A 10Gb/s Cryogenic Clock and Data Recovery System with Low Jitter

L. de Jong

A 10Gb/s Cryogenic Clock and Data Recovery System with Low Jitter

by

L. de Jong

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Tuesday October 19, 2021 at 14:00 PM.

This thesis is confidential and cannot be made public until October 19, 2023.

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

TUDelft

Abstract

A key issue in current quantum computing interfaces is the dense interconnect between electronics at cryogenic temperature (CT) and room temperature (RT). Recently, progress has been made to move more control electronics from RT to CT, reducing interconnect overhead. The next step towards minimal interconnect is a direct wireline interface between RT and CT. This work presents a full-rate 10 Gb/s clock-and-data recovery circuit for a high speed serial link receiver operating at CT.

A novel phase detector is utilized to reduce power consumption by removing the need for both a pulse generator at the input and, a buffer between the phase detector and voltage controlled oscillator. Additionally, a digital delay-locked loop is added to improve the retiming margin, achieving higher jitter tolerance. Implemented in 40-nm CMOS, post-layout simulation shows a core power consumption of 3.89 mW from a 1.1-V supply at 10 Gb/s, producing an rms-jitter of 84 fs and an estimated jitter tolerance of 1.1 $\mathrm{UI_{pp}}$ at 10 MHz.

Acknowledgements

Contents

1

Introduction

This chapter introduces the requirements of a clock and data recovery (CDR) system. It then presents the motivation and target specifications. The research contribution is discussed in the third section. Finally, the thesis outline is described in the last section.

1.1. CDR Requirements

Data transmission across a wireline from transmitter to receiver often consists of only one link. A basic block diagram of this wireline system is shown in Fig. 1.1. The data is clocked at the transmitter's frequency and phase prior to being transmitted across a channel. The receiver consists of an equalizer that compensates for channel losses and amplifies the signal. The receiver has no knowledge of the incoming frequency and phase of the transmitted data. A clock recovery block extracts this information from the data and recovers a clock signal which is then used to retime the signal at the receiver [1]. The combination of the clock recovery circuit and the decision circuit is called clock and data recovery.
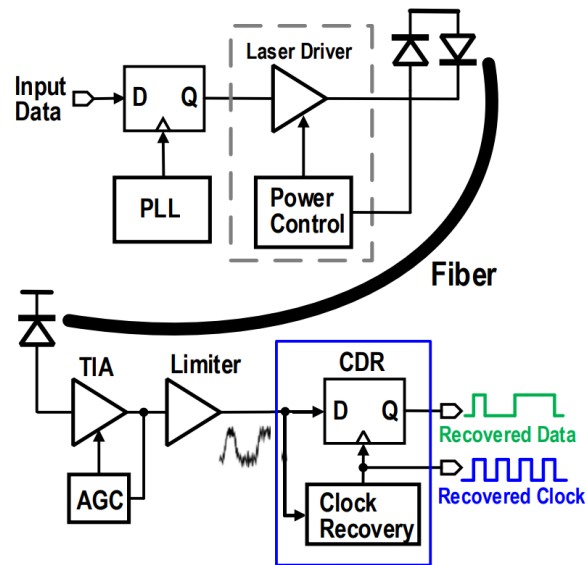
Figure 1.1: Block diagram of a wireline link.

Fig. 1.2 shows an example of an eye diagram seen at the receiver, which includes jitter, noise, and inter-symbol interference (ISI). To successfully resample each incoming bit, the sampling instance t_s must be far enough away from the transitions such that the voltage is above or below the decision voltage V_s for a '1' or '0' sample respectively. Based on this a few functional requirements can be determined for such a system:
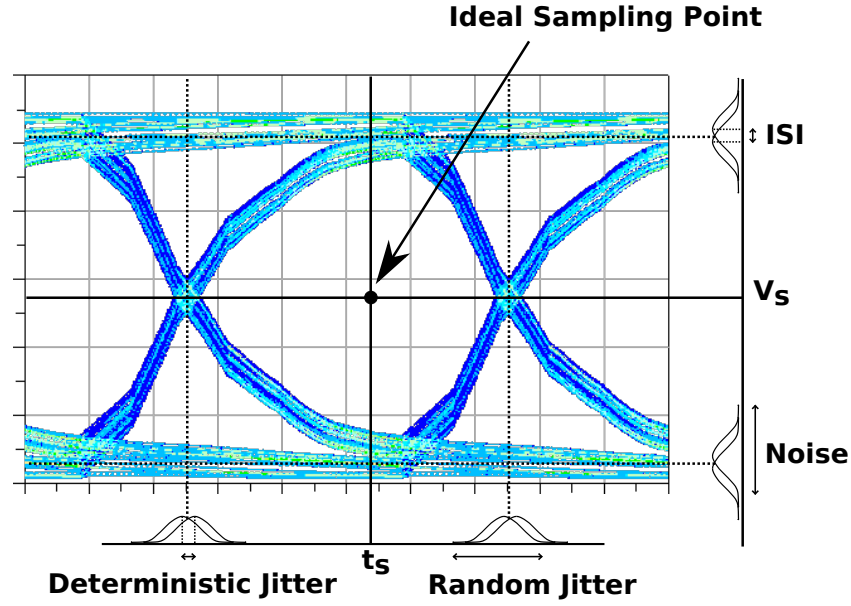
Figure 1.2: Example of a received wireline eye diagram with indicated optimal sampling point.

- The sampling clock must be synchronous with the received data to be able to assess each incoming bit. For a 10 Gb/s signal, a clock period of 100 ps is desired.

- The sampling clock itself should have low jitter as it directly impacts the quality of the retimed synchronous data.

- The timing distance between the incoming data transition edges and the sampling moments should be fixed to achieve the best possible bit error rate (BER).

1.2. Motivation and Target Specifications

A key issue in current quantum computing interfaces is the dense interconnect between electronics at cryogenic temperature (CT) and room temperature (RT). Practical quantum algorithms require thousands of quantum bits (qubits) which have to be controlled and read-out [2]. This issue is addressed by moving the control and read-out electronics to CT [3]. Recent developments have seen more and more control electronics moved to 4-K [4] [5] to reduce physical interconnect. Following this trend, the lowest achievable interconnection between CT and RT will be two cables as shown in Fig. 1.3, where only the quantum algorithm is executed at RT. To achieve such a system, numerous blocks have to be developed at CT, one of them being the high-speed receiver. With the goal of minimizing interconnect, the wireline link carrying the data consists of only a single cable. This calls for a clock recovery system at CT as the data will only gain meaning when referenced to a clock.

Power consumption at cryogenic temperatures is limited by the cooling power available in the fridge [3]. The 4-K stage of the system has only a few watts of cooling power and as a result, the power consumption of individual blocks must be limited. Considering both this constraint and the performance of current state-of-the-art CDRs operating at room temperature, the power consumption target for the CDR is 5 mW.

Phase noise performance is known to improve at 4-K [5] and while jitter is relevant to the achievable (BER), the jitter of all published designs is sufficiently low to meet arbitrary BER requirements [1]. To still obtain a metric of a CDR's performance, a slowly varying jitter is applied to the incoming data stream until the BER rises above a certain threshold, typically $10e-12$. At jitter frequencies inside the bandwidth, the system is able to track the movement. However, outside the bandwidth, it is no longer able to do so. Standards have been developed for wireline interfaces which define limits on what transmitters
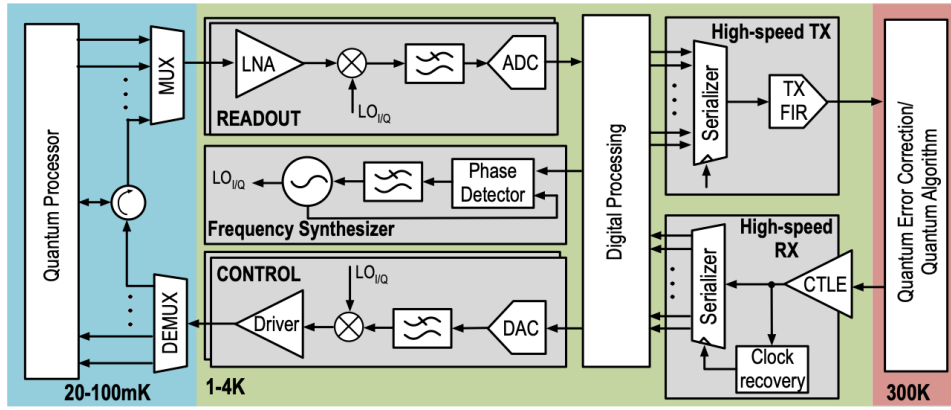
Figure 1.3: Block diagram of complete cryogenic interface.

and receivers must be able to handle. Meeting the jitter tolerance mask of such a standard guarantees that the CDR is compatible with wireline transmitters designed by others. Currently, the strictest mask for jitter tolerance is for Synchronous Digital Hierachy (SDH) which is the STM-256 shown in Fig. 1.4. The design should at minimum fulfill the mask requirement and the targeted jitter tolerance is $1\,\mathrm{UI_{PP}}$ at $10\,\mathrm{MHz}$.
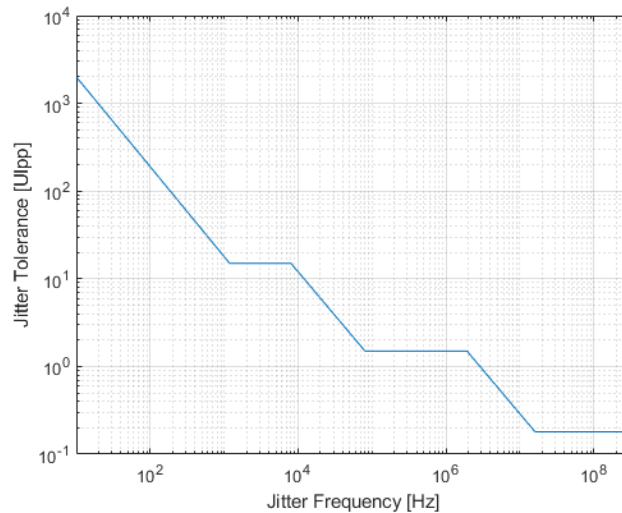
Figure 1.4: The STM-256 Jitter Tolerance Mask [6].

Additionally, earlier work within the group on high fidelity frequency synthesis [5] gives rise to another design target; to extract the clock with extremely low jitter. The current lowest reported recovered clock jitter is above $250\,\mathrm{fs_{rms}}$ [7]. Based on this and the synthesizer performance, the target jitter is set at $150\,\mathrm{fs_{rms}}$. Moreover, it should achieve this result without an external crystal reference as this would result in additional unwanted overhead.

The target specifications are summarized in Table 1.1.

1.3. Thesis Objective

Based on the motivations listed in Section 1.2, the objective of this thesis is to design and implement a clock and data recovery system that fulfills the specifications listed in Table 1.1. The system should operate both at room temperature as well as cryogenic temperature.

Table 1.1: Specifications for the design.

Parameters	Specifications
Data Rate	10 Gb/s
Recovered Clock Jitter	$<150\,\mathrm{fs_{rms}}$
Jitter Tolerance at 10MHz for BER 10e-12	1 UIpp
Power Consumption	$<5\,\mathrm{mW}$
External Reference	No

1.4. Thesis Outline

The thesis outline follows the project progress. Chapter 2 gives a background on CDR systems and discusses the proposed architecture as well as prior published art with their merits and shortcomings. In Chapter 3, modeling of the system is covered to find the required component parameters. Chapter 4 then discusses the design and simulation of various analog and RF blocks of the system. The RTL design and verification are presented in Chapter 5. Chapter 6 covers the top-level system simulations and performance. Chapter 7 concludes the thesis and discusses future work.

1.5. Research Contributions

The research contributions of this thesis are as follows:

- Introducing a novel CDR architecture (Chapter 2);

- Applying general knowledge of control theory to model a CDR (Chapter 3);

- Proposing a charge-sampling-based phased detector for CDR (Chapter 4);

- Proposing a delay-locked loop to calibrate the retiming margin (Chapter 4, 5);

- Design and implementation of a CDR in CMOS 40-nm technology with state-of-the-art jitter performance (Chapter 4, 6).

2

PLL-Based Clock and Data Recovery

This chapter aims to introduce a phase-locked loop (PLL)-based CDR architecture to meet the specifications listed in Table 1.1. It starts with an introduction to CDRs discussing their jitter metrics. This is followed by a comparison between CDR architectures. The two classical phase detector types, linear and bang-bang, are then analyzed. State-of-the-art developments and their shortcomings are addressed, and finally, a new CDR architecture is proposed.

2.1. Introduction to CDRs

In a wireline link, the receiver receives a signal containing random data that is asynchronous from its perspective. This signal is degraded due to numerous effects occurring in the transmission channel. To clean up this signal at the receiver, the clock and data recovery system utilizes either a forwarded clock signal or extracted clock signal, and resamples the data stream, resulting in synchronous data and clock outputs. There are, of course, a number of different applications where such a system is used, such as in repeaters and chip-to-chip connections. The choice of architecture is therefore highly dependent on the specification. For example, a CDR where both the clock and data are forwarded will employ a different architecture than a CDR receiving only data. To better understand and judge the performance of these architectures, we will discuss jitter in CDR systems in this section.

2.1.1. Jitter in traditional CDR systems

Jitter performance in CDR architectures is usually described by three metrics.

- jitter transfer, it consists of the jitter transfer from input to output, showing how much jitter on the incoming data will affect the output jitter.

- jitter generation, which includes jitter generated by the CDR circuit itself.

- jitter tolerance, which describes the amount of jitter the system can handle while maintaining a constant BER.

Each of these metrics will be discussed more in depth below.

