8mW/Gb/s 14Gb/s Serial Link Transceiv	Saurabh Saxena	65	14	2.8	39.2	12	
2018	ISSCC	A 32Gb/s 133mW PAM-4 Transceiver with I	Liangxiao Tang	65	32	4.16	133	23	8
2009	VLSI	A 21-Gb/s 87-mW transceiver with FFE/DFE	Wang, Huaide	65	20	4.35	87	21.2	
2017	ISSCC	A 16Gb/s 3.6pJ/b wireline transceiver with	shwin Ramachandra	65	16	4.37	69.9	27	8
2015	CICC	A 190mW 40Gbps SerDes Transmitter and	Ke Huang	65	40	4.75	190	15	
2017	ISSCC	A 60Gb/s 288mW NRZ transceiver with ada	Jaeduk Han	65	60	4.8	288	21	30
2015	ISSCC	1V 10Gb/s/pin single-ended transceiver wit	J. Song	65	10	4.9	49		
2007	VLSI	A Scalable 5-15Gbps, 14-75mW Low Powe	G. Balamurugan	65	15	5	75		
2017	CICC	Channel Adaptive ADC and TDC for 28 Gb/	Aurangozeb	65	28	5.71	160	30	7
2017	ISSCC	A 3-to-10Gb/s 5.75pJ/b transceiver with flex	resh Kumar Nandw	65	10	5.75	57.5		
2016	ISSCC	A 16Mb/s-to-8Gb/s 14.1-to-5.9pJ/b Source	Guanghua Shu	65	8	5.85	46.8		
2015	ISSCC	A 7Gb/s rapid on/off embedded-clock serial	T. Anand	65	7	9.1	63.7		
2011	ISSCC	A 40Gb/s TX and RX chip set in 65nm CMC	Chen, Ming-Shuan	65	40	11.43	457	18	
2012	CICC	An 8.5711.5Gbps SONET transceiver with r	Kocaman, N	65	11.5	12.26	141		
2011	ISSCC	11.3Gb/s CMOS SONET-compliant transce	N. Kocaman	65	11.3	18.94	214		
2009	ASSCC	A low latency transceiver macro with robust	Zhang Feng	65	3	30.33	91		
2010	ISSCC	A 5Gb/s transceiver with an ADC-based fee	H. Yamaguchi	65	5	56	280	15	
2009	ISSCC	A 500mW digitally calibrated AFE in 65nm C	Jun Cao	65	10	120	1200	26	
2009	ISSCC	A 40Gb/s multi-data-rate CMOS transceiver	Y. Amamiya	65	40	140	5600		
2006	ESSCIRC	A 6-Gbps/pin 4.2mW/pin Half-Deplex Pseu	Sua Kim	80	6	4.2	25.2		
2007	ASSCC	A PCI-express Gen2 transceiver with adapt	Tse-Hsien Yeh	80	5	23.3	116.5	20	
2011	ISSCC	A highly digital 0.5-to-4Gb/s 1.9mW/Gb/s se	R. Inti	90	4	1.8	7.2		
2007	ISSCC	A 14mW 6.25Gb/s Transceiver in 90nm CM	R. Palmer	90	6.25	2.24	14	15	
2005	CICC	A 0.8-1.3V 16-channel 2.5Gb/s high-speed	Y. Doi	90	2.5	9.2	23	10	
2008	ISSCC	A 20Gb/s Duobinary Transceiver in 90nm C	Jri Lee	90	20	9.75	195		
2006	ISSCC	A 100mW 9.6Gb/s Transceiver in 90nm CM	E. Prete	90	9.6	10.42	100		
2006	ISSCC	A 20Gb/s Forwarded Clock Transceiver in 9	B. casper	90	20	11.85	237	16	
2012	CICC	FEC-based 4 Gb/s backplane transceiver in	Adam C. Faus	90	4	12.5	50	18.2	
2006	CICC	A Multi-Standard Low Power 1.5-3.125 Gb/s	A. Yokoyama-Mart	90	3.125	15.58	48.7		
2006	ISSCC	A 20Gb/s Embedded Clock Transceiver in 9	B. casper	90	20	15.9	318	15	
2006	VLSI	Power/Performance/Channel Length Tradeo	Evelina Yeung	90	9.6	17.68	169.7	9.4	
2005	ISSCC	12Gb/s duobinary signaling with 72 oversam	K. Yamaguchi	90	12	19.17	230		
2008	VLSI	A 21-channel 8Gb/s transceiver macro with	A. Hayashi	90	8	20	160		
2004	VLSI	A 1-4 Gbps quad transceiver cell using PLL	Y. Frans	90	3.125	23.36	73		
2010	CICC	A 16 Gb/s four-wire CDMA-based high speed	Tzu-Chien Hsueu	90	16	24	384	5	
2007	ISSCC	A 250mW Full-Rate 10Gb/s Transceiver Co	T. Masuda	90	10	25	250		
2009	ISSCC	A 4-channel 10.3Gb/s backplane transceiv	Y. Hidaka	90	10.3125	25.21	260	35.8	
2009	CICC	A 373.8Gb/s four-wire high speed I/O link b	Tzu-Chien Hsueh	90	11.4	27	307.8	10	
2011	ISSCC	A 4-channel 10.3Gb/s transceiver with adap	Y. Hidaka	90	12.5	27.84	348	34.9	
2008	ISSCC	An 8Gb/s Transceiver with 3x-Oversampling	K. Fukuda	90	8	29	232	36.8	
2006	ISSCC	A 10Gb/s 5-Tap-DFE/4-Tap-FFE Transceiv	M. Meghelli	90	10	30	300	24	
2012	ISSCC	An 8Gb/s quad-skew-cancelling parallel tra	Kim, Young-Sik	90	8	32.13	257		
2008	ASSCC	An ASIC-Ready 1.2576.25Gb/s SerDes in 9	Y Nishi	90	6.25	36.8	230		
2007	ASSCC	A quad 1210Gb/s serial transceiver in 90nm	Han Bi	90	10	98	980		
2006	ISSCC	A 20Gb/s Bidirectional Transceiver Using a	Y. Tomita	110	20	13.15	263	5	
2004	VLSI	5-6.4 Gbps 12 channel transceiver with pre	H. Higashi	110	6.4	38.67	247.5		

Year	Publication	Title	First Author	Process	[Gbps]	[pJ/bit]	[mW]	Loss (dB)	Loss Frequency (GHz)
2003	ISSCC	A CMOS multi-channel 10Gb/s transceiver	H. Takauchi	110	10	41.5	415		
2004	VLSI	A 4-channel 3.125Gb/s/ch CMOS transceiver	Weixin Gai	110	3.125	80	250	30	
2008	ISSCC	A 10Gb/s IEEE 802.3an-Compliant Ethernet	S. Gupta	130	10	1.05	10.5		
2011	ISSCC	A 4.8Gb/s impedance-matched bidirectional	Shin, Woo-Yeol	130	4.8	14.24	68.352		
2005	VLSI	A quad 3.125 Gbps transceiver cell with all-	Bong-Joon Lee	130	3.125	15.68	49		
2003	VLSI	0.622-8.0 Gbps 150 mW serial IO macrocell	R. Farjad-Rad	130	8	18.75	150		
2005	ISSCC	A 0.6 to 9.6Gb/s binary backplane transceiver	K. Krishna	130	6.25	24	150		
2005	ASSCC	A Sub-1V CMOS 2.5Gb/s Serial Link Transceiver	Chih-Chien Hung	130	2.5	28	70		
2002	VLSI	A 0.4-4 Gb/s CMOS quad transceiver cell using	K. Y. K. Chang	130	4	32	128		
2004	ISSCC	An 8Gb/s source-synchronous I/O link with a	J E Jaussi	130	8	35	280		
2004	ISSCC	An 800 mW 10 Gb Ethernet transceiver in 0.13	S. Sidiropoulos	130	22.8	35.09	800		
2004	ISSCC	A fully integrated 0.13 μ m CMOS 10 Gb Ethernet	Hyung-Rok Lee	130	22.5	39.91	898		
2004	ESSCIRC	A 2.5 Gbps - 3.125 Gbps multi-core serial-link	T. Geurts	130	3.125	44.8	140		
2005	ISSCC	A 6.4Gb/s CMOS SerDes core with feedforward	M. Sorna	130	6.25	46.4	290	32	
2003	CICC	A low-power 0.13 μ m CMOS OC-48 SONET	R. Wadhwa	130	3.125	51.2	160		
2002	VLSI	A 0.13 μ m CMOS 5-Gb/s 10-meter 2	Y. Kudoh	130	5	60	300	9	
2006	ISSCC	A Quad 6Gb/s Multi-rate CMOS Transceiver	Yongsam Moon	130	6	64.33	386	22	
2003	ESSCIRC	A serial 10 gigabit Ethernet transceiver on a	Bin Wu	130	22.5	66.67	1500		
2006	ISSCC	A 9.95 to 11.1Gb/s XFP transceiver in 0.13 μ m	J. Kenney	130	11.1	71.44	793		
2002	ISSCC	A 10Gb/s transceiver with GND-VDD compatibility	K. Tanaka	130	5	75.6	378		
2001	ISSCC	A 2Gb/s 21CH Low-Latency Transceiver Cell	Tanahashi, T.	130	2	78.75	157.5		
2008	VLSI	A 40-Gb/s transceiver in 0.13- μ m CMOS technology	Jeong-Kyoum Kim	130	39.5	91.14	3600		
2004	ISSCC	A 10 Gb/s SONET-compliant CMOS transceiver	H. Werker	130	10.7	91.59	980		
2003	CICC	A fully-integrated 10.5 to 13.5 Gbps transceiver	G. Miao	130	10.88	91.91	1000		
2006	ESSCIRC	8Gbps Parallel Data Transmission with Adaptive	E. Takahashi	130	2	125	250		
2006	ISSCC	A 12.5Gb/s Single-Chip Transceiver for UTF-8	M. Callicotte	130	12.5	304	3800		
2005	ISSCC	A 5Gb/s NRZ transceiver with adaptive equalization	N. Krishnapura	130	5.15	407.77	2100	9.9	
2004	CICC	A 2.5-3.125 Gb/s quad transceiver with sector	A. L. Coban	150	3.125	20.8	65		
2003	VLSI	A 27-mW 3.6-Gb/s I/O transceiver	K. L. J. Wong	180	3.6	7.5	27		
2003	ISSCC	8Gb/s differential simultaneous bidirectional	A. Martin	180	8	15	120		
2006	VLSI	A Low Power 4.2Gb/s/pin Parallel Link Using	Zogopoulos	180	4.2	17.1	71.82		
2008	CICC	A 5-Gb/s/pin transceiver for DDR memory interface	Kwang-II Oh	180	5	20.4	102		
2002	ISSCC	A quad 3.125Gbps/channel transceiver with	Dong Deng	180	3.125	27.2	85		
2003	ISSCC	A 2.7 Gb/s CDMA-interconnect transceiver	Zhiwei Xu	180	2.7	27.41	74		
2003	ISSCC	A 2.5-10Gb/s CMOS transceiver with alternative	Bong-Joon Lee	180	10	54	540		
2004	ISSCC	A quad-channel 3.125Gb/s/ch serial-link transceiver	Jeongsik Yang	180	3.125	57.6	180		
2008	ASSCC	A 3.2-Gb/s transceiver with a quarter-rate line	Kyung-Soo Ha	180	3.2	70.31	225		
2008	ASSCC	A 8 GByte/s transceiver with current-balancing	Seon-Kyoon Lee	180	8	115.07	920.52		
2002	CICC	A 10Gbase Ethernet transceiver (LAN PHY)	Yoshimura	180	22.5	128.89	2900		
2001	ISSCC	Fully-integrated SONET OC48 transceiver in	Momtaz, A	180	2.666	187.55	500		
2005	ISSCC	A 3Gb/s 8b single-ended transceiver for 4-d	Seung-Jun Bae	250	3	50.75	152.25		
2001	ISSCC	A 0.6-2.5 GBaud CMOS tracked 3/spl times	Yongsam Moon	250	2.5	107.6	269		