Jitter transfer
The incoming data in every physical system will always contain jitter. The amount of this jitter that will be seen at the output highly depends on the architecture. The jitter transfer in a closed-loop feedback circuit can be expressed as

$$|\frac{\phi_{out}}{\phi_{in}}| = \frac{H_{ol}(s)}{1 + H_{ol}(s)} \tag{2.1}$$

Where $H_{ol}(s)$ is the open-loop transfer of the system. A first-order system will, therefore, directly transfer the jitter on the data to the recovered clock inside its bandwidth, but not amplify it. All jitter on the data outside of the loop bandwidth will be attenuated. In a second-order system, the jitter will experience some amplification due to the zero-pole-pole combination. Standards such as SONET require a limit on jitter amplification or 'peaking' of 0.1 dB. A more thorough analysis of the loop dynamics will follow in Chapter 3.

Jitter generation

Jitter generation in a CDR system originates from 4 main sources: 1) voltage-controlled oscillator (VCO) phase noise, 2) ripple on the control voltage, 3) coupling of data to the VCO through the phase detector and retiming, 4) supply and substrate noise. Each of these four will be elaborated below.

A VCO produces phase noise due to its non-ideal components. This phase noise shows three regions distinguished by their transfer slope centered around the oscillation frequency. The $1/f^3$ region closest to the fundamental tone, which has a 3rd order roll-off due to up-converted flicker noise. Followed by the $1/f^2$ region, which continues until it hits the noise floor, after which the noise floor becomes dominant [8]. Aside from the VCO, other components in the system also contribute to the total phase noise.

The second source relates to the control voltage, which tunes the VCO frequency based on the feedback in the loop. In the ideal case this voltage is constant so that it does not introduce additional noise and spurs. However, due to many nonidealities such as mismatch, feedthrough and noise, there will be some ripples present. Moreover, in CDR, the data is random and as such the ripple on the control voltage is also random [1].

The third source of jitter is also related to the random data. Switching in the retimer and PD occurs at random moments and these switching transients can couple back to the VCO introducing additional jitter. Buffers between the VCO and respective retimers and PD can be used to shield this effect [9].

The fourth source is random interference caused by supply and substrate noise. Since the circuit is highly integrated, there is interference from neighbouring sources that cannot be fully isolated. A CDR is often part of a larger system and one has to be conscious of the impact of parasitic coupling with other circuit blocks such as nearby digital blocks.

Jitter tolerance

As described earlier, a closed-loop CDR has a low-pass jitter transfer, accurately tracking the phase error during slow variations while ignoring those outside of the bandwidth. When this jitter occurs inband the resampling of data in the CDR circuit occurs exactly in the center of the bit. If the jitter is not within this band, the sample will be taken off center, effectively moving the sides of the eye closer to the center. If this jitter becomes too severe, the eye will become compromised and the BER will suffer. Jitter tolerance describes how much sinusoidal jitter can be applied at a frequency before the BER degrades below a certain threshold, i.e., 10^{-12}.

Intuitively it can be understood that a higher bandwidth increases the jitter tolerance of the system. This shows a fundamental tradeoff between jitter transfer and jitter tolerance in typical CDR system; a higher bandwidth will allow more jitter to pass to the output, but is also able to track higher frequency variation.

The jitter tolerance is an important specification in wireline standards such as SONET and SDH. The required jitter tolerance is defined by a mask as shown in Fig. 2.1. Each standard will have its own corner frequency values and corresponding amplitudes that the CDR's jitter tolerance must meet. If it crosses this standard line, it fails. Jitter tolerance is tested using a dedicated BERT such as the KeySight N4903B.

2.2. Top-level CDR Architectures

Existing CDR systems can be divided into feedback-based and non-feedback-based structures. Feedback-based approaches are dominant due to their better jitter performance and suitability in 'classical' CDR applications such as SONET and high-speed serial links operating at a continuous rate. Within the feedback-based category, distinctions can be made between different structures and, to find the most suited architecture, a comparison between them will follow below.

Figure 2.1: Typical jitter tolerance mask for wireline standards.

2.2.1. PLL based CDR

The 'classical' method of CDR is based on PLLs, where the reference input is replaced by the incoming data signal and a phase detector specifically for random data is used. Within this category, a few architectural differences are present. There are architectures with and without an external frequency reference for frequency tracking, as well as shared and separate loop filters for frequency and phase acquisition. Developments towards All-Digital CDR created another set of architectures [10]. Each of these will be briefly examined in this subsection.

Referenceless CDR The classical referenceless CDR is shown in Fig. 2.2a. It consists of two loops, a phase tracking loop and a frequency tracking loop. During system initiation, the VCO frequency and incoming data rate could be too far away to achieve phase locking with only the PLL. Hence, a separate frequency detector is used to pull in the VCO to the desired frequency. When it is close enough for the PLL to achieve lock, the PLL will dominate the system and drive it to phase lock [11].

Two potential issues can arise from this topology. First, the frequency-tracking loop and phase-locking

(a)

(b)
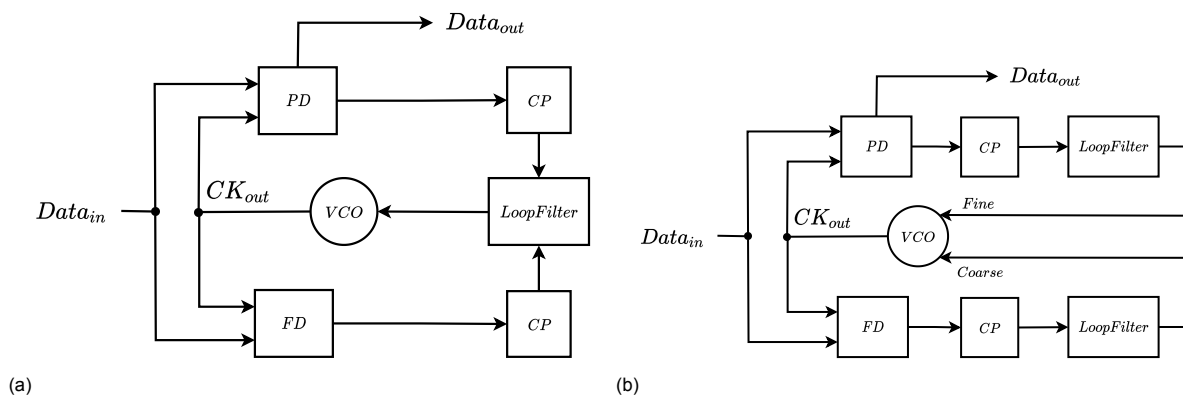
Figure 2.2: Referenceless CDR architectures with (a) Shared loop filter and, (b) Separate loop filters.

loop can interfere with each other, possibly preventing both a lock and generating ripples on the control line. Second, the frequency detector can be disturbed by spectral lines near the VCO frequency that can appear due to the randomness of the input data. Due to this, the frequency-tracking loop bandwidth should be a fraction of the PLL bandwidth, slowing down locking [12].

To overcome this clash between the two loops, it is possible to use the architecture shown in Fig. 2.2b.

Here, two independent loops are used, the frequency tracking sets the coarse tuning of the VCO while the PLL controls the fine tuning allowing an independent choice of bandwidth, thus speeding up locking [13]. The downside to this alternative is the need for two loop-filters, a block that occupies a large area, subsequently increasing cost.

CDR with external reference   The dual loop approach in Fig. 2.2a can also be adjusted to function with an external frequency reference, shown in Fig. 2.3a. The coarse tuning voltage is derived from a second PLL where $f_{data} = N \cdot f_{ref}$ which, in theory, guarantees that VCO_1 has the correct frequency. This allows for very little disturbance on the control line and therefore low jitter.

In practice, however, this approach has quite a few issues. First, the reference frequency must be exactly N times smaller than the incoming data frequency, something that will not be the case due to mismatch between the references at the transmitter and receiver [12]. Frequency pulling could then drive VCO_1 away from the data frequency towards $N \cdot f_{ref}$. Second, there will be mismatch between VCO_1 and VCO_2. This causes their frequencies to differ even with the same coarse input. The PLL must therefore have enough locking range to guarantee lock under mismatch. The matching issue is exacerbated by the layout of the two VCOs and loop filters which require a large area due to their inductors, capacitances and respective power routing.
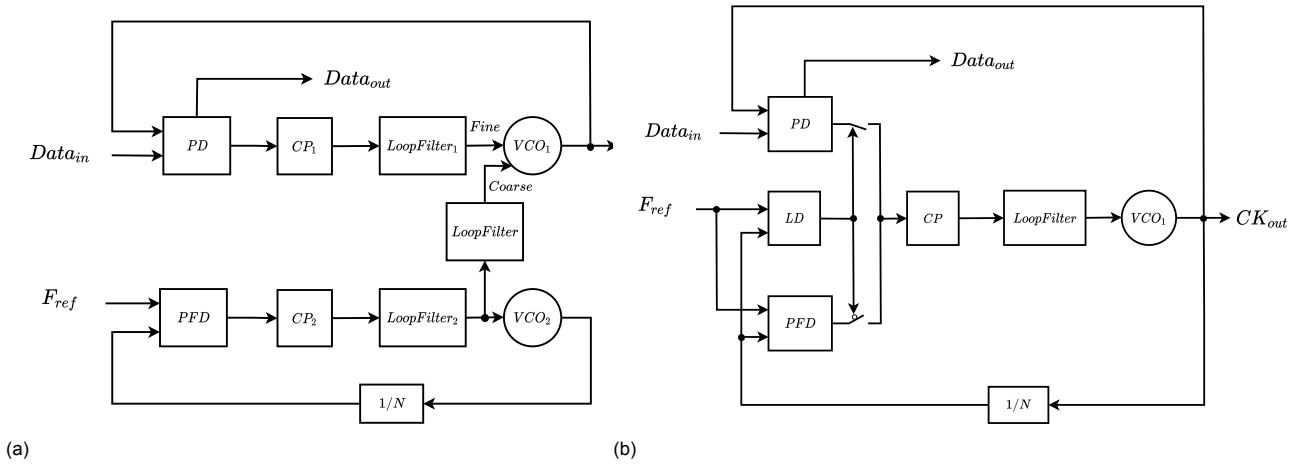
Figure 2.3: Two main CDR architectures with frequency reference. (a) Shared loop filter. (b) Separate loop filters.

The design can be altered to alleviate some of these issues with a lock detector as shown in Fig. 2.3b. Only one loop is active at a time, requiring only a single VCO and loop filter. First, the frequency tracking loop will be active, bringing the VCO frequency close to that of the incoming data. The lock detector then recognizes the frequency lock and toggles to the PLL which then acquires phase lock. Important here is to ensure that the control voltage does not jump significantly when switching between the two loops because the frequency could jump away and the system will fail to function [12].

Digital CDR   Recently, there have also been developments towards All-Digital CDR (ADCDR), similar to All-Digital PLLs. They aim to reduce area and add flexibility due to the use of digital loop filters (DLF) and ease of programming. An intermediate step towards such an architecture is shown in Fig. 2.4a, where the phase and frequency detector outputs are digital and fed to the digital loop filter. The outputs are then converted back to analog voltages to tune the VCO [14]. This brings two issues. First, the loop delay introduced by the digital loop filter and digital-to-analog converter (DAC) degrade the stability and can cause the system to fail. Second, the finite resolution of the DACs will add frequency drift and additional jitter to the system. One proposed way to ease the specification of the DAC is by using a delta-sigma modulator. However, this is limited to the integral path only since the proportional path is very sensitive to loop delay and must run at the data rate. [14] [15].

This first structure is only an intermediate step to the All-Digital CDR in Fig. 2.4b. Such a system

directly feeds the loop filter output to switched capacitors inside the oscillator, now a DCO [10]. The flexibility, resolution and loop latency remain in this architecture. ADCDRs are therefore mostly used in mid-range applications where their flexibility shines and the issues can be tolerated. The DLF can be made fully synthesizable by extensive use of subsampling after the PD, reducing the manual workload caused by the manual design and layout of the proportional path running at RF [16].
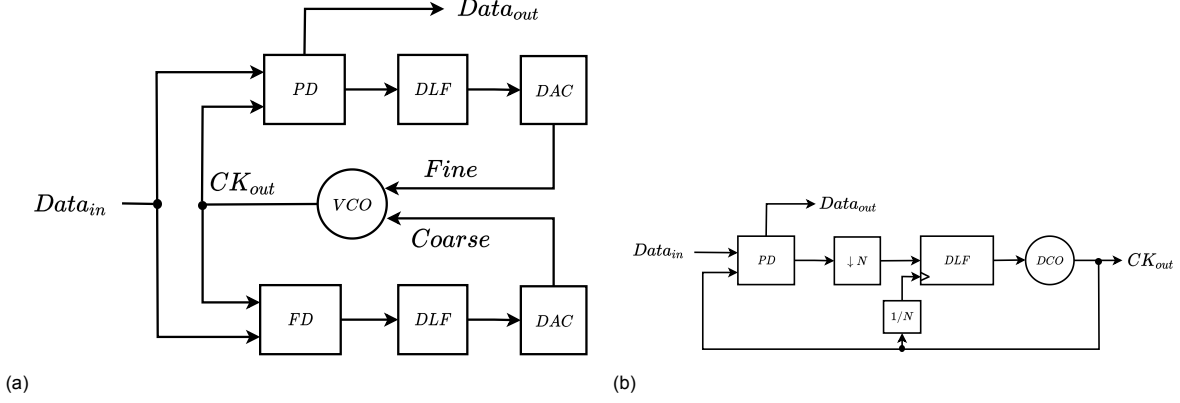
Figure 2.4: Digital CDR architectures. (a) Partially digital CDR with analog VCO. (b) All-Digital CDR with DCO.

2.2.2. DLL based CDR

Delay-locked-loop (DLL)-based CDR is very similar to its PLL counterpart. The main difference is the VCO being replaced by a voltage-controlled delayline (VCDL) as illustrated in Fig. 2.5. The second loop with external reference now supplies the clock signal to the delayline instead of the coarse voltage in PLLs. The main delay-locked loop then locks the phase using the control voltage coming from the feedback loop.

The two main benefits of the DLL structure over a PLL relate to the change in loop dynamics due to the VCDL. In a PLL, the tuning voltage controls the frequency, thus indirectly controlling the phase, adding a pole to the transfer. In the phase domain model it is described as:

$$\frac{\phi_{out}}{V_c} = \frac{K_{VCO}}{s}$$

The VCDL on the other hand directly tunes the phase and can be described by:

$$\frac{\phi_{out}}{V_c} = K_{VCDL}$$

Where K_{VCO} and K_{VCDL} are the oscillator gain and delayline gain in $\mathrm{rad/V}$ respectively. It can be observed from the expressions that a DLL-based structure will have one pole less in the loop transfer and, as a result, is a first order system. A first-order system has the benefits of being inherently stable and having no peaking their transfer function, meaning that no jitter peaking occurs [17].

The main drawbacks of this architecture relate to the procurement of the supplied frequency to the VCDL. The incoming data rate must be known to match the frequency to allow locking since the locking range of the system is small. As such, this structure is best suited for synchronous applications where the transmitter data rate is known, such as chip-to-chip interconnections [18].

2.2.3. Hybrid DLL/PLL CDRs

It is possible to combine the flat jitter transfer of the DLL structure and fast locking by employing the hybrid DPLL structure in Fig. 2.6a. It consists of two loops where the VCDL and VCO are controlled by the same voltage. Analysis of the loop dynamics shows that the jitter transfer bandwidth and jitter tolerance can be set independently through the choice of K_{VCO} and K_{VCDL} [19]. The use of a shared control voltage can make the loop unstable if the VCDL is near its delay limits and starts to function as an open-loop system [18].
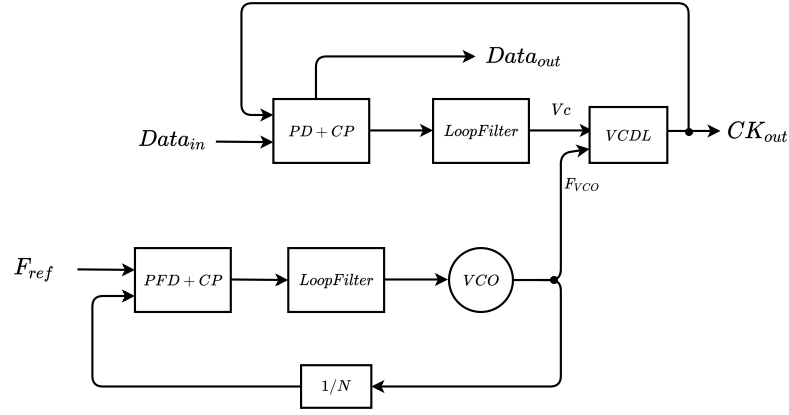
Figure 2.5: Typical delay-locked loop CDR structure.

The instability concern can be avoided by separating the two loops as illustrated in Fig. 2.6b. It appears similar to the DLL structure of Fig. 2.5, where now the reference input is replaced by the incoming data. This difference maintains the benefits of both loops and removes the open-loop stability risk. The cost of this is expressed in the two loop filters and charge pumps required [20]. Another factor to consider is the high power consumption of the VCDL operating at RF.

(a)   (b)
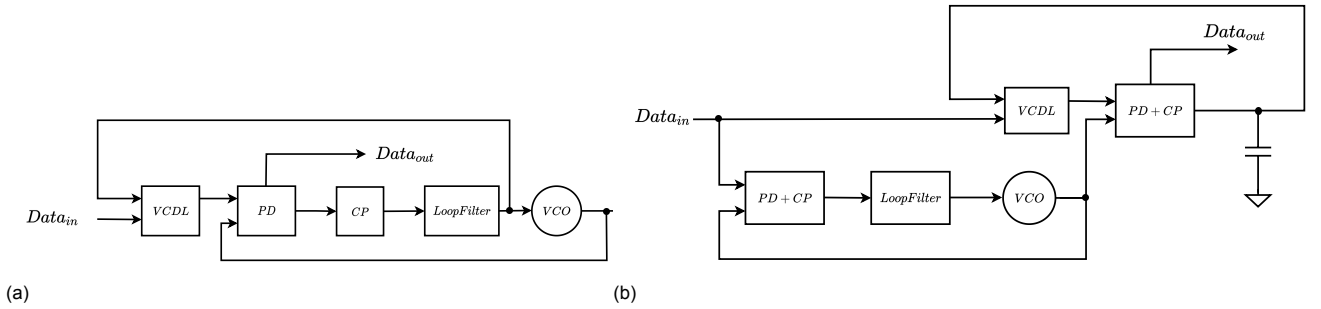
Figure 2.6: DPLL CDR architectures. (a) Double loop structure with shared control voltage. (b) Separated control for both loops.

2.2.4. PI based CDR

Phase interpolators (PI) are used in serialize-deserialize (SerDes) transceivers due to their relatively small footprint [21] and are relevant at a system-level architecture of a wireline receiver. The main principle is that only one PLL, locked to an external reference, is used to generate the multiphase receiver clock. These clock phases are then distributed to each CDR. Multiple phases are required for each deserialized lane, and subsequently connected to a PI. Fig. 2.7 shows the structure with a single lane. The phase interpolator has a discrete number of settings that can select and shift the phase of the clock from 0 to 360 degrees. Each sub-CDR then has a feedback loop to control the setting of the PI rather than the PLL [21].

The main benefit of this approach is area. A PLL with an LC-VCO occupies a significant amount of area due to its inductor and loop filter, so reducing the number of PLLs from the number of lanes to only one is a major advantage. However, this area save does have three drawbacks. First, it requires a large clock distribution network with long interconnect for all the phases. Second, the finite resolution of the PIs can introduce significant jitter into the system. Third, it requires an external reference or forwarded clock to lock the PLL, such an additional lane or crystal is relatively expensive.
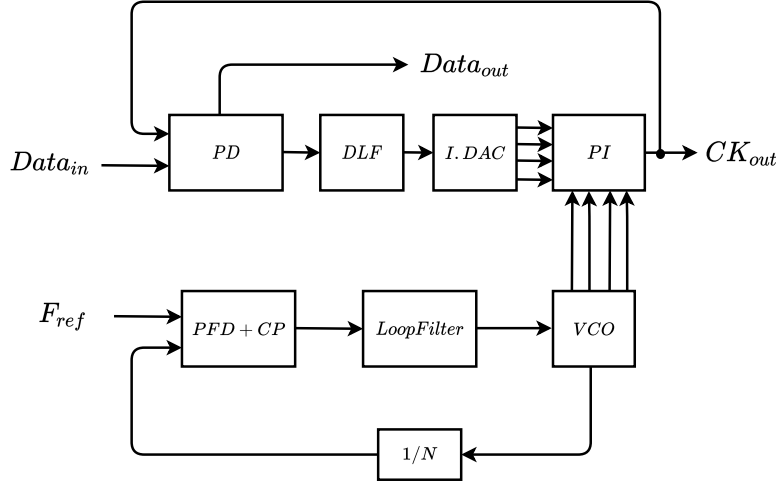
Figure 2.7: Phase interpolator structure attached to a single lane.

2.2.5. Injection-Locked CDR

Similar to the PI structure is the injection-locked CDR shown in Fig. 2.8. The PI and I-DAC are replaced by a phase selector, injection driver, and slave oscillator. These fulfill the same roles as the PI and I-DAC for the PI structure. The slave oscillator is locked to the injection driver which is controlled by the phase selector. The use of a slave oscillator instead of PI brings the advantage of additional filtering, smoothing out the discrete phase transitions present in its counterpart and filtering out large duty-cycle distortions [22]. Drawbacks include the limited locking range of the slave oscillator and potential unwanted coupling to the slave oscillator due to layout parasitics [18].
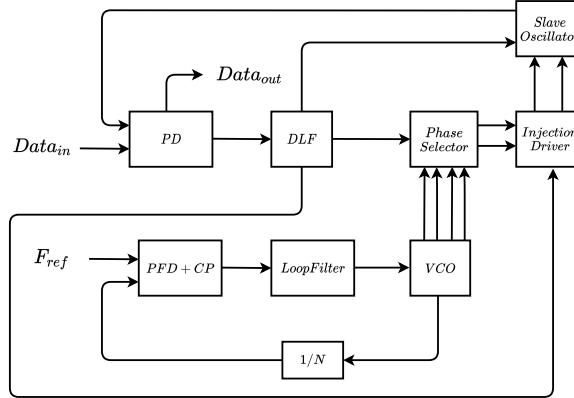
Figure 2.8: Injection locked CDR structure.

2.2.6. Comparison

Table 2.1 lists the pros and cons of each of the discussed architectures as well as some of their reported applications. The architectures all have their own strengths and weaknesses depending on the application. They are all suited for continuous-mode operation where the system is operating at all times. All, except for the DLL structure, are suitable for source-asynchronous systems where only a data signal is received from the transmitter side. The DLL structure requires the transmitter frequency to properly lock and is therefore only suited for source-synchronous operation. Both the PI, IL and all-digital architectures will have higher jitter due to their inherent quantization and are therefore not selected. While the D/PLL architecture does decouple the jitter tolerance and jitter transfer bandwidth to a certain extent, it comes at the cost of higher power making it unsuitable to meet the power requirement. Therefore, the selected CDR architecture for this project is analog PLL based. The next section will discuss different phase detection methods for PLL-based CDR.

Table 2.1: Overview of the discussed CDR architectures with reported applications listed.

	Pros	Cons	Applications
PLL	Input Jitter Rejection Input Frequency Tracking Fast Locking	Jitter Peaking Large Loop Filter Multichannel Crosstalk	Source-Asynchronous/Synchronous High Speed Serial Links [11] SONET/Gigabit Ethernet [23]
ADPLL	Input Jitter Rejection Input Frequency Tracking Fast Locking	Jitter Peaking Loop Latency Quantization Error	Source-Asynchronous/Synchronous High Speed Serial Links [24] SONET/SDH [25][16]
DLL	Stable, 1st Order System No Jitter Peaking Single Clock for Multichannel	Large Loop Filter Limited Locking Range Synchronous Only	Source-Synchronous Chip-to-Chip Interconnection [26] Intra-Panel Interface [27]
D/PLL	No Jitter Peaking Fast Locking Small Loop BW	Multichannel Crosstalk Dual Loop Analysis	Source-Asynchronous/Synchronous High Speed Serial Links [28] SONET/Gigabit Ethernet [29]
PI	Single Clock for Multichannel Area Efficient	PI Quantization Error Clock Routing	Source-Asynchronous/Synchronous SATA/PCIe/Gigabit Ethernet [30] Chip-to-Chip Interconnection [21]
IL	High Jitter Tolerance Single Clock for Multichannel Duty-Cycle Correction	Quantization Error Oscillator Layout	Source-Asynchronous/Synchronous SONET [22]

2.3. Linear or Bang-Bang Phase Detection

Phase detection in a CDR differs from conventional phase detectors found in PLLs due to the nature of their input signal. The most common modulation method in wireline communication is non-return-to-zero (NRZ), shown in Fig. 2.9. It consists of a random bitstream at a rate R_b where each bit has a

Figure 2.9: NRZ data waveform.

duration of $1/R_b = T_b$, also known as a unit interval (UI). This random sequence can be expressed mathematically by:

$$x(t) = \sum_k b_k p(t - kT_b) \tag{2.2}$$

where $b_k = \pm 1$ and $p(t)$ is the pulse shape of a single pulse. It can be shown that the spectrum of this random sequence is expressed by:

$$S_x(f) = T_b \left(\frac{\sin(\pi f T_b)}{\pi f T_b} \right)^2 \tag{2.3}$$

From this expression, we find that if $f = n/T_b$, the function evaluates to 0, which means that at the original clock frequency and its harmonics, there is no spectral component, as can be seen in Fig. 2.10. Therefore, the phase detector in a CDR has the additional task of recovering a spectral component at the clock frequency.

There are two main types of phase detectors in CDR, linear and bang-bang, referencing to the relation between their output and the phase error. Their functionality along with both strengths and weaknesses will be discussed below.
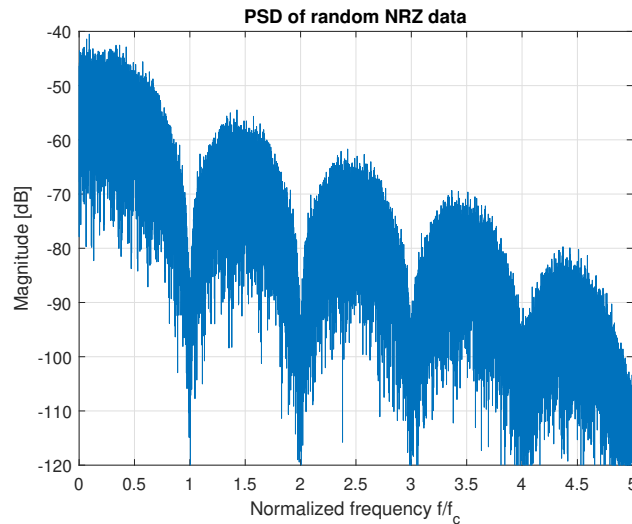
Figure 2.10: Example of the PSD of random NRZ data.