A.3 Wireline Interface Standard Survey

The following pages contain tables that show our information collected on wireline interface standards.

Standard	Application	Max. Lane rate in Gbps	Max. BER	Top level		Channel/profile	Functions			Data	Signal (nominal)			Subsystem specs		Reference/Comment
				Signal Power Efficiency	Clock		Coding functions	EQ functions	Other functions		Common mode in V	Differential swing in V	Min. rise/fall time in ps	Other	Subsystem specs	
LOW-LEVEL	LVDs	Backplane, cables	0.65	No spec	0.96154	No spec	No spec	No spec	No spec	No spec	1.2	0.25	260	Dynamic offset	TX + RX currents	Standard worries about EMI, limiting datarate. 0.65 Gbps is with 500 ps tr degradation.
	CML	SerDes, High-speed point-to-point, HDMI	10	No spec	0.2025	No spec	No spec	No spec	No spec	No spec	2	0.45	50	No spec	No spec	More power draw, worse input cm range, higher swing and not well standardized! But towards the target speed though.
LVPECL	CML is not well standardized, took it from hotlink spec															
	Clock distribution on PCB, Ethernet	10	No spec	1.28	No spec	No spec	No spec	No spec	No spec	No spec	2	0.8	150	No spec	No spec	Comparison of LVDS, CML, ECL. Good read. https://www.cypri.lvds-vs.-hollink-ii-cml-standard
																https://www.ti.com/lit/zt/tps62a01-q1
																Very nice comparison of CML, PECL. https://e2e.ti.com/forums/processors/1182111/pecl-lvpecl-explanation-and-drivers-by-ti
ETHERNET																https://ieee802.3.org
CAUI-4 (after XLAUI/CAUI-10, by IEEE)	Ethernet small distance	25	1.00E-15	0.32400	Recovered	Specified	Yes	FFE, CTLE adaptive	Many on protocols (ethernet)	Encoding	AC	0.9	Quite some 10 options	COM		General 2015 Ethernet standard.
OIF CEI-56G (XSR-NRZ)	Optical electronics interfacing, chip-to-chip	58	1.00E-15	0.02759	No spec	Specified	No spec	No spec	No spec	No spec	No spec	0.4	4 Basics	COM		https://ieee802.3.org
																Amendment to Ethernet describing CAUI-4.
OIF CEI-56G (VSR-PAM4)	Optical electronics interfacing															General 2015 Ethernet standard (see ch. 93). NRZ for extra short range (PCB). But, awesome example of PAM4 in the industry, they set the start of other specs in difficult media.
MID-LEVEL	Smartphone card display (CSI/DSI)	1.5	1.00E-12	0.26667	Forwarded	Specified	8b9b opt.	No spec	Multiple lanes, protocols etc.	AC encoding opt.	0.2	0.2	100	https://www.xilinx.com	Little	High-level standard, very low power due to low swing, feeds line LVDS but at lower cm input.
																https://www.xilinx.com
	Same as D-PHY	2.5	1.00E-12	PHY?	Embedded	Specified	Mapper for 3-wire	No spec	Weird signalling, CDR	3-wire	X	D-PHY ish	?	3-wire protocol	D-PHY ish	Another Xilinx D-PHY solution guide. https://files.xilinx.com
																Official 2011 spec incl circuit
C-PHY																General MIPI tech brief. https://mipi.org
M-PHY	Very high-speed mobile chip-to-chip interfaces	11.6	1.00E-10	0.05368	Forwarded	Specified	8b10b	Adaptive equalizer in HS gear 4	Protocol, state machine	8b10b	0.16	0.25	9	Slew, jitter	Specs on jitter	MPI's high-end standard for mobile. N.B. that like most MIPI specs, this link is suited for very low supply voltages, resulting in the low power operation.
																https://www2.digic.com
	Processor to chip interfaces on PCB	2.8	1.00E-12	2.892857143	Forwarded	Specified	Some, not clear	DE	DDR, many layers, multiple lanes	8b10b optional, scrambling	0.6	0.9	30	Basics	Specs on jitter	M-PHY v2 description (current is v4), incl. EMI mitigation techniques.
																M-PHY v2 specs VERY extensive, can learn a lot about the electrical specs. https://ieee802.3.org
Kandou Bus	Chip-to-chip, in-package transmission	16	1.00E-15													M-PHY equalizer: details on jitter specs. https://ieee802.3.org
																M-PHY driver in 65 nm (2012).
HIGH-LEVEL																A big load of other specifications, quite low rate for this high swing. Probably conservative because of important data.
JESD204b	Chip-to-chip dataconverters	12.5	1.00E-15	0.10368	Forwarded	Specified	8b10b	PE, EQ	Sync, error correction	8b10b scrambling	0.8	0.36	24	Eye mask, more		https://www.ti.com
																JESD204b overview by TI. https://www.ti.com
																Physical interface description by TI. https://ieee802.3.org
																JESD204b driver paper. https://ieee802.3.org
																Even better JESD204b driver paper.

Standard	Application	Max. Lane rate in Gbps	Top level			Functions			Data	Signal (nominal)			Subsystem specs	Reference/Comment
			Max. BER	Signal Power Efficiency	Clock	Channel/profile	Coding functions	EQ functions		Common mode in V	Differential swing in V	Min. rise/fall time in ps		
PCIe (v4)	Chip to chip between cards	2	1.00E-12	3.2	Forwarded	Specified	128b/130b + scrambling	FIR, DFE	128b/130b + scrambling	AC	0.8	20	Many specs (incl. logic)	https://en.wikipedia Encoding consumes overhead, net = 2 Gbps.. and has special connectors.
														https://pci-sig.org Slideset on electrical basics of PCIe 4.
														https://www.youtube Video on electrical basics of PCIe.
RapidIO (v1 IV)	HPC equipment	25	1.00E-15	CAUI-4	CAUI-4	CAUI-4	Embedded	FFE, EQ		AC	1.2	CAUI-4	CAUI-4	Official 3.0 base specs http://www.rapidio Official 3.0 base specs Very big open standard, but on PHY-level is just chip-to-chip ethernet.
														https://ieee.org The electrical specifications for the short-reach 25-78125 Gbaud PHY shall be the same as those specified in Annex 63D.3 of the IEEE Standard 802.3bn-2015[3] for CAUI-4.
Infiniband (EDR) JESD204c	Chip to chip, HPC	25	1.00E-12	0.064	ODR	Specified	64b/66b	FIR, DFE	A lot of overhead	AC	0.4	13	Eye mask, more	https://ieee.org General 2015 Ethernet standard (see ch. 93). https://www.infini Hard to find, but 2016 specs. It refers to Ethernet, but eye mask does not correspond.
	Chip-to-chip dataconverters	32	The	0.2	Forwarded	Specified	Efficient coding	PE, EQ	Efficient coding, scrambling	AC	0.8	?	COM	<i>Comment</i> Physical level is probably equal to JESD204b. This standard specifies more strict channel requirements, allowing for higher data rate.
														Assume CEI-28G-SR
														Article by AD on the differences compared to JESD204b. http://www.ll.com Article by TI on the differences compared to JESD204b.

B Wireline Transmitter Auxiliary Circuit Designs

In this appendix, we provide more details on auxiliary circuit design for the wireline transmitter chip presented in this thesis. Some parts of this appendix already appeared in the main text.

B.1 Aligned Divided Clock Generation

This clocking system is designed with the three functional requirements below.

- The clocking subsystem takes a complementary, rail-to-rail 10 GHz nominal clock signal (i.e. each 1 or 0 is 50 ps) from a single-ended to balanced transformer¹.
- The system must output four divided, complementary clocks, namely 1/2, 1/4, 1/8 and 1/16, while buffering the 10 GHz clock.
- The system must ensure alignment of the low-to-high transitions of all clocks to the rising edge of the the slowest clock. That is, the alignment of $CK/16$ down to $CK/2$ is critical. This is a special, critical requirement to guarantee a predictable serialization behavior by the design multiplexer.

Shown in Fig. B.1 is an illustration of the desired signals.

Quadrature clock dividers² (based on C2MOS latches) provide the clock division in this clock generation block. The dividers incur some delay, so the three remaining design questions are:

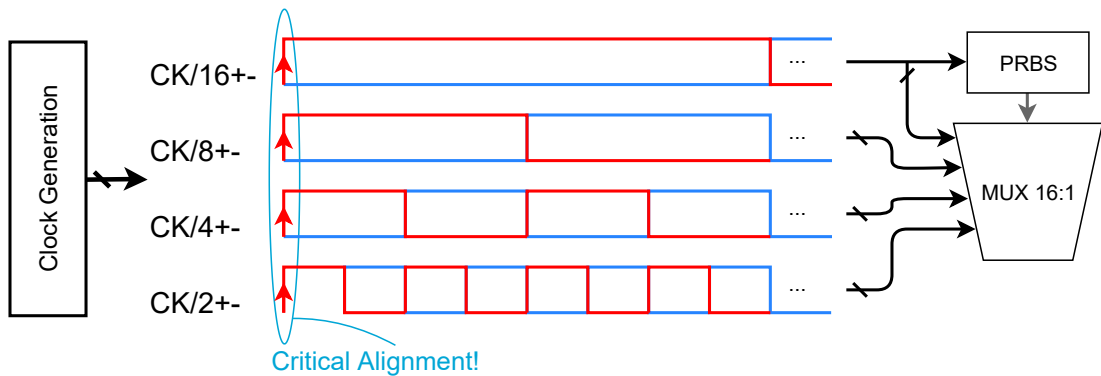


Figure B.1: Required clock signals for the multiplexer.