2.3.1. Hogge Phase Detector

The linear phase detector (PD), also known as the Hogge PD [31] consists of a positive edge-triggered D flip-flop (DFF), a negative edge-triggered DFF and two XOR gates. Its operation can be explained using Fig. 2.11: FF_1 samples the incoming data on the receiver clock resulting in signal B. B is then again sampled, now on the negative clock edge. As a result, signal A is delayed by half a period from signal B. Now, by performing two XOR operations with signal B, one on the incoming data and one on signal A, we obtain two pulses Y and X. $(Y - X)$ is now proportional to the phase error. The gain of this phase detector is linear from $-\pi$ to π and described as:

$$K_{PD} = \frac{1}{\pi}\alpha \tag{2.4}$$

Where α is the transition density of the incoming data, for random data $\alpha = 0.5$. Since the PD is linear, phase domain analysis of such a system is quite straightforward. In the locked state, node B contains the retimed data which is used as output of the CDR, a very useful property of the PD.
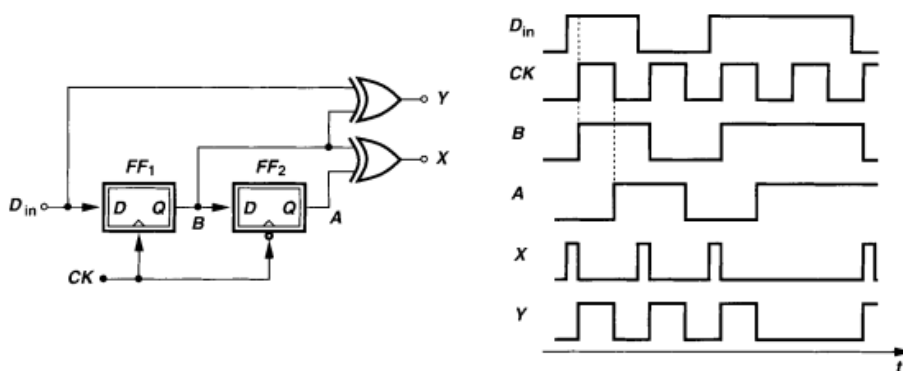
Figure 2.11: Hogge PD with internal waveforms [1]

In practice, there are a number of non-idealities and properties that limit the effectiveness. Firstly, the PD outputs are pulses at RF frequencies. This requires the connected charge pump to operate at gigahertz rates. Generating the pulses themselves also becomes difficult when the data rate is in the

tens of gigabits per second where pulsewidths are half of the bitperiods. Secondly, the propagation delays inside the PD result in skew. CK-to-Q delay in FF_1 results in a ΔT between B and the input data which widens the Y pulse by ΔT. At significantly high speeds, this delay becomes a relevant fraction of the clock period and as a result, the locking point will move to compensate for the extra pulsewidth, degrading phase margin and jitter tolerance. The delay can be compensated for inside the PD by placing a delay between the input data and the XOR logic gate that mimics the skew. Near the locking point, there is barely any activity in the Hogge PD and as such it suffers less from data-dependent jitter compared to its main counterpart.

A typical implementation of this PD is shown in Fig. 2.12. As mentioned, the PD output pulses are converted to current using a charge pump which feeds into the loop filter. This control voltage then tunes the VCO and consequently the CK in the PD, completing the loop.
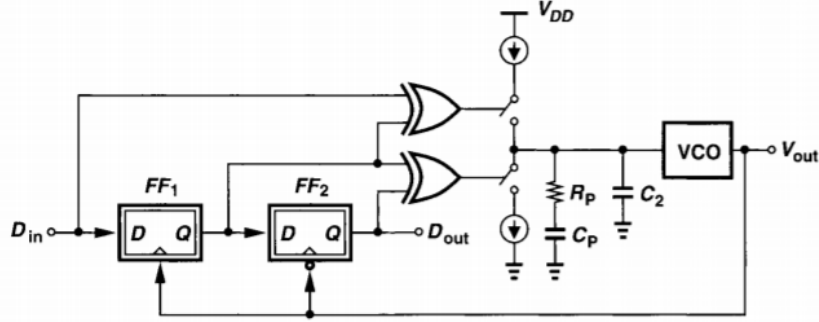
Figure 2.12: Typical CDR implementation with Hogge PD [1].

2.3.2. Bang-Bang Phase Detector

The bang-bang PD, also known as the Alexander PD, uses samples at a consecutive rising edge, falling edge and rising edge to detect if the clock is either leading or lagging compared to the data or, if there is no data transition at all. As its name suggests, the PD only provides information on the sign of the phase error not the magnitude [32].

Its operation is explained from Fig. 2.13 as follows: the incoming data is sampled by FF_1 on the
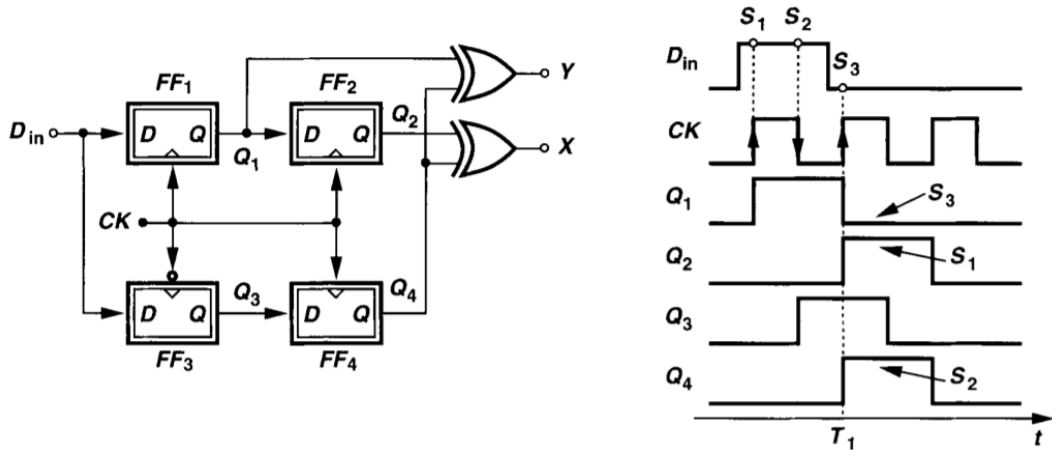
Figure 2.13: Alexander PD with waveforms [1].

rising clock edge, S_1, and by FF_3 on the falling edge, S_2, resulting in waveforms Q_1 and Q_3. Both signals are then sampled on the rising edge again, S_3, resulting in waveforms Q_1 to Q_4. Now, we consider two general cases, either S_2 leads the data edge or it lags behind.
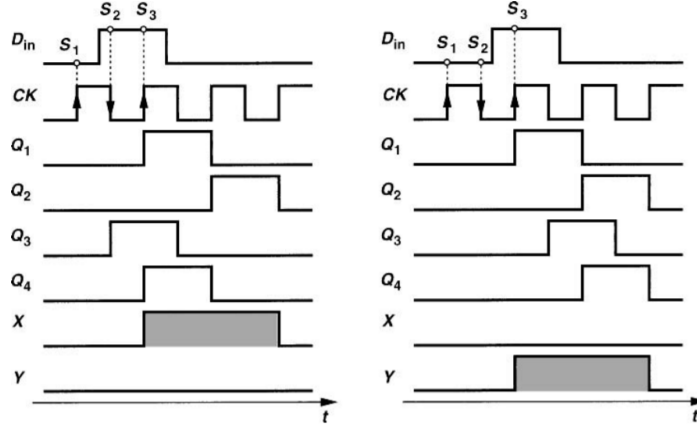
Figure 2.14: Alexander PD waveforms for early and late scenarios [1].

Output Y is generated by an XOR operation between Q_1 and Q_4, while X is generated from Q_2 and Q_4. Consider the first scenario in Fig. 2.14 where S_2 is lagging behind the data transition. In this case, both Q_1 and Q_4 sample high on S_3 while Q_2 is sampled one period after S_3. As a result, X will be high when lagging behind while Y will remain low. Hence, the output $(X - Y)_{avg}$ is positive.

In the second scenario, we can observe that the sampling point S_2 is leading the data transition. Due to this leading edge, the bit is sampled one clock period later in Q_3 and Q_4 compared to the lagging scenario. Consequently, X will remain low while Y becomes high. The output $(X - Y)_{avg}$ is negative.

Lock is achieved when S_2 coincides with the data transition, the output $(X - Y)_{avg}$ is zero. In this state, FF_3 and FF_4 sample the input at mid-rail resulting in meta-stable outputs Q_3 and Q_4. This then propagates to the XOR gates, which at high speeds are typically realized using Gilbert cells. If one of the two differential inputs is small and the other large, the resulting differential output is near zero. As a result $(X - Y)_{avg}$ will also be small.

In theory, this phase detector exhibits high loop gain and loop activity around the locking point due to its bang-bang nature. Loop dynamics are also more complex as it is a nonlinear component. However, there is a key advantage that makes this PD variant more popular compared to its linear counterpart: the V/I stage following the PD does not have to operate at RF, greatly easing design. It only has to sense the average of the output and as a result, can operate at high data rates.

2.4. State-of-the-Art Techniques

Literature aims to increase the speed and reduce power consumption using various circuit techniques. These techniques are mainly focused on the phase detector. This section will examine the improvements made over the classical phase detectors and some of the shortcomings therein.

To overcome the obstacle of high-speed latches for the linear PD in the tens of gigabit range, a mixer-based linear PD was proposed in [11]. The described architecture of the 20 Gb/s CDR is shown in Fig. 2.15a. Edge detection is performed by 4 inverters and an XOR gate which generates a pulse of roughly $T_{CK}/2$ for each data transition. Applying both this pulse and the clock to a mixer results in an output that is proportional to the phase difference. This output, while proportional, does not have a reference pulse attached to it and as such the current injected into the loop filter would not be zero during long runs. This architecture supplies the reference pulses by applying the $\overline{CK}$ to the charge pump. As can be seen in Fig. 2.15b, combining both the CK_{out} signal and the mixer output results in a zero net current when no transition occurs and, an up-down current proportional to the transition when it is present. This design was fabricated in 90-nm CMOS technology and, combined with the high data rate, required a circuit level implementation of current-mode logic with inductive peaking to achieve sufficient bandwidth. As such, the power consumption of this design is one of the highest reported at

154 mW, but with a low recovered jitter of 480 fs_{rms}.

(a)   (b)
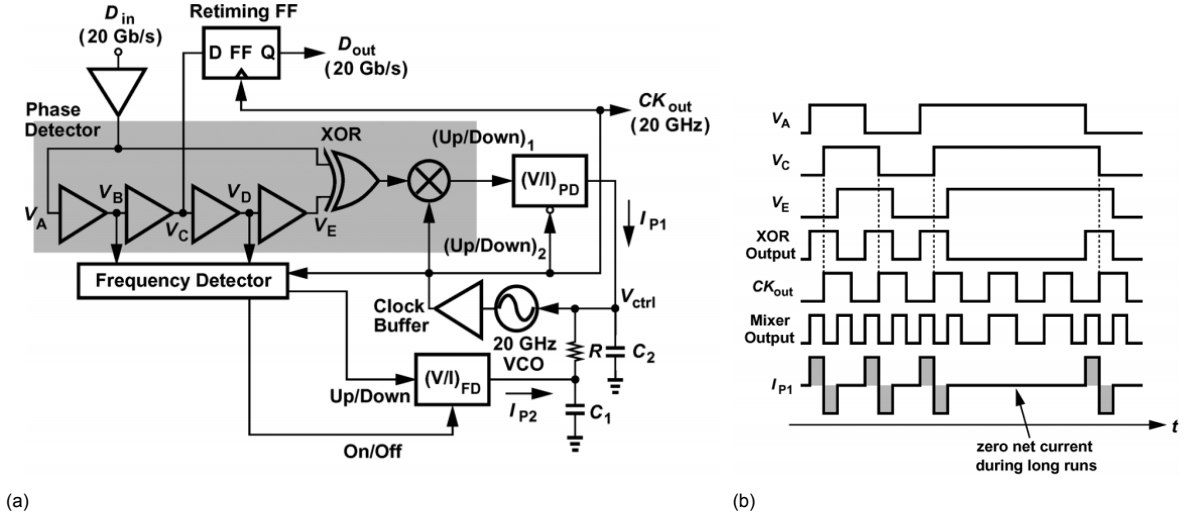
Figure 2.15: Proposed architecture for a 20 Gb/s linear CDR in [11].

The power consumption distribution is 65 mW in the PD and 66 mW in the VCO and clock buffers with an additional 23 mW in frequency tracking. In the PD, the delay chain draws 24 mW while the actual mixer consumes less than 5 mW. It should be noted that the pulse generator consumes far more power than the mixer itself. The architecture also puts a significant load on the clock as it has to drive the mixer, frequency detector, V/I and retiming FF. Such a clock buffer adds to the power budget. As can be deduced from the described design, more power efficient circuit techniques are required to implement linear CDR at such a speed.

One of the highest reported bit efficiency, described by data rate/power, is featured in [33]. It achieved a speed 25 Gb/s with a power consumption of only 5 mW by using charge-steering logic. Derived from the aforementioned current-mode logic, charge-steering logic operates through steering charge instead of current. A standard CML cell and charge-steering logic cell are shown in Fig. 2.16. Current-mode logic has a constant current I_T flowing through the cell at all times. The charge-steering logic cell, on the other hand, only draws current during a brief moment after the clock switches to charge the tail capacitor. This difference in current draw between the two cells results in a roughly 4× lower power consumption in the charge-steering cell.
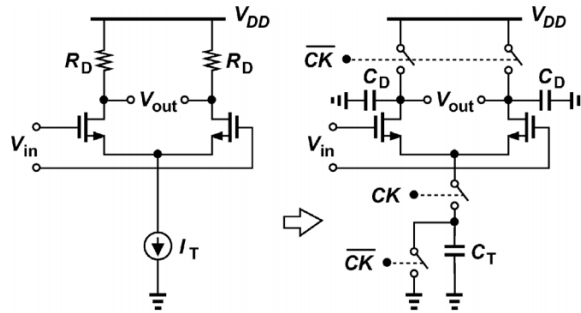
Figure 2.16: Current-mode logic and charge-steering logic comparison as described in [33].

Its operation is illustrated in Fig. 2.17. When CK is low nodes X and Y are reset to V_{DD} and C_T is discharged. When in evaluation C_T is connected to the tail node and C_D is disconnected from V_{DD}, con-

sequently V_P increases due to the differential input until both M_1 and M_2 turn off. There are, however, some drawbacks of using this logic style. Firstly, it requires a rail-to-rail clock to function properly. This requires extra attention when designing as both the clock generation and logic have to be optimized simultaneously [33]. A reduced clock swing will introduce ISI in the latches. Secondly, the output of the logic is return-to-zero (RZ) instead of NRZ owing to the two phases in its operation. This can be both hinder or be beneficial depending on the design. The benefit relates to the reset phase of the latch, as it reduces ISI by resetting all nodes. However, in classical CDR design NRZ is used due to its lower required bandwidth and as such this logic requires extra circuitry to convert RZ back to NRZ.
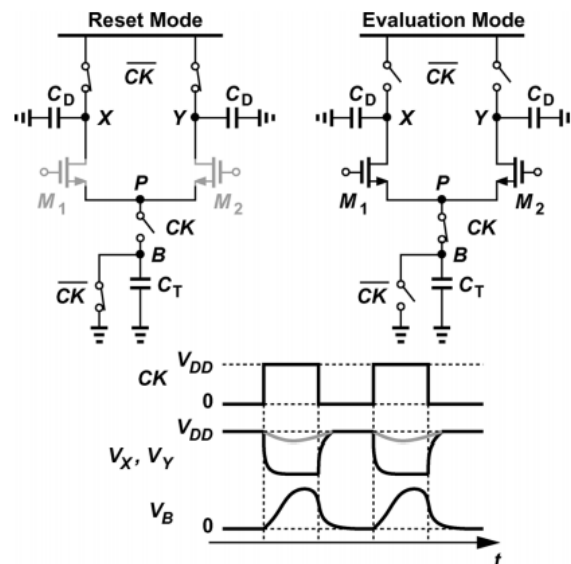
Figure 2.17: Current-mode logic and charge-steering logic comparison as described in [33].

A paper by Kong [9] proposes a new phase detector using a switched capacitor. The topology is shown in Fig. 2.18. This switched capacitor setup is known as a voltage sampling infinite impulse response (IIR) discrete-time low-pass filter when swapping D_{in} and V_{CK}. Its transfer function is described by the ratio of its capacitors and the sampling frequency, both of which are highly accurate in current processes. In this case, the sampling frequency is determined by the 'clock' period of the random data. On average, the transition density of a random NRZ signal is 0.5, meaning that the odds of a data transition occuring between two given bits is 50%. One 'clock' period requires both a one-to-zero and zero-to-one transition, which, on average, occurs 25% of the time. Consequently, the effective sampling frequency is 1/4th that of the bitrate: $1/(4T_b)$. Further analysis also shows that this PD produces,

Figure 2.18: Proposed bang-bang PD [9].

theoretically, no ripple in the control voltage when locked. This result enables a high bandwidth which will result in high jitter tolerance without compromising in data-dependent jitter generation. Fig. 2.19 shows the complete CDR architecture. The high bandwidth permits the use of a ring-VCO instead of a conventional LC-VCO while maintaining reasonable jitter generation since the oscillator phase noise is heavily suppressed, saving considerable area and power. Additionally, no pulse generator is required in the PD, reducing power consumption further. There are two buffers present at both VCO outputs to

shield it from the data transitions, that consume 33% of the total power.

The use of a first order loop also results in static phase offset under locked condition which dependents on PVT, this is partly compensated for by selecting different phases in the ring-VCO for feedback and retiming. The static phase offset can be resolved by moving to a Type-II structure such as in [34]. It does, however, reveal a second issue, the locking point of this PD even in a Type-II structure only has a 0.25 UI phase difference whereas a 0.5 UI shift is desired for optimal retiming. A solution such as using a delayed phase from the input pulse generator, as was done in Fig. 2.15a, is not possible since there is no longer a pulse generator.
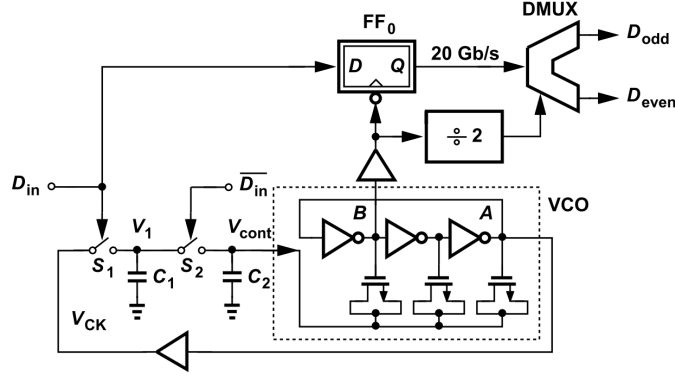
Figure 2.19: Proposed inductorless CDR structure [9].

The discussed techniques have shown areas where considerable power can be saved. Charge-steering can be used to reduce the time where a circuit is actively drawing current and the lack of a pulse generator at RF speeds saves an additional high power block.

2.5. Phase Rotation
The drawback of the phase detection method in Fig. 2.19 is the locking point. A mixer operates with a phase shift of 90 degrees between inputs to generate its output. This is not an issue in regular PLLs as the exact phase shift between reference and clock is not important. In CDRs, it is desired to sample at the center of a bit, where one has the highest SNR and the lowest BER. In order to bridge the gap between the mixer's phase shift and this optimum sampling point, [9] utilized different phases in a ring-VCO to achieve a roughly 180 degree phase shift between incoming data edges and the retiming clock, varying by 90 degrees over PVT. [11] makes use of a delayed phase inside the pulse generator at the input of the mixer, such a pulse generator has a high power consumption at RF speeds and is therefore not preferred. [35] has an LC-VCO in combination with a Self-Biased-PLL to generate a multiphase clock for a PAM4 CDR. Such a solution requires an additional loop with a ring-VCO and consequently has worse phase noise performance and requires additional power of buffers.

2.6. Proposed Architecture
The observations made in this chapter now bring us to the proposed architecture in Fig. 2.20.

Clock recovery is achieved with the use of a Type-II PLL, it has no steady state offset and tracks both frequency and phase errors. A proposed phase detector based on charge-sampling will receive complementary input data and oscillator waveforms to measure the phase error. It produces an output that is linearly related to the phase error. To reduce power, the PD topology does not require a pulse-generator to operate. The phase detector output voltage is subsequently converted to a current by a V/I stage. The phase detector, unlike the classical 'linear' type, allows for a low frequency V/I stage. The current then feeds into the loop filter, where the resulting voltage controls the VCO frequency. The clock recovery loop is fully differential to minimize disturbance from the supply.

The advantage of the phase detector is its power consumption, low jitter and low frequency output. However, it will lock the input data and VCO with a static phase offset of roughly 0.25 UI. Hence, an additional 0.25 UI shift is introduced to align the clock. The VCO signal is fed to a buffer with tunable delay that calibrates said delay to the desired 0.25 UI over PVT variations with the use of an additional feedback loop. The system then retimes the input data using the buffer output.

To enable testing and verification of the design, the core circuit requires several buffers and drivers. The incoming data is single ended and terminated to 50Ω. While the PD utilizes complementary data, it is preferred to generate this on-chip to prevent the need for accurate cable matching during measurement. Consequently, the single ended data is converted to complementary near the PD. To measure the performance of both the recovered clock and data, two drivers are used. The recovered clock is divided by 4 to allow use of more readily available test equipment. The recovered data is fed to a 50Ω differential driver to allow for testing while limiting the impact on the on-chip supply.
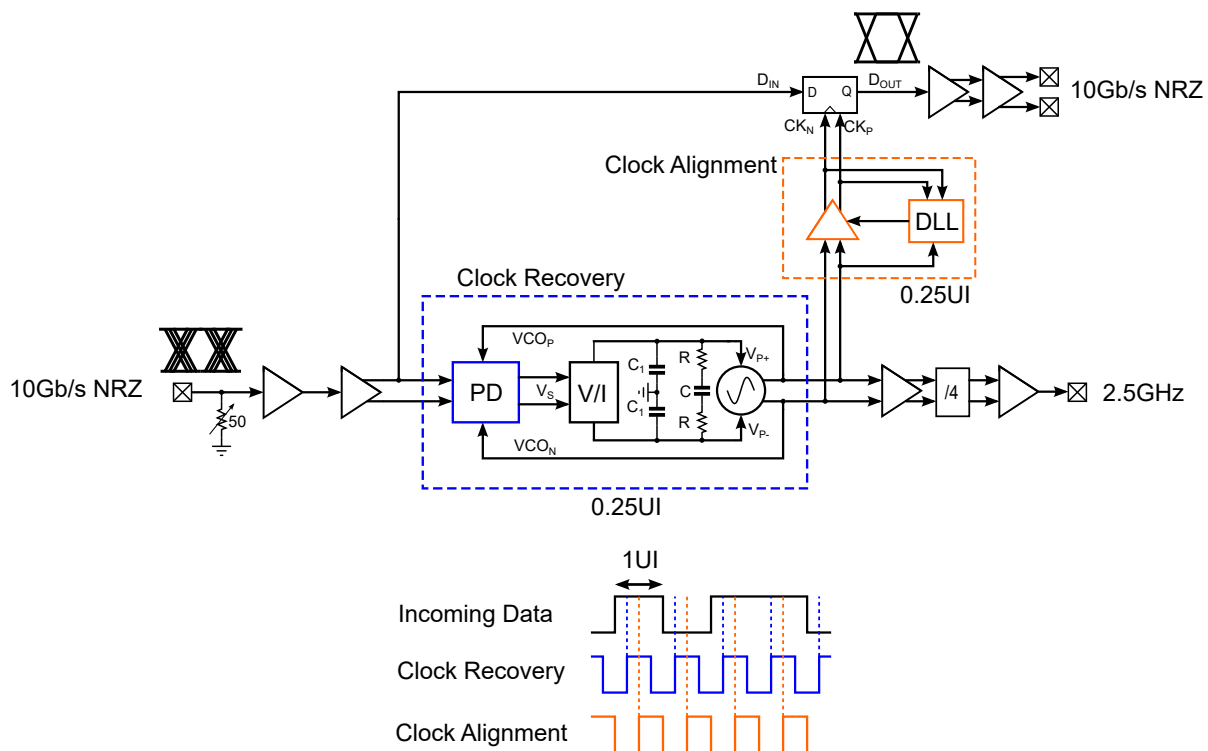
Figure 2.20: Block diagram of the proposed CDR.

3

System-Level Modeling and Simulation

This chapter will derive the required loop parameters to achieve the specifications listed in Chapter 1. The dynamics of the proposed architecture of Fig. 2.20 are evaluated in the S-domain of the phase transfer. Based on this model, different transfer functions for jitter transfer, tolerance, and generation are derived along with the phase margin. Non-idealities in the model are also discussed. Additionally, the phase noise contribution of each block is also derived.

3.1. Loop Transfer Modeling

The phase-domain model of the proposed clock recovery architecture is shown in Fig. 3.1, where $K_{V/I}$ is the gain of the V/I stage and $F_{LPF}(s)$ the transfer function of the loop filter. Unlike PLLs, there is no multiplication factor N at the input due to the incoming data and VCO operating at the same frequency. This model is accurate for linear phase-detection methods. In the case of bang-bang phase detection, the model does not predict certain stability effects if the input jitter becomes too large [36].

Figure 3.1: Linear phase-domain model of the CDR loop.

With the use of the model we can describe the closed-loop transfer function by

$$H_{cl}(s) = \frac{H_{ol}(s)}{1 + H_{ol}(s)}, \tag{3.1}$$

Where the open-loop transfer $H_{ol}(s)$ is can be expressed as

$$H_{ol}(s) = K_{PD}(s) \cdot K_{V/I} \cdot F_{LPF}(s) \cdot \frac{2\pi K_{VCO}}{s} \tag{3.2}$$

$$= \frac{2\pi \cdot K_{PD}}{1 + s \cdot R_D C_S} \cdot K_{V/I} \cdot (R + \frac{1}{s \cdot C}) \left\| \frac{1}{s \cdot C_1} \cdot \frac{K_{VCO}}{s} \right. \tag{3.3}$$

To get a more practical insight into the system dynamics, C_1 will be ignored and K_{PD} will be considered a constant. Their impact will be commented on further in Section 3.1.3. The simplified system can be formulated into a canonical form

$$H(s) = \frac{s \cdot R K_{PD} K_{V/I} K_{VCO} + K_{PD} K_{V/I} K_{VCO}/C}{s^2 + s \cdot R K_{PD} K_{V/I} K_{VCO} + + K_{PD} K_{V/I} K_{VCO}/C} = \frac{2\zeta\omega_n s + \omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2} \tag{3.4}$$

Where the damping constant and natural frequency are described by Eq. 3.5 and Eq. 3.6, respectively.

$$\zeta = \frac{R}{2}\sqrt{K_{PD} \cdot K_{V/I} \cdot K_{VCO} \cdot C} \tag{3.5}$$

$$\omega_n = \sqrt{\frac{K_{PD} \cdot K_{V/I} \cdot K_{VCO}}{C}} \tag{3.6}$$

The bandwidth of the loop can be calculated using both the damping constant and natural frequency as

$$f_{-3dB} = \frac{\omega_n}{2\pi}\sqrt{((2\zeta^2 + 1) + \sqrt{(2\zeta^2 + 1)^2 + 1}}. \tag{3.7}$$

With the use of Eqs. 3.5, 3.5 and 3.7, we can make a few observations about the system.
The bandwidth is determined by both ω_n and ζ. However, there are constraints that set bounds on their values. The damping constant has a lower bound at $\sqrt{2}/2$ to ensure loop stability. For now, let us consider using ω_n to control the bandwidth. While it might be clear from Eq. 3.6 that the three main gain factors K_{PD}, $K_{V/I}$ and K_{VCO} each contribute equally, some can be set more independently than others.