¹This transformer is kindly provided by PhD student Mohammadreza Beikmirza.

²Unit divider cells were kindly provided by PhD student Mohammadreza Beikmirza. The design and overall layout of the presented clock solution was the author's (Jun Feng) own work.

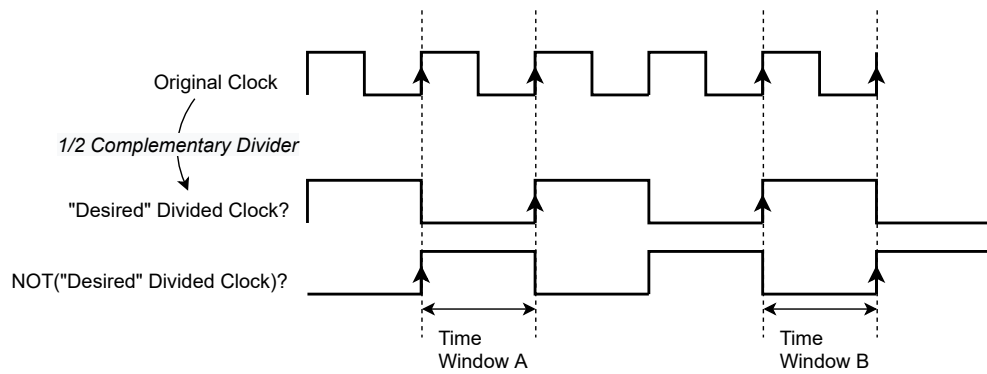


Figure B.2: Down stream positive selection is ambiguous.

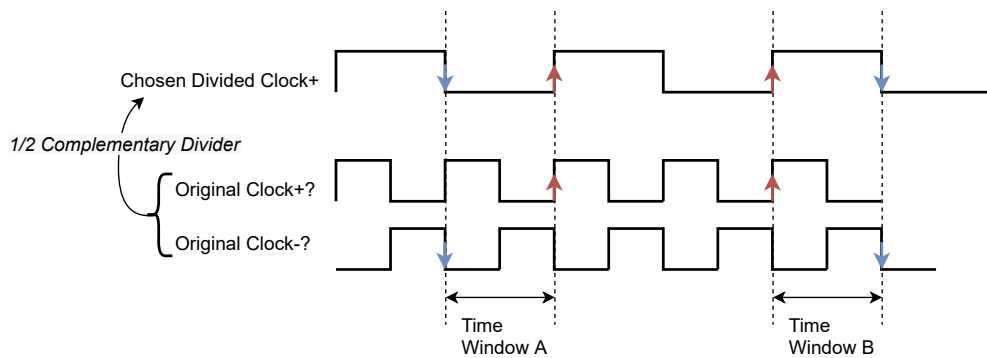


Figure B.3: Up stream positive selection is necessary.

1. How can the correct "positive" clock signals out of the generated complementary/quadrature signals be determined? E.g. which is $CK/8+$?
2. How do we realize the wanted alignment of output clocks?
3. What are important timing constraints for retiming of clock signals?

Positive Clock Selection The dividers are rising edge triggered. Looking at Fig. B.2, it becomes clear that from the perspective of the undivided clock (CK), it is impossible to distinguish $CK/2+$ from $CK/2-$. Fundamentally, this is because it takes a full phase of CK before both of the divided clocks swap values.

On the other hand, from the point of view of the $CK/2$ signals, there is a correct $CK+$ and $CK-$ to be defined³! As Fig. B.3 illustrates, it is most logical to choose $CK+$ to be the input clock signal that triggers a change in both $CK/2+$ and $CK/2-$ with its rising edge. In short, naming the positive phases should be done starting from point of view of the slowest clock in the system.

Moreover, an inspection of the proposed quadrature divider shows that, deterministically, one pair of the divided outputs ($CK/2+$ and $CK/2-$) reacts to the rising edge of $CK+$, while the other pair ($CK/2_{90}$ and $CK/2_{270}$) reacts to the falling edge of $CK+$.

A deterministic reaction of the output clocks to the input clock edges is critical, so we can indeed conclude that this divider can be used to create an arrangement of divided clocks for the MUX with correctly chosen positive clock phases.

³We must recognize that the input of our divider is complementary too.

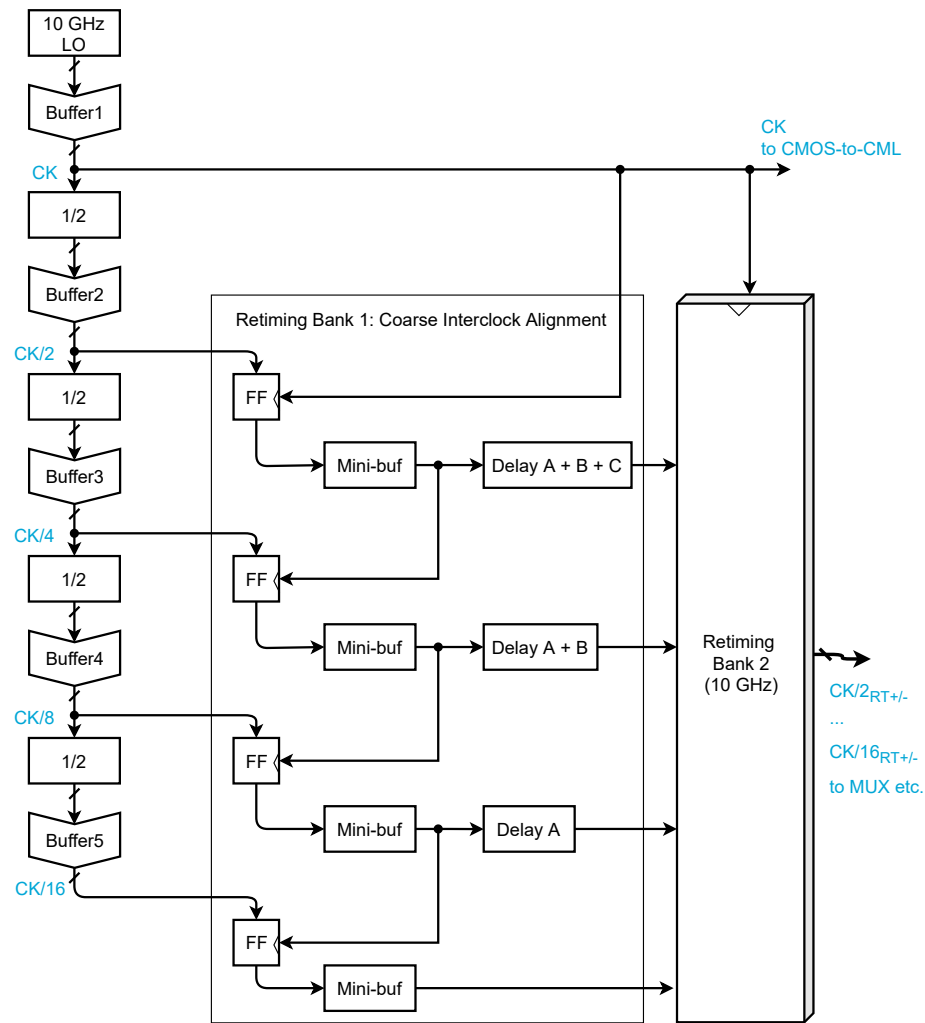


Figure B.4: Overview of designed clock generation architecture.

Self-Retimed Clock Division Shown in Fig. B.4 is an overview of the self-retimed divided clock generator. The design reasoning is as follows:

1. The dividers can form a chain and the correct positive phases will be identified starting from the slowest clock $CK/16+$.
2. The divider chain has propagation delays in reality (more than 15 ps per stage). All outputs should be retimed again with one “retimer bank” at 10 GHz.
3. The divider chain delays are actually quite severe because the delays add up for each division. There must be some coarse *interclock* alignment first to compensate for these different delays before the final retiming bank.

The interclock alignment architecture consists of three components:

- Retiming flip-flops after buffering the divided outputs with the original undivided clock: to reduce the accumulated delays.
- Mini-buffers after the flip-flops to generate enough drive strength after retiming, but invoke minimum extra delay in the flip-flop to flip-flop path.

- Delays lines of different lengths to compensate for the remaining delays (t_{CQ} of the flip-flops and mini-buffers) in the retiming chain.

At the output of the delay lines, the clock signals have the desired phase relationships with less than 100 ps of error. Then, the final retiming bank can successfully remove any skew for the final edge alignments.

Timing Constraint Consideration We must address the possible concern that this scheme fails timing constraints at the interface between the clock divider chain and the first retiming stage. Namely, the divider chain incurs significant delays due to the dividers and buffers while the clocks at the flip-flops do not. As such, it seems that the data path arrives too late at these flip-flops.

However, all "data" and clocks of the flip-flops in the proposed divider are *periodic* (repetitive). Consequently, all "data" values are sampled eventually. The resulting clocks remain aligned as intended and from the system's point of view, all requirements are met.

B.2 Timing Delay Trimmers

The main issue regarding the clock divider architecture is that it does not guarantee to work at different input clock frequencies. With a different input clock, the time period between retiming clock edges changes, while the delays due to the divider chain, buffers and clock-to-Q remain unchanged. Input clock frequencies exist that result in violation of setup and hold conditions.

The most evident location where this problem occurs is at the second retiming bank. Suppose that everything works well at 10 Gbps and that we choose $CK+$ such that we have good hold time margin at the second retiming bank, but a weak setup time margin. If the input clock increases to 11 GHz, the clock period lowers by 10 ps. It becomes risky to properly sample $CK/2$ now! This could trigger metastability of the second retiming stage.

After submission of the preliminary full chip layout, we decided that we wanted to implement some variable delay in locations critical to the timing-related functionality of the system. These trimming features were designed in very short time frame for the following critical timing-sensitive parts of the system:

1. For the latch in the CMOS-to-CML interface (skew-less complementary signal generation) of the driver.
2. For the second retiming bank of the clock generation circuit.

Fig. B.5 shows the context of both delay trimmers.

CMOS-to-CML Trim For the CMOS-to-CML latch, a two-bit variable delay line with 4 steps of 30 ps in the data path at the output of the CMOS multiplexer is implemented as shown in Fig. B.6. Thus, the total delay range is 120 ps to be able to cover for low operational frequencies too.