For example, in a classical Hogge PD described in Chapter 2, the gain K_{PD} is fixed and hence cannot be controlled. In a more recently published work, K_{PD} is set at a relatively large value to suppress the phase noise of the V/I stage; lowering the gain will result in increased in-band phase noise and increasing the gain above 0.4 V/rad is not feasible due to limited supply voltages. For example, a phase error of π rad with a gain of 0.4 V/rad would require an output voltage of 1.2 V. With a supply voltage of 1 V, it would have to be differential and even still, require a single-ended swing of 0.6 V. Therefore, a practical target for K_{PD} is the highest achievable gain to lower the noise constraint of the subsequent stage and mostly unrelated to bandwidth.

The oscillator gain K_{VCO} also affects the loop dynamics. A higher sensitivity allows for more disturbances to be picked up, thereby degrading performance. In LC-VCOs, the value of K_{VCO} is determined based on the maximum tolerable frequency disturbance of the oscillator. E.g., if $K_{pushing} = 100\,MHz/V$, a 100 mV drop in V_{DD} changes the oscillator frequency by 10MHz. K_{VCO} should be at least 3x-4x higher than that, however, further increasing the varactor size becomes undesirable due to large 1/f noise upconversion.

Ring-VCOs, in contrast to LC-VCOs, can be controlled by tuning their supply voltage, achieving a much higher gain. Varactor-based tuning also results in a higher gain than in LC-VCOs due to the typically small total capacitance and footprint compared to its LC counterpart. Nevertheless, reported values for LC-VCO gain range from 10 MHz/V up to 1 GHz/V [33].

This leaves the gain $K_{V/I}$ of the V/I stage as a prime candidate to control the bandwidth. Phase noise contribution of this stage is suppressed by the phase detector which also allows for a relatively low power consumption.

Let us now return to the damping constant. It not only has the previously mentioned lower bound for stability, it also determines the jitter peaking of the system. To illustrate this, we derive the poles and zeros present in the transfer function of Eq. 3.4

$$\omega_Z = \frac{1}{RC} \tag{3.8}$$

$$\omega_{PL} \approx \frac{1}{RC} \tag{3.9}$$

$$\omega_{PH} \approx K_{VCO}K_{PD}K_{V/I}R \tag{3.10}$$

Where ω_Z is the zero, ω_{PL} the low frequency pole and ω_{PH} the high-frequency pole. The bandwidth can be approximated by this high frequency pole. The magnitude Bode plot is shown in Fig. 3.2. The

amount of peaking is determined by the ratio of ω_z and ω_{PL} and can be approximated [1] by

$$20\log(J_p) \approx \frac{2.172}{\zeta^2}, \tag{3.11}$$

where J_p is the magnitude of peaking observed in the transfer function. From this expression it becomes apparent that the damping constant directly determines the amount of peaking. In optical repeater standards, such as SONET, there is a maximum allowed peaking of only 0.1 dB to prevent excessive peaking when cascading multiple repeaters. In these systems, the lower bound on the damping constant is set by the peaking specification rather than stability. Since the work is not related to repeaters, the damping constant is set at 1.5 to ensure stability.

Figure 3.2: CDR transfer function model.

The required bandwidth of the system is determined by the jitter tolerance specification. The jitter tolerance can be modeled as

$$\phi_{err} = |\phi_{data} - \phi_{out}| \le \frac{h}{2} \ \text{UI} \tag{3.12}$$

If the phase difference between the incoming data and the retiming clock exceeds $h/2$ a bit error can occur. The value of h is related to a number of factors and will be discussed more in-depth later. Using $H_{cl}(s) = \phi_{out}/\phi_{data}$, the expression can be rewritten

$$\phi_{data} \le \frac{h}{2} \cdot \frac{1}{|1 - H_{cl}(s)|} \ \text{UI} \tag{3.13}$$

Jitter tolerance is defined with respect to peak-to-peak input data jitter and therefore expressed as

$$G_{JT}(s) = 2\phi_{data} = \frac{h}{1 - H_{cl}(s)} \text{UI}_{\text{pp}} \tag{3.14}$$

Similar to the original closed-loop transfer, the jitter tolerance can be formulated using the natural frequency and damping constant,

$$G_{JT}(s) = h \cdot \frac{s^2 + 2\zeta\omega_n s + \omega_n^2}{s^2} \text{UI}_{\text{pp}}. \tag{3.15}$$

The jitter tolerance function consists of two poles and two zeros. Both poles are located at the origin, creating a 40 dB/dec downward slope. The two zeros correspond to the two poles present in the closed-loop transfer function. As a consequence, a higher jitter tolerance requires a larger bandwidth. The factor h also impacts the jitter tolerance and sets the floor of out of band performance. Outside the loop bandwidth, no tracking is possible and as such the only protection against bit errors is the innate distance between the data transition and the resampling time t_s. As shown in Fig. 1.2. The value of $h/2$ is often assumed at 0.5 UI [1], [37]. The next section will discuss the estimation of h more in-depth.

3.1.1. Jitter Tolerance Floor

The value of h can have a considerable impact on the jitter tolerance while affecting virtually no other system metric. A higher h value will result in a higher jitter tolerance without increasing the bandwidth, meaning that a higher jitter tolerance can be achieved while maintaining the same jitter transfer. Let us now consider the setup in Fig. 3.3 to find an estimate of h. We define the retiming flip-flop's setup and hold time, T_{su} and T_{ho}, the deterministic jitter T_{dj} and random jitter σ_{rj} of the retiming clock and the fixed distance between data transition and clock transition T_m. The setup and hold time here are defined as the minimum time between the clock edge and data edge where the flip-flop's output is correct. The edges are considered ideal. The random jitter on the clock is unbounded and therefore needs a reference BER, in this design $10\mathrm{e}{-}12$, corresponding to $14\sigma_{rj}$.


Figure 3.3: A retiming scenario considering setup and hold time, random and deterministic clock jitter and intrinsic timing margin.

No bit error will occur as long as the clock edge does not cross into either the setup or hold region, based on this we can formulate an expression for h,

$$h = 2 \cdot \min(T_m - T_{su}, T_m - T_{ho}) - T_{dj} - 14\sigma_{rj} \quad \mathrm{UI_{pp}} \tag{3.16}$$

From Eq. 3.16, it becomes clear that to achieve the maximum h, one needs T_m to be centered between the setup and hold time. If one is not equal to the other, T_m should be moved accordingly or jitter tolerance performance will degrade. The location of T_m is often fixed depending on the architecture. In the classical Hogge and Alexander CDRs, it is located directly at the center as was discussed in Chapter 2. Type-I loops, on the other hand, possess a static phase offset that is also dependent on PVT [9] meaning that T_m and consequently h has a significant dependence on PVT.

3.1.2. System Bandwidth

Using the established model, we can now find the required loop bandwidth to achieve the targeted jitter tolerance of $1\,\mathrm{UI_{PP}}$ at 10MHz offset. Fig. 3.4 plots the modeled jitter tolerance for different loop bandwidths and values of h. For the first plot, h is fixed at $0.4\,\mathrm{UI_{PP}}$. It can be observed that to meet the target specification marked in red, a bandwidth of over 30MHz is required in combination with an h of at least $0.4\,\mathrm{UI_{PP}}$. The second plot shows the jitter tolerance for a loop with 30MHz bandwidth and different values of h, it can be seen that for a relatively large value of h, the specification is met, while for a lower value it significantly undershoots. In order to have a comfortable margin, a loop bandwidth of 40 MHz is targeted. This allows the specification to be met with $h = 0.3\,\mathrm{UI_{PP}}$.

3.1.3. Impact of Phase Detector Pole

The previous analysis omitted additional factors such as frequency dependence of K_{PD}. In this subsection, we will examine the effect of an extra pole on the various transfer functions. Let us now consider Eq. 3.17 where an additional pole is present in the phase-detection gain

$$K_{PD}(s) = \frac{K_{PD}}{1 + R_D C_S}. \tag{3.17}$$

Figure 3.4: Jitter tolerance for different values of the loop bandwidth and h.

The location of the pole relative to the loop bandwidth directly affects the phase margin of the system and its degradation can be approximated by

$$-\tan^{-1}(\omega_u R_D C_S), \tag{3.18}$$

where ω_u notes the frequency where $|H_{ol}(s)| = 1$. The effect of this pole on the transfer function is shown in Fig. 3.5. It introduces additional jitter peaking near the cut-off frequency if the pole is too close to the unity-gain frequency. This is undesired, as such peaking will increase the jitter of the recovered clock unnecessarily. To minimize this degradation, the pole must be at least 5 times higher than the loop bandwidth.
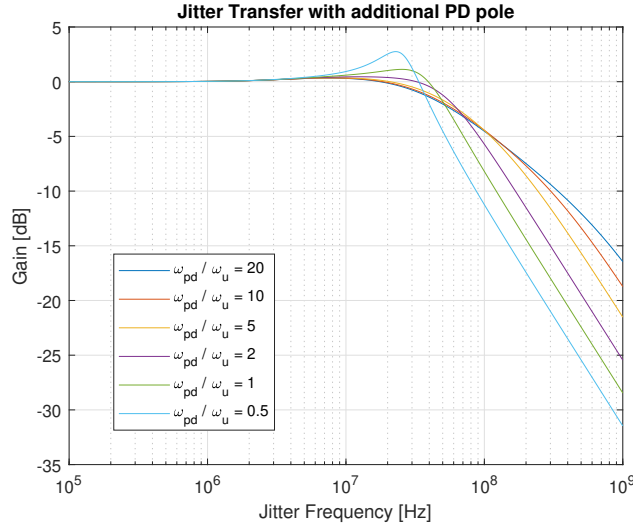
Figure 3.5: Jitter transfer for various ratios of the phase detector pole and bandwidth.

The jitter tolerance is also affected by the pole, as illustrated in Fig. 3.6, introducing a dip at its corner frequency. If this dip is too severe, the system will violate the mask. Therefore, the pole location must also be considered with respect to this metric. Similar to the jitter transfer case, the dip only becomes prevalent when the ratio between poles drops below a factor of 5.

In short, the degradation due to an additional pole in the phase detector is limited as long as $\omega_u <<$ ω_{pd}. While a factor of 5 might seem enough based on the previous discussion, it does degrade the phase margin by $11°$. Therefore, in a realistic design, a factor of 10 to 20 is desired where the phase margin degradation is limited to $5.7°$ and $2.9°$ respectively. If this condition holds the simplification in Eq. 3.4 will also be reasonably accurate.
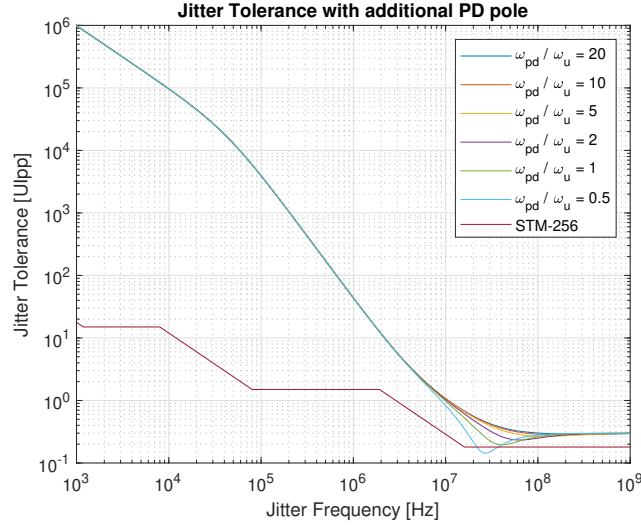
Figure 3.6: Jitter tolerance for various ratios of the phase detector pole and bandwidth.

3.1.4. Impact of loop latency

An extra pole in the phase detector gain is not the only additional source of disturbance to consider. Let us examine how the loop dynamics change if there is a delay of τ. It can originate due to the finite bandwidth of paths connecting each block in the model. We add the additional delay to Eq. 3.1 which now becomes

$$H_{cl}(s) = e^{-s\tau} \frac{H_{ol}(s)}{1 + e^{-s\tau}H_{ol}(s)}, \tag{3.19}$$

where τ is the delay in seconds. Fig. 3.7 shows the effect of this added latency. If the loop bandwidth is high, in this case 40 MHz, the CDR transfer function starts to show peaking, even with only a few nanoseconds of loop delay. Careful attention must be paid to prevent this effect in the loop.
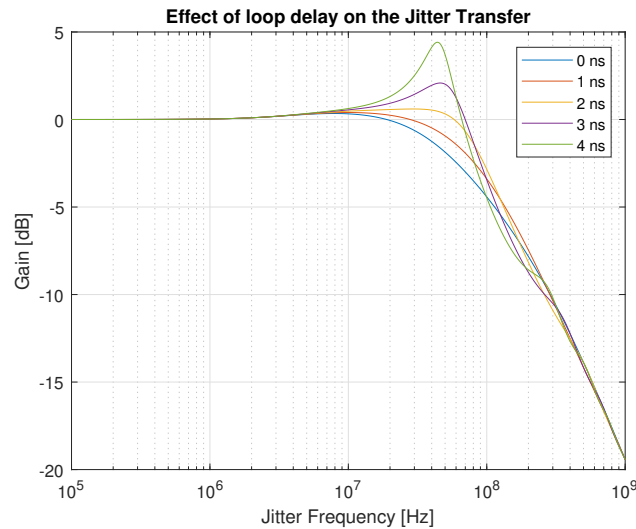
Figure 3.7: Effect of loop delay on the jitter transfer in a system with 40 MHz bandwidth.

A similar plot is shown for the jitter tolerance in Fig. 3.8. A similar dip appears in this case as with the phase detector pole, only with the addition of ringing. Again, both the jitter transfer and jitter tolerance require a low loop latency, ideally limited to less than 1 nanosecond.

Figure 3.8: Effect of loop delay on the jitter tolerance in a system with 40 MHz bandwidth.