The addition of variable delays in the clock path might be a more obvious solution. However, the MUX output data is not complementary and has a transition rate twice as slow as the clock. Therefore, the data path implementation is relatively small in area and consumes less power than a clock path solution.

Retimer Bank Trim There was only one I/O signal left in the pad ring. Therefore, a *controllable* inversion in the nominal 10 GHz clock signal path in the second retiming bank implements a configurable 180° phase shift. Just like a delay, a phase shift yields additional timing budget.

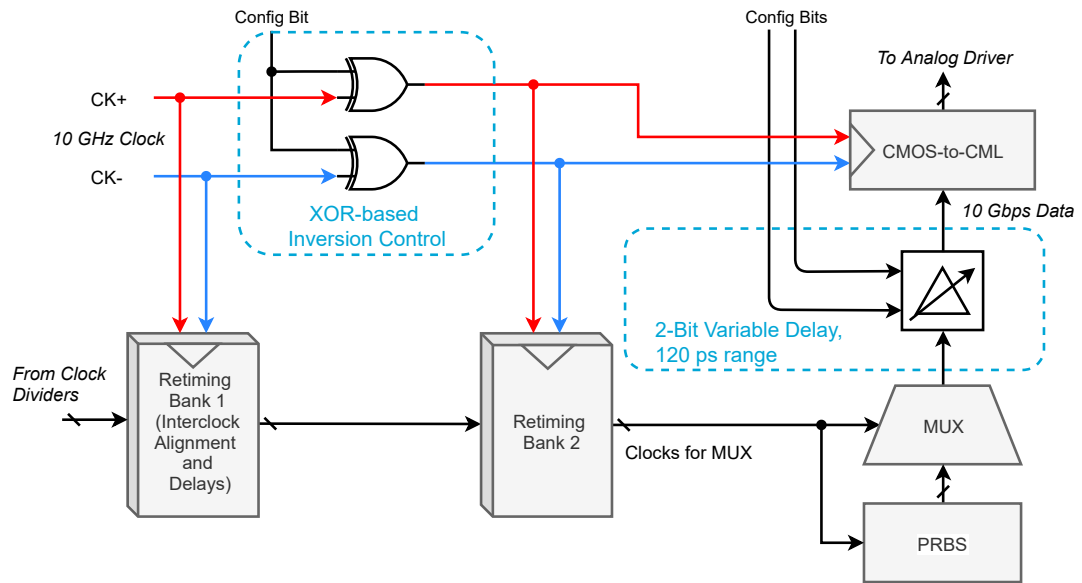


Figure B.5: Overview of delay trimmers in the transmitter chip.

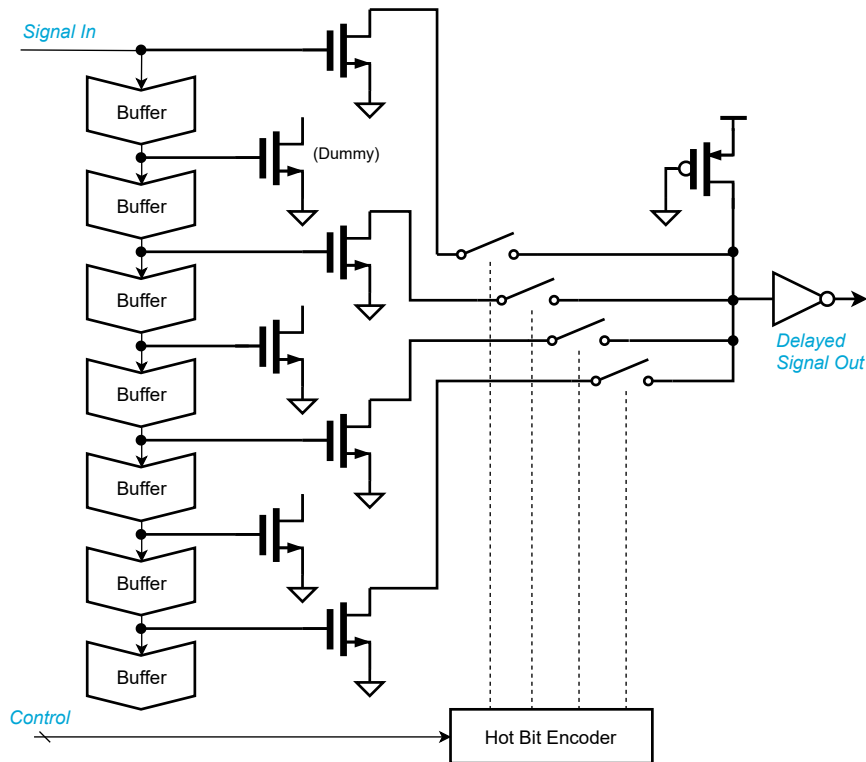


Figure B.6: Implementation of 2-bit variable delay element.

A custom four-transistor XOR circuit is utilized as shown in Fig. B.7. The standard library XORs are too slow for 10+ Gbps operation and the inversion operation has to be with minimum delay compared to the non-inverting operation. We must recognize for this solution that the low-speed control signals should act like a good V_{DD} and ground when the inverter is active. Terminal A must be used for the high-speed signal.




Figure B.7: Four-transistor XOR.

B.3 Pseudo-Random Bitstream Generator

The PRBS generator is critical but quite straight forward. The inclusion of a PRBS generator avoids the need for on-chip storage of transmit data for this chip. A basic PRBS generator consists of a string of cascaded flip-flops with a XOR that feeds some of the intermediate nodes back to the input [6]. In our proposed transmitter design, we incorporate multiple selectable feedback options for a PRBS generator that can generate a random bit stream that is non-periodic up to the $2^{16} - 1$ th bit as shown in Fig. B.8. As we need a 16-bit word for the multiplexer, the PRBS solution remains compact as we feed all nodes in the loop to the multiplexer as opposed to using only the output node (*D15*).

We use the following Matlab code to simulate the PRBS sequence:

```

1 startSeed = [1 zeros(1,15)];
2 statesLFSR = zeros(2^16,16);
3 statesLFSR(1,:) = startSeed;
4 for period = 1:2^16 %compute next state, known feedback seq is 16, 14, 13, 11
    (XOR is mod2 addition)
5     newBit = mod(statesLFSR(period,16) + statesLFSR(period,14) + ...
        statesLFSR(period,13) + statesLFSR(period,11), 2);
6     newState = [newBit statesLFSR(period,1:15)];
7     %if same as startSeed, stop and report %if different, save and continue to
        next period
8     if newState == startSeed
9         period
10        break;
11    else
12        statesLFSR(period + 1,:) = newState;
13    end
14 end
15
16 %% Serialize Data (if ever needed)serialPRBS =
    reshape(statesLFSR(1:period,:),[],1);

```

B.4 Current Reference

The CML tails only need high-side current sources to perform reference voltages. Therefore, we designed for a current source external to IC. Near the IC's pad, we copy the current with a NMOS

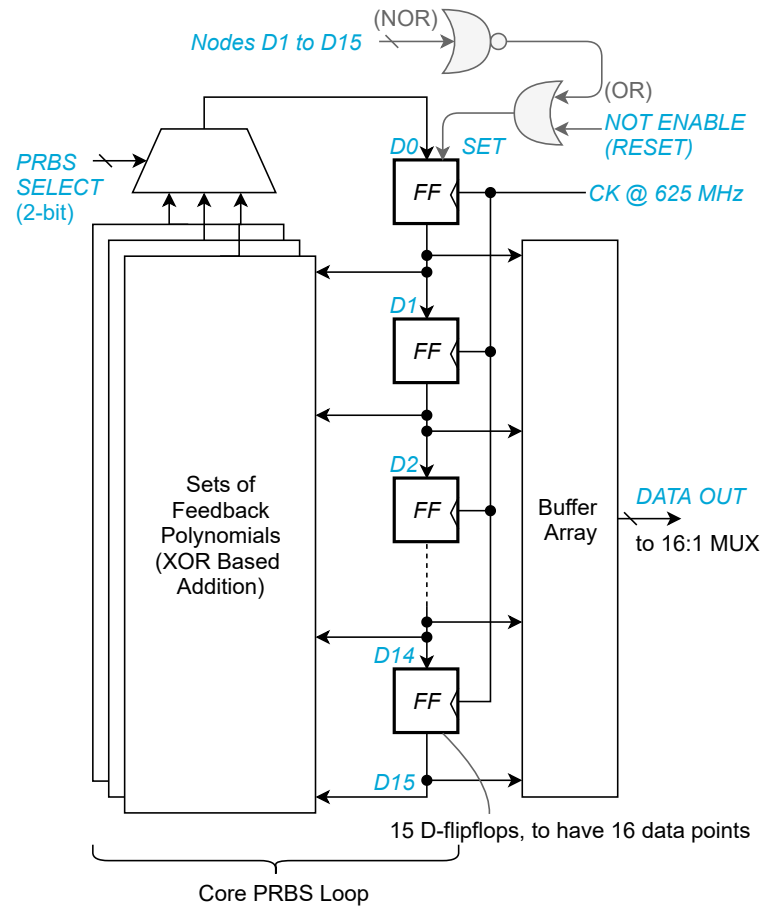


Figure B.8: Overview of 16-bit PRBS generator in this work. The XOR feedback determines the periodicity of the sequence while the NOR feedback resets the loop if all D1 – 15 become zero in an unfortunate case (this would halt generation of new pseudo random bits).

mirror. In current domain, we distribute the copied currents to local PMOS mirrors to act as the necessary high-side sources (as suggested by [10] to minimize return path resistance). An input resistance of $3\text{ k}\Omega$ ensures a pad voltage 1 V at $100\text{ }\mu\text{A}$ input. The result is illustrated by Fig. B.9.