3.2. Phase Noise Contributions

The linear model in Fig. 3.1 also contains the various noise sources present in the system. Each block has a specific transfer to the output which is derived in this section. The input buffer, PD, V/I, LPF and VCO all contribute noise to the model. Additional contributions consist of the VCDL and its analog front-end. The total phase noise of the recovered clock is
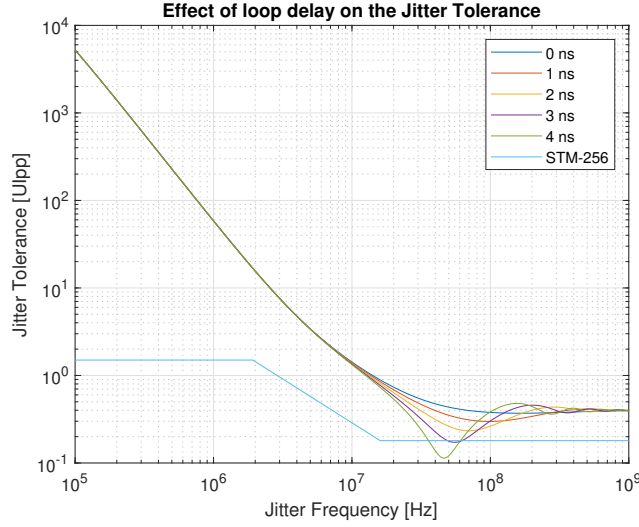
$$\phi_{out,n}^2 = |H_{cl}(s)|^2 \phi_{data,n}^2 + |H_{cl}(s)|^2 \frac{V_{pd,n}^2}{K_{PD}^2} + |H_{cl}(s)|^2 \frac{i_{v/i,n}^2}{(2K_{V/I}K_{PD})^2} \tag{3.20}$$

$$+ |H_{lpf}(s)|^2 v_{lpf,n}^2 + |H_{osc}(s)|^2 \phi_{vco,n}^2 + \phi_{dcdl,n}^2 + \phi_{div,n}^2, \tag{3.21}$$

where $\phi_{data,n}$, $\phi_{vco,n}$, $\phi_{div,n}$ and $\phi_{dcdl,n}$ are the phase noises of the input data, VCO, divider and DCDL; $V_{pd,n}$ and $V_{lpf,n}$ the voltage noise spectrum of the phase detector and loop filter; $i_{v/i,n}$ the current noise spectrum of the V/I stage; $H_{osc}(s)$ the transfer function of the oscillator phase noise and $H_{lpf}(s)$ the loop filter transfer function to the output.

As expected, the input data jitter follows the closed-loop transfer function and experiences a low-pass characteristic. The phase detector's voltage noise is suppressed by its own gain when referred to the input and consequently, a higher gain will result in a lower contribution. Similarly, the current noise of the V/I stage is suppressed not only by the phase detector gain but also by its own transconductance. High phase detection gain is therefore very beneficial. The loop filter voltage noise has its own unique transfer to the output which is derived from the phase noise model as

$$H_{lpf}(s) = \frac{s^2}{s^2 + 2\zeta \omega_n s + \omega_n^2}. \tag{3.22}$$

It has a bandpass transfer characteristic. The oscillator's phase noise spectrum sees a high-pass transfer function to the output described by

$$H_{osc}(s) = \frac{1}{1 + H_{ol}(s)} = \frac{s^2}{s^2 + 2\zeta \omega_n s + \omega_n^2}, \tag{3.23}$$

suppressing the high phase noise at low frequency offsets. With the use of these equations, we can find an estimate for the phase noise spectrum of the recovered clock.

The targeted rms-jitter for the design is $150\,\text{fs}_{\text{rms}}$, with the bandwidth fixed at 40 MHz it is possible

to calculate the in-band phase noise required to meet the specification. The relation between phase noise and jitter is described by [1]

$$\sigma_{\phi,rms} = \sqrt{2 \int_{f_{min}}^{f_{max}} \mathcal{L}(\Delta f) \cdot d(\Delta f)},$$ (3.24)

where $\mathcal{L}(\Delta f)$ is the phase noise spectral density at offset frequency (Δf) and f_{min} and f_{max} denote the integration window of interest. In published CDRs $f_{min} = 100\,Hz$ and $f_{max} = 1\,GHz$ to make sure that all relevant phase noise is captured.

For a flat in-band spectrum with a bandwidth of 40 MHz at 10 GHz, we can calculate the required in-band phase noise as

$$\mathcal{L}(\Delta f) = 10 \log_{10}(\frac{\sigma_{rj,rms}^2}{2} \frac{f_c^2}{f_{bw}^2}) = -119.5 \, \text{dBc/Hz}.$$ (3.25)

Based on this target, we can calculate the maximum allowed phase noise of the various blocks in the system.

Starting with the VCO, we can find the required phase noise at 1 MHz offset by first calculating the suppression a 40 MHz loop provides. At 40 MHz, the oscillator phase noise should be -119.5 dBc/Hz. Extrapolating to 1 MHz, we find $-119.5 + 20\log_{10}(40\text{MHz}/1\text{MHz}) = -87.5 \, \text{dBc/Hz}$. Leaving an additional margin of 3 dB results in a required phase noise of -90 dBc/Hz at 1 MHz offset.

On the low-pass side of the transfer, the phase detector and V/I stage should also remain below -119.5 dBc/Hz. The loop filter's contribution can be directly calculated, the voltage noise from the resistors is

$$v_{r,n}^2 = 2 \cdot 4kTR.$$ (3.26)

Even for a resistance of 10 kΩ, the resulting PSD is only -150 dBc/Hz at room temperature and lower at cryogenic temperature, contributing negligible jitter.

This leaves the phase detector and V/I stages as main contributors for in-band phase noise. Their combined contribution must be -119.5 dBc/Hz. The V/I can have relatively higher noise due to suppression by the phase detector gain. Other auxiliary components, such as the IO drivers, should also contribute negligible noise compared to the loop components.

3.3. Summary

In this chapter, the parameters for the clock recovery loop were derived. To meet the jitter tolerance specification, the loop must have a bandwidth of 40 MHz, allowing it to be achieved with an h of 0.3. To ensure stability and minimize peaking, the targeted damping constant is set at 1.5 to 3. Tunability is included to allow for optimization in the actual measurement. While an h of 0.3 might suffice with the targeted bandwidth, we aim for a higher value to ensure the jitter tolerance specification is met. The ideal phase detector gain is the highest feasible at transistor-level, which is roughly 0.3 V/rad. The combination of V/I gain and VCO gain should be enough to reach the bandwidth. As discussed earlier, $K_{V/I}$ is more suited for tunability compared to K_{VCO}. However, setting $K_{V/I}$ arbitrarily high will needlessly increase power consumption. Based on the aforementioned parameters, the loop filter values can be calculated using Eq. 3.5 and 3.6. The derived parameters are summarized in Table 3.1.

Additional considerations include the location of a phase detector pole and loop latency, which must both be limited to prevent additional peaking in the phase noise spectrum.

Table 3.1: Targeted loop parameters for the CDR.

Parameter	Target
Bandwidth	40 MHz
Damping Constant	1.5-3
Phase Margin	65°
h	0.4
K_{PD}	high
K_{VCO}	100 MHz/V
$K_{V/I}$	100-150 uS
C_{lpf}	2 pF
R_{lpf}	22-30 kΩ
$C_{1,lpf}$	$C_{lpf}/20$

4

Analog/RF Design

This chapter covers the transistor-level design (schematic and physical layout) of the analog and RF blocks of the CDR. These include the phase detector, V/I stage, VCO, retimer, DCDL, input buffer, and output driver, as well as top-level layout.

4.1. Phase Detector

The phase detector is the error detection block of the CDR. It takes the incoming random data and the RF clock as input and generates a voltage proportional to their phase difference at its output. The uniqueness of this task makes it a defining feature in a CDR. The analysis and design of the proposed phase detector regarding gain, phase noise, power consumption and jitter are presented in this section.

4.1.1. Edge Detection

Edge detection is a required operation to extract the frequency of the incoming data. Typically this detection is done using a pulse generator consisting of an XOR operation on the data and a delayed version of it. However, at 10 GHz, this will result in high power consumption, even with relatively low capacitance. For example, a pulse generator in a PLL for a 100 MHz reference has to consume 100 μW to meet its phase noise requirement [38], by extrapolating this to 10 GHz, it would consume two orders of magnitude more power. An alternative method based on complementary switching rather than pulse generators has been used in IL-PLLs [39] to circumvent the need for a pulse generator. This structure can be modified to accommodate random data as shown in Fig. 4.1. When no transition is present, the switches are fixed and V_{tail} sees both capacitors C_P and C_{tail} to ground. When a data transition occurs, the polarity of C_P is flipped and as a result V_{tail} is lowered according to a capacitive ratio.


Figure 4.1: Complementary switched injection for random data.

The injected charge can be described by:

$$Q_{inj} = C_P \cdot V_{tail} \tag{4.1}$$

The voltage drop ΔV due to the injection can then expressed as:

$$\Delta V = 2V_{tail} \cdot C_P/(C_P + C_{tail}) \tag{4.2}$$

By sizing $C_P >> C_{tail}$, the tail node is almost fully discharged. The symmetrical switches operate in the same manner for both low-high and high-low transitions and can therefore detect all data transitions. Depending on the slew-rate of the data, there is a brief moment where both D_{in} and $\overline{D_{in}}$ conduct and a short circuit current can flow from the tail to ground, illustrated in Fig. 4.2. However, this window is small (1-3 ps) and the on-resistance of both devices is still high during this time, consequently, the short circuit current is negligible. The increased threshold voltage at cryogenic temperature [3] further reduces this window, possibly eliminating it completely.


Figure 4.2: Short circuit window at room temperature and cryogenic temperature.

4.1.2. Phase Detection

The voltage drop can be used for phase detection by combining it with the charge-sampling principle [40]. In Chapter 2, the merit of using charge rather than current was discussed. Charge-sampling has been shown to achieve both low-power and low-jitter in PLL phase detection [38]. By combining the edge detection structure and the charge-sampling structure, we arrive at a charge-sampling based phase detector for random data, shown in Fig. 4.3.

Figure 4.3: Circuit diagram of the proposed phase detector.

Its operation principle can be understood using Fig. 4.4. When a data transition occurs, the tail node will be discharged turning on M_1 and M_2. The tail node remains low for approximately half a bit period before it is recharged, we will return to this in more detail later. In the locked state, charges Q_{SP} and Q_{SN} are equal, resulting in an output voltage V_S of 0 V. When there is a phase difference, the charges are no longer equal resulting in a non-zero differential output voltage, whose magnitude is proportional to the phase error.

During the sampling operation, the common-mode (CM) voltage of V_S will drop and should subsequently be reset to properly detect the next data transition. Resistors R_D are added to recharge the CM voltage of V_S. In order to prevent the main differential current from flowing through R_D, it is sized such that $R_D >> 1/(\omega_{VCO} C_S)$.

Figure 4.4: Conceptual operation of the charge-sampling based phase detection.

4.1.3. Gain

In order to find design trade-offs and gain more insight into its operation, we derive an expression for the phase detector gain, K_{PD}, with the help of Fig. 4.5.

Figure 4.5: Waveforms in the phase detector.

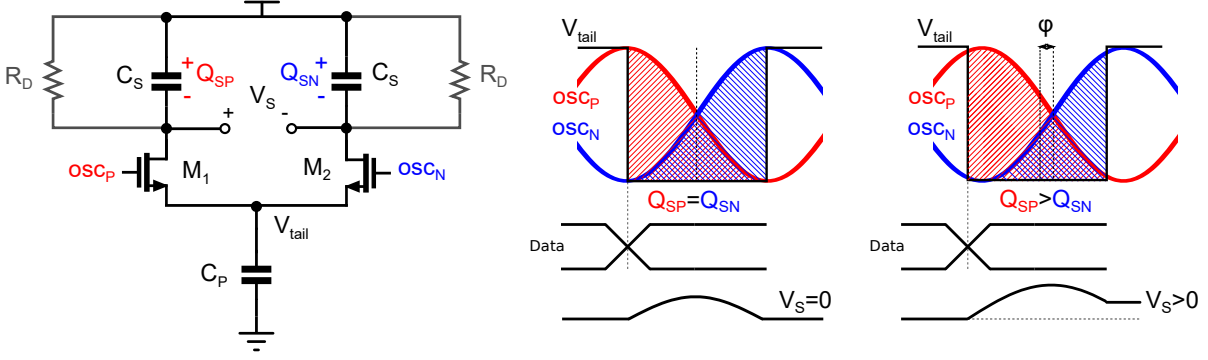First, we approximate the tail discharge by a periodic sampling function $p(t)$, that samples a continuous-time RF current $I(t) = G_M A_{VCO} \sin(\omega_{VCO} t + \phi)$ for $T_P = 0.5 T_{bit}$. The odds of a data transition occurring is α, 0.5, resulting in a train of pulses spaced by $T_{bit}/\alpha = 2 T_{bit}$

$$I_S(t) = G_M A_{VCO} sin(\omega_{VCO} t + \phi) \cdot p(t), \tag{4.3}$$

where G_M is the large-signal transconductance of $M_{1,2}$ and

$$p(t) = \begin{cases} 1, & -\frac{T_P}{2} + n \cdot \frac{T_{bit}}{\alpha} \leq t \leq \frac{T_P}{2} + n \cdot \frac{T_{bit}}{\alpha} \\ 0, & \text{otherwise.} \end{cases} \tag{4.4}$$

A phase error ϕ between the center of T_P and the VCO crossing gives an output voltage

$$\Delta V_S = V_{S,END}[n] - V_{S,INI}[n] = \frac{2}{C_S} \int_{-0.5T_P}^{0.5T_P} I_S(t) dt \tag{4.5}$$

$$= \frac{4 G_M A_{VCO}}{\omega_{VCO}} \cdot \sin(\omega_{VCO} t) \cdot \sin(\phi) \tag{4.6}$$

at the end of every pulse. Where $V_{S,INI}[n]$ is the voltage at the start of the pulse and $V_{S,END}[n]$ the voltage at the n-th data transition. Between the pulses, there is a discharge time $T_{DIS} = T_{bit}/\alpha - T_P$

where $p(t)$ is zero. After the pulse $V_{S,END}[n]$ discharges exponentially through R_D and C_S, lowering it to

$$V_{S,INI}[n+1] = V_{S,END}[n] \cdot e^{-k} \tag{4.7}$$

at the start of the next cycle. Here k is given by $T_{DIS}/(R_D C_S)$. Combining Eqs. 4.6 and 4.7, we arrive at an expression for $V_{S,END}[n]$

$$V_{S,END}[n] = \Delta V_S \sum_{k=0}^{n-1} e^{-n \cdot k} \tag{4.8}$$

with its steady state value being $\Delta V_S/(1 - e^{-k})$. In the steady state, V_S is a periodic function of T_{bit}/α, with its average value estimated as

$$\overline{V_S} \approx \frac{\alpha}{T_{bit}} \int_{0.5T_P}^{T_{bit}/\alpha - 0.5T_P} \Delta V_S/(1 - e^{-k}) \cdot e^{-t/(R_D C_S)} dt \tag{4.9}$$

$$\approx \frac{2\alpha G_M A_{VCO} R_D}{\pi} \cdot \sin(0.5\omega_{VCO} T_P) \cdot \frac{\sin(\phi)}{\phi} \tag{4.10}$$

$$\approx \frac{G_M A_{VCO} R_D}{\pi} \cdot \frac{\sin(\phi)}{\phi} \tag{4.11}$$

This derivation resembles the one in [38] and is valid under the assumption that $T_{DIS} >> T_P$. The gain is proportional to R_D which can be understood intuitively: a higher resistance will discharge the differential voltage more slowly, resulting in a higher average output voltage. R_D cannot be increased indefinitely to achieve higher gain, if it becomes too large, it will not recharge the common-mode enough to allow proper operation of M_1 and M_2, bringing them into the triode region near the end of T_P. Surprisingly, C_S does not influence the gain, as its contributions to the peak voltage and recharging cancel each other.