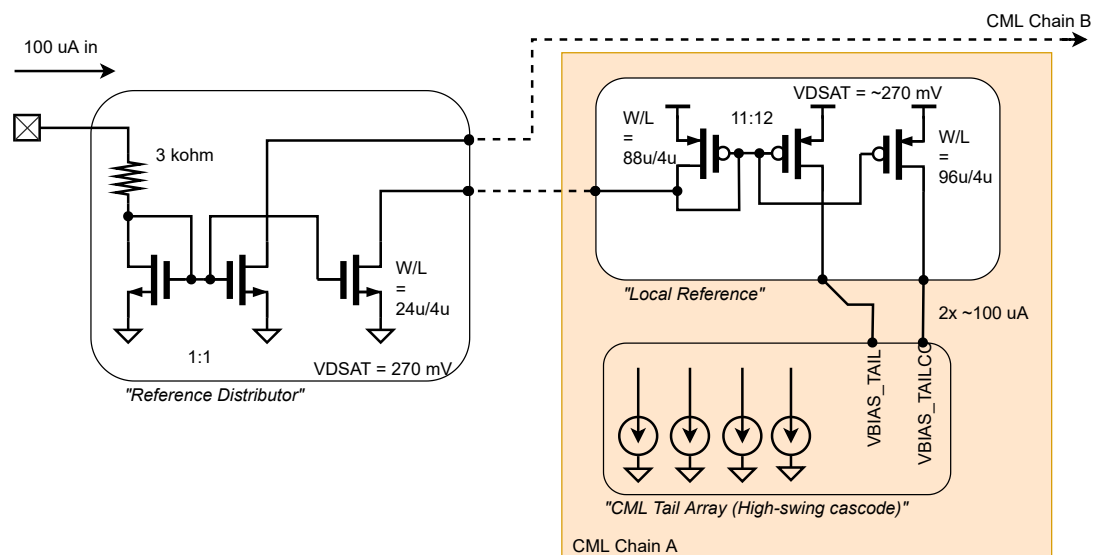


Figure B.9: Overview of current reference distribution.

Journal of Management Education 36(7) 809–824

Table C.1: Overview of pin numbers and functions.

Pin	Name	Cell	Purpose/Notes
1	VSS_A	PVSS3AC	Analog ground
2	TXP	PBD1A	200 fF pad! Transmitter positive signal
3	TXN	PBD1A	200 fF pad! Transmitter negative signal
4	VSS_A	PVSS3AC	Analog ground
5	CKP	PBD1A	200 fF pad! Output clock positive signal
6	CKN	PBD1A	200 fF pad! Output clock negative signal
7	VSS_A	PVSS3AC	Analog ground
8	GND_ESD	PVSS2AC	Ring ground
9	VDD_A	PVDD3AC	Analog supply
10	VSS_A	PVSS3AC	Analog ground
11	VDD_A	PVDD3AC	Analog supply
12	I_REFA	PDB3AC	Ref. Current
13	CTRL5	PDB3AC	Control pin 5
14	VSS_CK	PVDD3AC	Clocking circuit ground
15	VDD_CK	PVDD3AC	Clocking circuit supply
16	VBIAS_CK	PDB3AC	Balun DC bias secondary side
17	VSS_LO	Custom RF Pad	Ext. clocking LO ground
18	CK_IN	Custom RF Pad	Must be 700 fF! 10 GHz input clock
19	VSS_LO	Custom RF Pad	Ext. clocking LO ground
20	ENABLE	PDB3AC	Digital enable pin
21	CTRL1	PDB3AC	Control pin 1
22	CTRL2	PDB3AC	Control pin 2
23	CTRL3	PDB3AC	Control pin 3
24	GND_ESD	PVSS2AC	Ring ground
25	CTRL4	PDB3AC	Control pin 4
26	VDD_D	PVDD3AC	Digital supply
27	VSS_D	PVSS3AC	Digital ground
28	VDD_D	PVDD3AC	Digital supply
29	VSS_D	PVSS3AC	Digital ground
30	VDD_A	PVDD3AC	Analog supply
31	VSS_A	PVSS3AC	Analog ground
32	VDD_A	PVDD3AC	Analog supply

Internally, this current is mirrored and distributed to be applied to the tails of the CML drivers.

For control, the chip has one enable signal to enable PRBS generation. Three more control pins can be used at any time to fine-tune internal delays trimmers.

The generated serial data output is driven to differential pins TXP and TXN. The forwarded clock is also driven to differential pins, namely CKP and CLKN. The outputs are placed on the opposite side from the clock input to ease PCB design. In addition, the differential outputs are shielded by placing VSS pads next to and in between them.

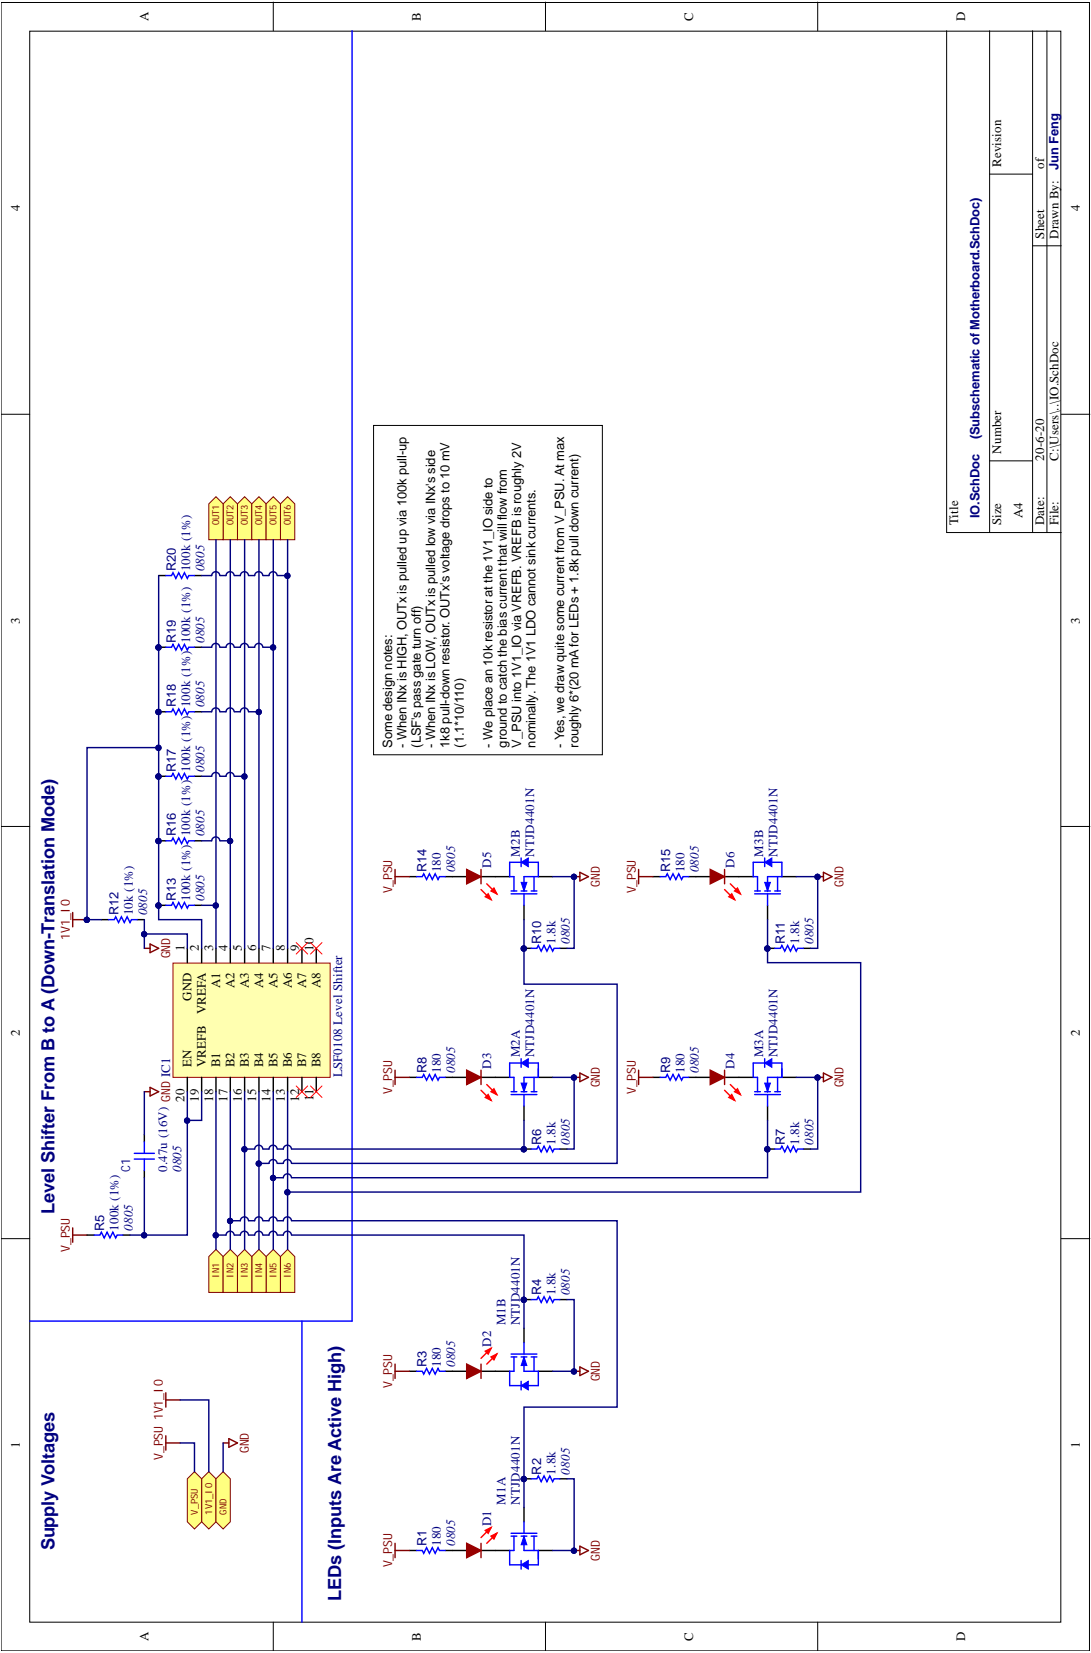
Power supplies are separated for digital, clock generation and analog. In addition, the GSG transformer input is isolated from these domain. As such, the clock input does not have ESD protection. We alternate between VSS and VDD as much possible to minimize bondwire inductance. The analog circuits are sensitive to power supply variation, and rightfully have the most power pins. In hindsight, the clock generation should have enjoyed more VDD and VSS pins than the digital domain as it draws more current.

D Wireline Transmitter Measurement Setup PCBs

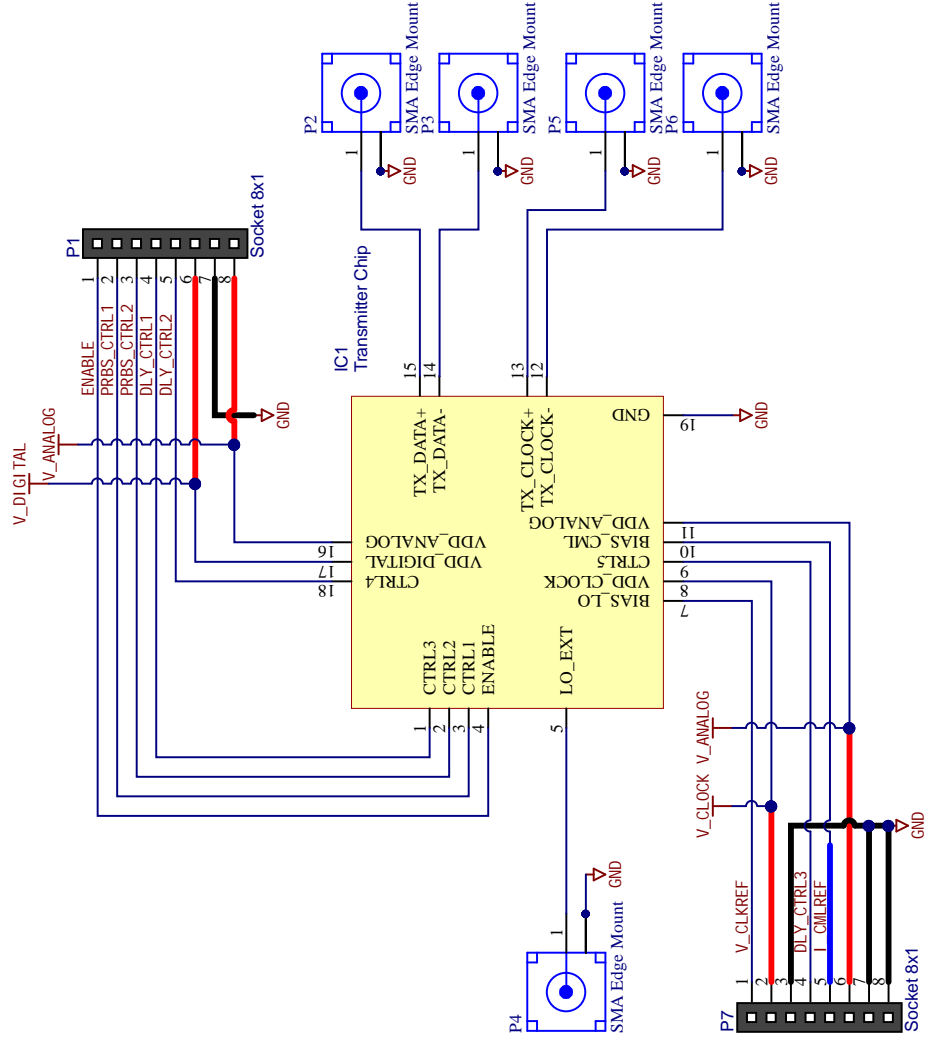
Fig. D.1 shows the measurement PCBs and the measurement setup. The chip sits on a small custom designed test board. The test board is made from FR4 material with a 1 mm thickness to be able to realize the 100 ohm differential characteristic impedances. The testboard mounts on another larger “mother” testboard, also custom designed to provide the necessary power supplies, current measurement and I/O configuration bits.

The full setup is kept simple by virtue of the flexible motherboard. A single power supply, a 10 GHz RF source and high-speed oscilloscope are essential. Optionally, a 100 uA external reference current can be applied to closely observe the change of output swing as function of the reference current. This first daughterboard version uses SMA exclusively.


The following pages show the full PCB schematics.



Chip and Connectors



Decoupling / Filtering



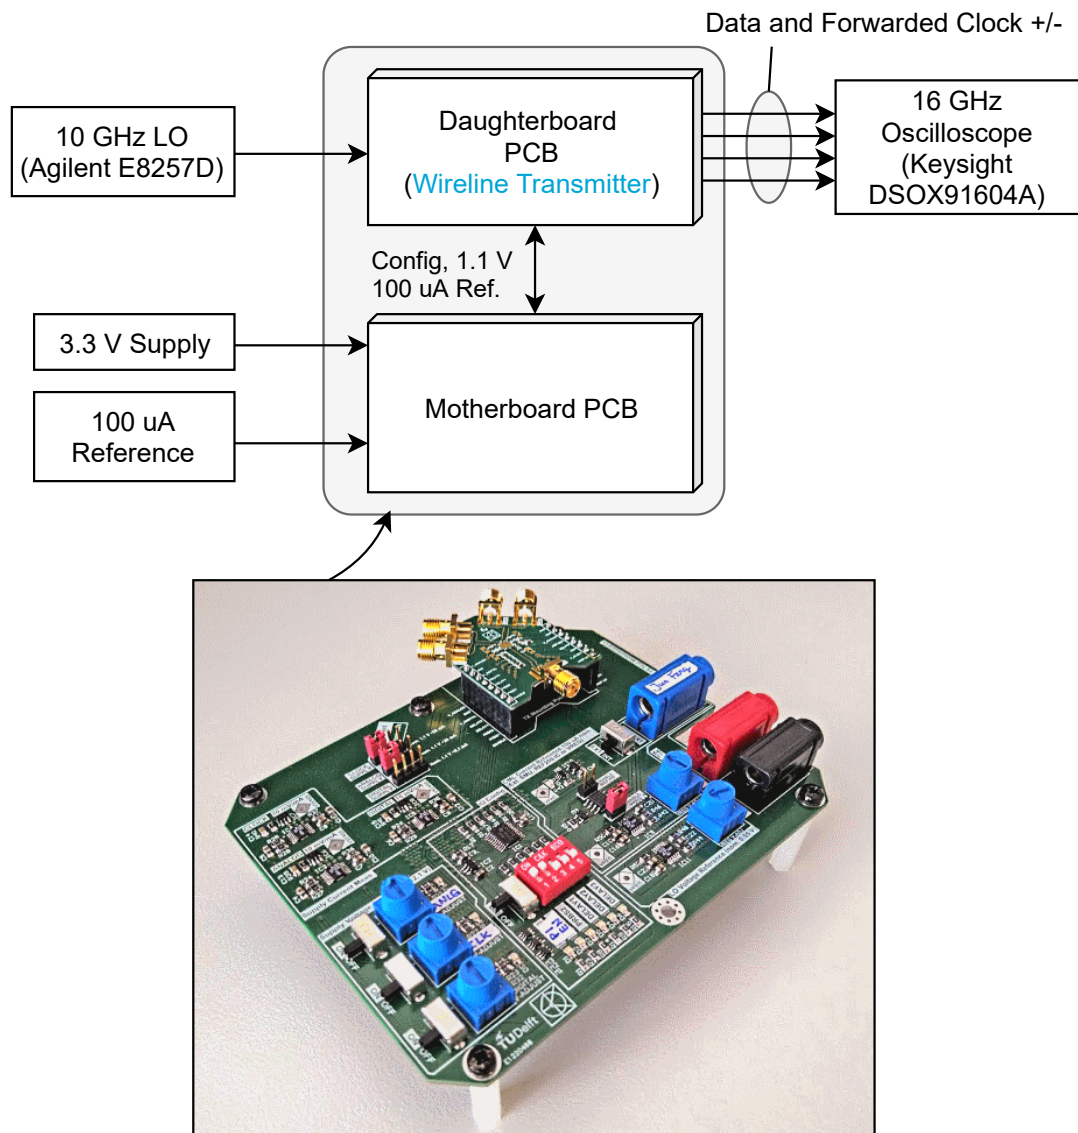


Figure D.1: Overview of ideal measurement setup and fabricated PCBs for this wireline transmitter. The oscilloscope should be AC-coupled.

E BER Derivation for PAM-M Flash Comparators

This appendix includes a slideset as part of this thesis work that derives the BER for detecting a PAM-M signal with a flash ADC. A (simple) flash ADC comprises a number of comparators in parallel, each connected to a different reference voltage [11], as shown on the second slide.

The identical conclusion as derived in the main text for PAM-2 follows from this derivation. Namely,

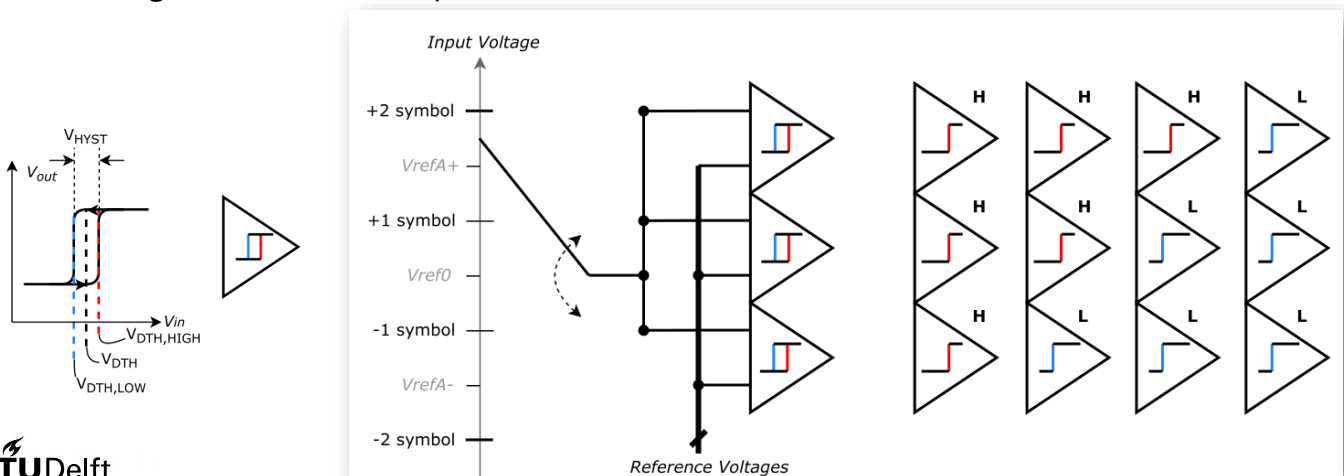
1. For equal-rate detection of random PAM-M signals, we should minimize hysteresis to maximize BER.
2. To compute the effects of a varying detection threshold (hysteresis) on BER, it is required to know the original SNR *and* peak eye opening $V_{PK,EYE}$ to be able to evaluate the new Q_{EYE} .
3. At higher original Q_{EYE} , the relative impact of hysteresis on BER worsens.

Nevertheless, it is worthwhile to have the exact equations for the general solution in one place, that is, slides 9 and 11 below.

BER for PAM-M With Hysteresis

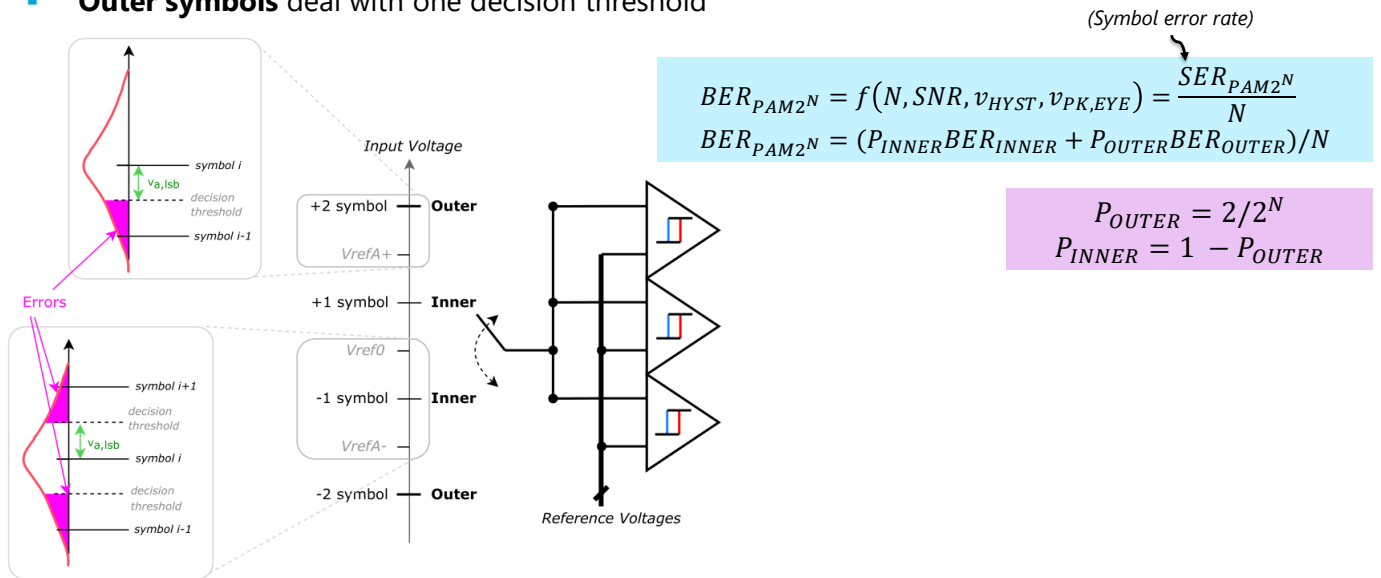
General PAM-(2^N) with Hysteresis

- For general PAM- 2^N the situation is more complex ($2^N - 1$ comparators)
 - Comparator bank's hysteresis state** depends on previous bit: 2^N states
 - E.g. with $4 = 2^2$ comparators \rightarrow 4 states: HHH, HHL, HLL, LLLL



General PAM-(2^N) with Hysteresis


- **Inner symbols** deal with two decision thresholds (two tails spilling over)
- **Outer symbols** deal with one decision threshold



3

PAM-(2^N) with Hysteresis: Inner Symbols

- Symbol lands **between either HH, HL or LL** combination of comparators in the hysteresis state of the comparator bank
- First, BER due to reception **inner symbol** BER_{INNER}
 - **HH or LL case:** one tail shrinks, other tail grows due to hysteresis;
 - growing effect on this "bad" tail is stronger due to Gauss. distr. → BER performance worsens



4


PAM-(2^N) with Hysteresis: Inner Symbols

- **HH or LL case:** one tail shrinks, other tail grows due to hysteresis; growing effect on "bad" tail is stronger → BER performance **worsens**
- Two tails of the same noise distribution → we must add them together

$$BER_{HH,LL} = \int_{x_{good}}^{\infty} PDF_{Gauss}(x) dx + \int_{-\infty}^{x_{bad}} PDF_{Gauss}(x) dx$$

$$x_{good} = \frac{v_{PK,EYE} + v_{HYST}/2}{v_{n,rms}} = \frac{v_{PK,EYE} + v_{HYST}/2}{v_{PK,EYE}/\sqrt{SNR/L_M}}$$


$$x_{bad} = \frac{v_{PK,EYE} - v_{HYST}/2}{v_{n,rms}} = \frac{v_{PK,EYE} - v_{HYST}/2}{v_{PK,EYE}/\sqrt{SNR/L_M}}$$



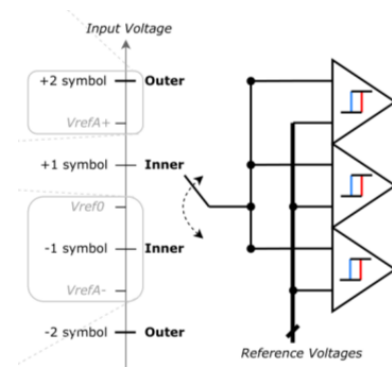
$$\frac{v_{sig,rms}^2}{v_{a,lsb}^2} = \sum_{x=0}^{M/2-1} \frac{2(1+2x)^2}{M} = L_M \rightarrow Q_{eye} = \frac{v_{PK,EYE}}{v_{n,rms}} = \sqrt{\frac{v_{sig,rms}^2}{L_M} \frac{1}{v_{n,rms}^2}} = \sqrt{\frac{SNR}{L_M}}$$

PAM-(2^N) with Hysteresis: Inner Symbols

- **HL case:** both tails shrink! → BER is **better** when symbol lands between HL
- But **HL almost never happens**, P_{HL} is tiny! (for a **random signal**, large N)



$$BER_{HL} = 2 \int_{x_{good}}^{\infty} PDF_{Gauss}(x) dx$$

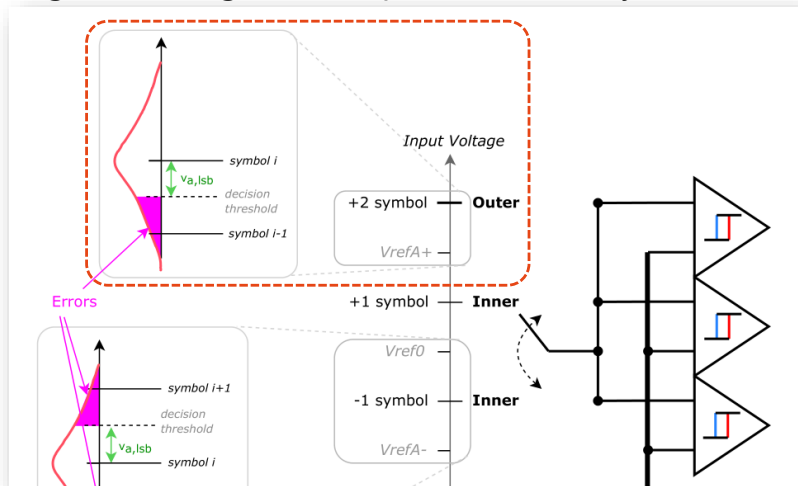


$$BER_{INNER} = P_{HH,LL} BER_{HH,LL} + P_{HL} BER_{HL}$$

$$P_{HL} = \frac{1}{2^N}, P_{HH,LL} = 1 - P_{HL} = \frac{2^N - 1}{2^N}$$


PAM-(2^N) with Hysteresis: Outer Symbols

- BER due to reception **outer symbol** BER_{OUTER}
 - Simple single tail integration, top and bottom symbol have equal conditions



PAM-(2^N) with Hysteresis: Outer Symbols

- Let's look at outer top symbol
 - Low hysteresis **L** would be **good**, high hysteresis **H** would be **bad**



- For a fully random signal, **P_{L,top} << P_{H,top}**! (one out of 2^N states)
 - The exact same scenario unfolds at the bottom symbol (**P_{H,bot} << P_{L,bot}**)

$$BER_{OUTER} = P_{L,top} \int_{x_{good}}^{\infty} PDF_{Gauss}(x) dx + P_{H,top} \int_{x_{bad}}^{\infty} PDF_{Gauss}(x) dx$$

$$P_{L,top} = \frac{1}{2^N}, P_{H,top} = 1 - P_{L,top} = \frac{2^N - 1}{2^N}$$

PAM-(2^N) with Hysteresis: Total BER Random Input

$$BER_{PAM2^N} = f(N, SNR, v_{HYST}, v_{a,LSB}) = \frac{SER_{PAM2^N}}{N}$$

$$BER_{PAM2^N} = (P_{INNER}BER_{INNER} + P_{OUTER}BER_{OUTER})/N$$

$$P_{OUTER} = 2/2^N$$

$$P_{INNER} = 1 - P_{OUTER}$$

$$BER_{OUTER} = P_{L,top} \int_{x_{good}}^{\infty} PDF_{Gauss}(x) dx + P_{H,top} \int_{x_{bad}}^{\infty} PDF_{Gauss}(x) dx$$

$$P_{L,top} = \frac{1}{2^N}$$

$$P_{H,top} = 1 - P_{L,top} = \frac{2^N - 1}{2^N}$$

$$BER_{INNER} = P_{HH,LL}BER_{HH,LL} + P_{HL}BER_{HL}$$

$$P_{HL} = \frac{1}{2^N}, P_{HH,LL} = 1 - P_{HL} = \frac{2^N - 1}{2^N}$$

$$BER_{HH,LL} = \int_{x_{good}}^{\infty} PDF_{Gauss}(x) dx + \int_{x_{bad}}^{\infty} PDF_{Gauss}(x) dx$$

$$BER_{HL} = 2 \int_{x_{good}}^{\infty} PDF_{Gauss}(x) dx$$



$$M = 2^N$$

$$x_{good} = \frac{v_{PK,EYE} + v_{HYST}/2}{v_{n,rms}} = \frac{v_{PK,EYE} + v_{HYST}/2}{v_{PK,EYE}/\sqrt{SNR/L_M}}, x_{bad} = \frac{v_{PK,EYE} - v_{HYST}/2}{v_{n,rms}} = \frac{v_{PK,EYE} - v_{HYST}/2}{v_{PK,EYE}/\sqrt{SNR/L_M}}, L_M = \sum_{x=0}^{M/2-1} \frac{2(1+2x)^2}{M}$$

9

PAM-(2^N) with Hysteresis: Total BER Random Input

- **Random input:** hysteresis hurts BER



PAM-(2^N) with Hysteresis: Total BER Slow Input

$$BER_{PAM2^N} = f(N, SNR, v_{HYST}, v_{a,LSB}) = \frac{SER_{PAM2^N}}{N}$$

$$BER_{PAM2^N} = (P_{INNER}BER_{INNER} + P_{OUTER}BER_{OUTER})/N$$

$$P_{OUTER} = 2/2^N$$

$$P_{INNER} = 1 - P_{OUTER}$$

$$BER_{OUTER} = P_{L,top} \int_{x_{good}}^{\infty} PDF_{Gauss}(x) dx + P_{H,top} \int_{x_{bad}}^{\infty} PDF_{Gauss}(x) dx$$

$$P_{L,top} = P_{same} > \frac{1}{2^N}$$

$$P_{H,top} = 1 - P_{L,top}$$

$$BER_{INNER} = P_{HH,LL}BER_{HH,LL} + P_{HL}BER_{HL}$$

$$P_{HL} = P_{same} > \frac{1}{2^N}$$

$$P_{HH,LL} = 1 - P_{HL}$$

$$BER_{HH,LL} = \int_{x_{good}}^{\infty} PDF_{Gauss}(x) dx + \int_{x_{bad}}^{\infty} PDF_{Gauss}(x) dx$$

$$BER_{HL} = 2 \int_{x_{good}}^{\infty} PDF_{Gauss}(x) dx$$


$$M = 2^N$$

$$x_{good} = \frac{v_{PK,EYE} + v_{HYST}/2}{v_{n,rms}} = \frac{v_{PK,EYE} + v_{HYST}/2}{v_{PK,EYE} / \sqrt{SNR/LM}}, \quad x_{bad} = \frac{v_{PK,EYE} - v_{HYST}/2}{v_{n,rms}} = \frac{v_{PK,EYE} - v_{HYST}/2}{v_{PK,EYE} / \sqrt{SNR/LM}}, \quad L_M = \sum_{x=0}^{M/2-1} \frac{2(1+2x)^2}{M}$$

11

PAM-(2^N) with Hysteresis: Total BER Slow Input

- Just like PAM-2, hysteresis possibly helps when input is **slow**
- E.g. PAM-8 and input speed ~250x slower:



12

F Fast Comparator Simulation Methods

The fast development of the high-speed two-tail comparator would not have been possible without a set of quick and accurate comparator simulation methods. Obviously, a transient simulation reveals the speed of the comparator, but offset and noise simulations easily consume a prohibitive time span using traditional simulation methods.

F.1 Noise Testbench

AC noise simulations are not applicable to the simulation of noise for clocked comparators. After all, these comparators perform charge steering over time and cannot be represented by a single linearized operating point for the whole comparison cycle.

As suggested in [31] this non-linear and time varying noise development can be captured by performing many *transient* noise analyses. For a sufficiently large number of simulation runs, the distribution of comparator decisions (1's vs. 0's) versus input amplitude can be fit to a gaussian cumulative distribution function. Intuitively, for small negative input voltages, the input-referred noise will cause the comparator to still decide for a 1 for some share of the runs as depicted in Fig. F.1.

While possibly accurate, the major downside of this transient noise simulation is the huge number of simulations to obtain for a single noise figure. To achieve simulation accuracy (confidence) for high SNR comparators, a sufficient amount of runs must be ran per applied input amplitude level for and these amplitudes should be finely spaced. For example, 500 transient simulations for 20 input amplitudes spaced 0.5 mV from each other results in 10000 runs.

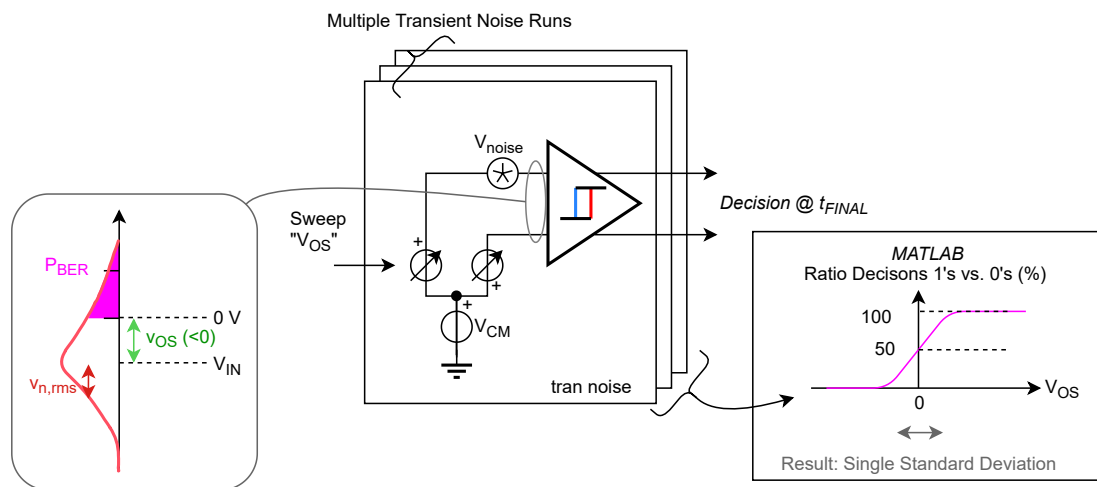


Figure F.1: Traditional, slow noise estimation method employing transient noise simulations.

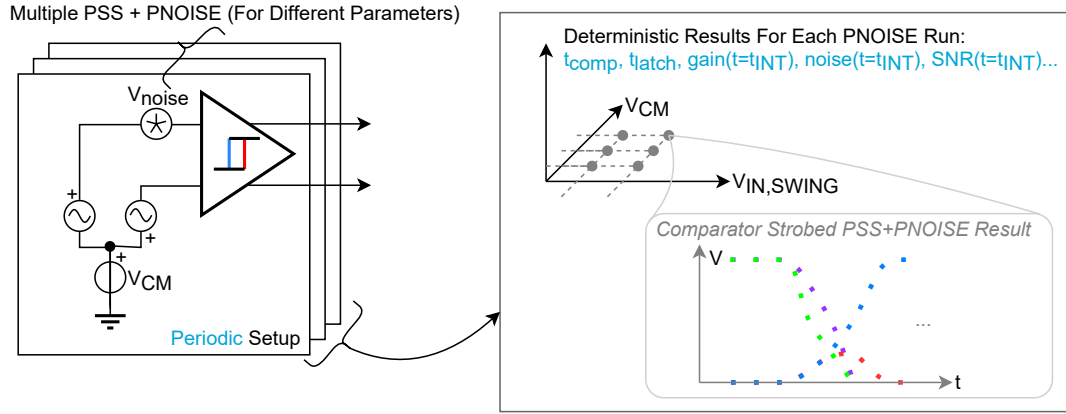


Figure F.2: Fast PSS + PNOISE simulation method in this work, producing deterministic results for a single run.

A much faster method is proposed by [32] based on the periodic steady state (PSS) analysis. In short, these authors recognize that the comparator can also be forced to run in a periodic fashion. Consequently the Spectre simulator can employ its special periodic noise (PNOISE) analysis that accurately measures the noise spectrum evolution over time. The analysis accuracy stems from a combination of a transient simulation and time-varying circuit linearization over the short periodic window. The author proposes to measure the output noise when the comparator differential output voltage crosses $V_{DD}/2$. Then a periodic AC (PAC) analysis should reveal the gain to refer the output to the input.

We agree that the PNOISE analysis is a fast and accurate way to measure noise, but believe that the above method has two problems:

- The PAC analysis should be avoided to compute the voltage gain at the decision event. Our analysis motivates that the gain evolves over the course of two time-windowed integrations. For the comparator simulations in this thesis' work, the gain is obtained from the periodic transient simulation result saved by PSS instead.
- The decision by the latch is not initiated when the latch reaches $V_{DD}/2$. The decision is formed after t_{PRE} for the two-tail comparator, when the latch is about to turn on. Also, the previous method suggests to catch the single decision moment by observing the output voltage nodes. We believe it is much more instructive to continuously strobe the comparator output noise and gain¹.

We depict our methodology in Fig. F.2. A comparison with transient noise simulation results confirm that the input-referred noise number match.

To summarize, a good PSS + PNOISE based setup can yield the evolution of noise and gain accurately over the full operation period of a comparator for a specific input signal with just one simulation (typically less than one minute). A transient noise based simulation requires hundreds of simulations for multiple fine spaced input amplitudes and a curve fit in MATLAB to obtain only the input referred noise at the moment of decision for small input amplitudes (typically 15 minutes at minimum).

¹In technical terms, the previous author proposes to use the event-based *jitter* Noise Type calculation in Spectre, while we prefer the strobed *timedomain* Noise Type calculation. This results in a number of noise analysis results evenly spaced over the periodic time window. Even if we do not catch the exact decision moment, Cadence can interpolate the noise/gain values at other time instances well.

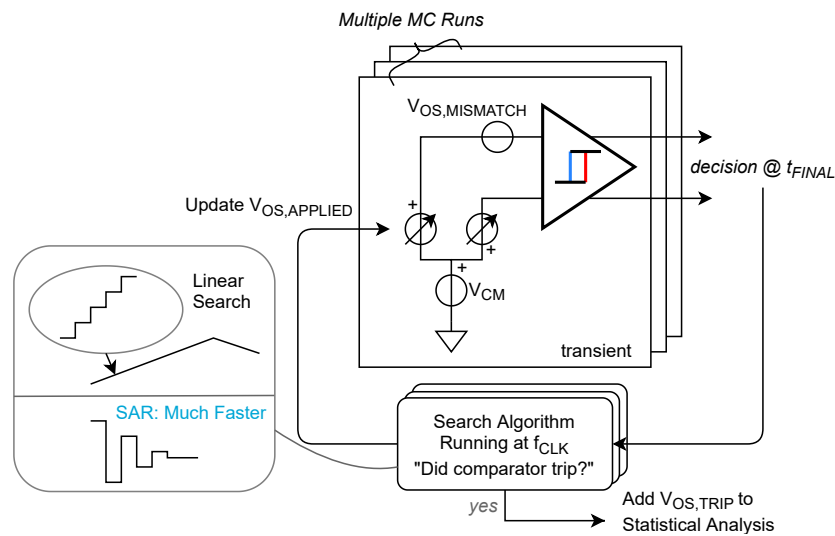


Figure F.3: Offset simulation testbench that can speed up exponentially with the right search algorithm.

F.2 Offset Testbench

The typical simulation method to obtain the input-offset of a comparator is to apply a finely and slowly stepped staircase (or triangular wave) signal to the comparator. Then, for each step the comparator makes a decision and is saved by the testbench. If mismatch is enabled, the comparator will switch its decision centered around a non-zero input voltage (e.g. decision 0 at -26 mV, decision 1 at -25 mV). This trip voltage represents the offset.

Recently, [33] reports a much faster method to obtain the offset. Instead of applying a linear staircase, a binary search algorithm implemented in Verilog-A applies the optimum next input voltage to find the trip point like a successive approximation algorithm as illustrated by Fig. F.3. In this thesis work a successful adaptation results in an $\sim 10\times$ increase of simulation speed when looking for a trip point with 1 mV accuracy.