(a)   (b)
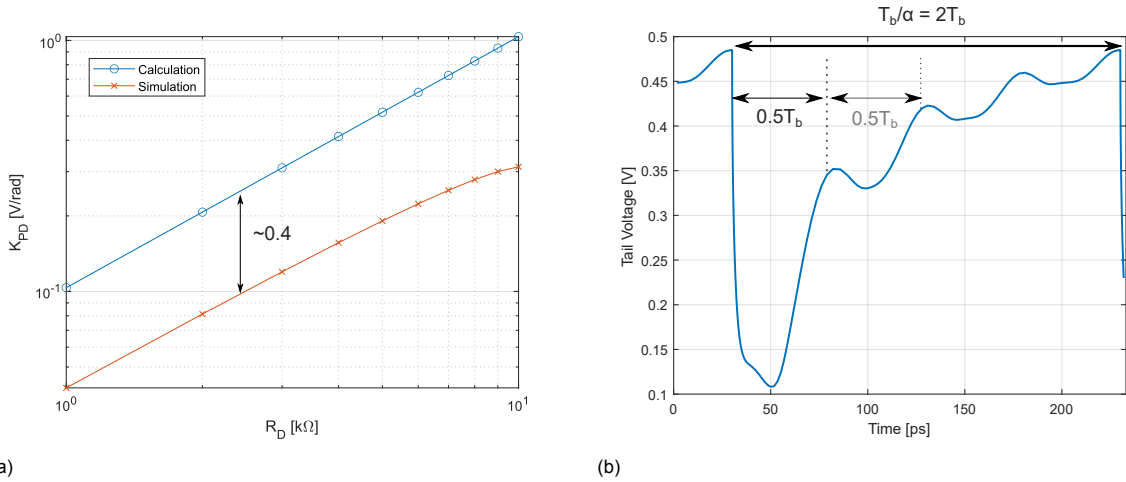
Figure 4.6: a) Comparison of the calculated K_{PD} with simulation and b) the simulated tail voltage.

Fig. 4.6a shows a simulation using $R_D = 3.5\,k\Omega$, $A_{VCO} = 0.45\,V$ and $G_M = 1.15\,mS$. We can observe that there is a scaling factor of roughly 0.4 between the calculated and simulated gain. The system has been simplified too much, neglecting the relevant dynamics. To improve it, we take a closer look at the actual tail voltage in Fig. 4.6b. We observe that T_P is indeed roughly $0.5T_{bit}$, however, the tail voltage does not rise high enough in the following periods to completely set $I_S(t)$ to 0. Instead, there will be more pulses, albeit with a lower G_M due to the decreased V_{GS}. Now, we change the model such that $p(t)$ consists of two sampling pulses with the magnitude of the second sample being a fraction of the first sample. ΔV_S can now be rewritten as

$$\Delta V_S = V_{S,END}[n] - V_{S,INI}[n] = \frac{2}{C_S}\left(\int_{-0.5T_P}^{0.5T_P} I_S(t)dt - (1-\beta)\int_{0.5T_P}^{1.5T_P} I_S(t)dt\right) \tag{4.12}$$

$$= \beta \frac{4G_M A_{VCO}}{\omega_{VCO}} \cdot \sin(\omega_{VCO}t) \cdot \sin(\phi), \tag{4.13}$$

where β is the ratio between the first and second pulse. Based on simulation, the value of β is roughly 0.48. The steady state value of V_S becomes

$$\overline{V_S} \approx \frac{\beta G_M A_{VCO} R_D}{\pi} \cdot \frac{\sin\phi}{\phi}. \tag{4.14}$$

With this extra factor added to the expression, we again plot the calculated and simulated gain over different phase detector parameters in Fig. 4.7. The estimation is reasonably accurate, with the remaining difference in gain due to a number of non-idealities. For example, V_S also discharges during the actual charge sampling, reducing gain. The function $p(t)$ also does not fully approximate the actual tail node, as there are more pulses after the second, decaying in magnitude. Of note here is that in Fig. 4.7c, we observe an optimum value for C_P. Using a larger capacitance severely degrades the gain, whereas a smaller value has less of an impact.
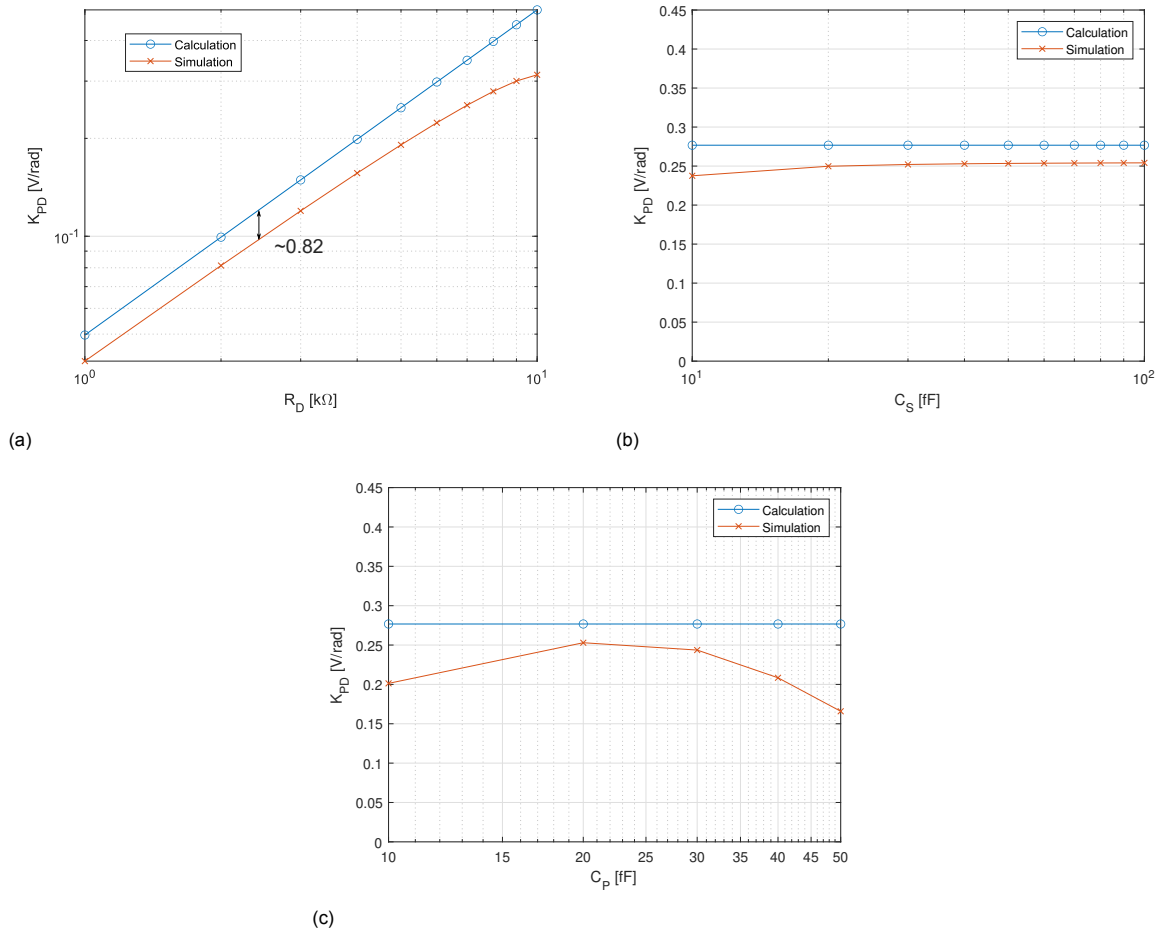

(a)

(b)

(c)

Figure 4.7: Phase detector gain plotted over a) resistance R_D, b) capacitance C_S and c) capacitance C_P.

4.1.4. Locking Point

Due to the previously discussed 'additional' sampling pulses, the locking point of the PD is not exactly $t_{transition} + 0.25 * T_{VCO}$, but slightly further at 27.8ps. The components within the phase detector

all experience mismatch, however, passive components such as resistors and capacitors are matched much more accurately compared to transistors. The differential pair M_1, M_2 will have a threshold voltage mismatch ΔV_{TH}. This translates into a current mismatch ΔI through the device g_m, resulting in a phase offset. When using small devices $(1.2\text{um}/40\text{nm})_{\text{W/L}}$ for $M_{1,2}$, the 3σ variation in locking point is ± 4 ps, which can degrade h by nearly 0.1 in the worst case. For this reason the area of $M_{1,2}$ was increased by a factor 16, resulting in a 3σ variation of ± 1.25 ps, degrading h by only 0.025.

4.1.5. Maximum Frequency

The phase detector is functional at 10 GHz, but is able to operate at much higher data rates. It can still produce a gain of 0.1 rad/V at 40 GHz in 40-nm by scaling down all devices. Fig. 4.8 shows a sweep of the gain at 40 GHz, displaying that the phase detector can be utilized at high speeds. The component values are listed in Table 4.1 and the incoming data was modelled as a square wave with 15 ps rise and fall times. $M_{1,2}$ are smaller which will increase the locking point offset over PVT, degrading h.

Figure 4.8: Simulated phase detector gain while operating at 40GHz.

Table 4.1: PD component values when simulated at 40Gb/s.

Parameter	Value
R_D	7kΩ
C_S	20fF
$M_{1,2}$	2.4um/40nm
S_{1-4}	4.8um/40nm
C_P	5fF
A_{VCO}	0.45V

4.1.6. Data Dependent Effects

The previous analysis only considered an input consisting of a '0011' repeating bit sequence. In reality, this data is random and the duration of T_{DIS} can vary from half a bit period to the maximum run length, i.e., $32T_b$. The overall gain will remain accurate, however, each individual cycle consisting of a charge sampling and discharging phase will have a different $\overline{V_S}$ as shown in Fig. 4.9. This variation leads to a random voltage ripple on the control line even when the system is locked, introducing additional jitter. The variation is low-pass filtered and therefore proportional to the bandwidth of the loop.

Figure 4.9: Illustration of the change in V_S due to long runs.

When the CDR is locked, there will be offset voltages when $\alpha \neq 0.5$. For the two most extreme cases, $\alpha = 1$ and $\alpha = 1/64$, the simulated peak-to-peak offset voltage is $\sim$15 mV. This voltage disturbance has a similar spectral density as the NRZ data itself and is suppressed by the loop filter. With a K_{VCO} of 100 MHz/V, this results in a simulated peak-to-peak jitter of 120 fs.

4.1.7. Common Mode Voltage
The common mode of V_S experiences a ripple during nominal operation due to charge-sampling and subsequent recharging of the V_{Sp} and V_{Sn} nodes. The amplitude of this ripple is proportional to C_S and can therefore be suppressed by increasing its size. However, C_S cannot be arbitrarily large as it will lower the frequency of the phase detector pole and reduce the phase margin. Consequently, with a C_S of 40 fF, the CM ripple amplitude is 75 mV. This ripple can propagate to the VCO in two ways: by being converted to a differential signal through the common-mode to differential gain, A_{CM-DM}, of the V/I, affecting the control line and, by continuing as a CM ripple through A_{CM} of the V/I and the CM rejection of the VCO. Both A_{CM-DM} and A_{CM} of the V/I must therefore be sufficiently low. Additionally, the different run lengths of the random data increase the maximum CM ripple, in the most extreme case to 450mV, when the CM fully recharges to V_{DD}.

4.1.8. Post-Layout Simulation Results
The layout of the phase detector is shown in Fig. 4.11, it consumes 80 μW while operating at 10 GHz. The gain is plotted in Fig. 4.10a, revealing a K_{PD} of 0.3 V/rad. The locking point varies 0.3 rad over corners which corresponds to 27.8$\pm$2.5ps. By inserting a '0011' data stream at the input, we can simulate the resulting output voltage noise spectrum of the PD and calculate the phase noise spectrum by dividing by K_{PD}^2. The result is plotted in Fig. 4.10b, the K_{PD} used in this setup is 0.3 V/rad. As expected, the spectrum shows a pole related to $R_D C_S$ at >300MHz and has a very low integrated jitter due to the absence of data dependent effects.
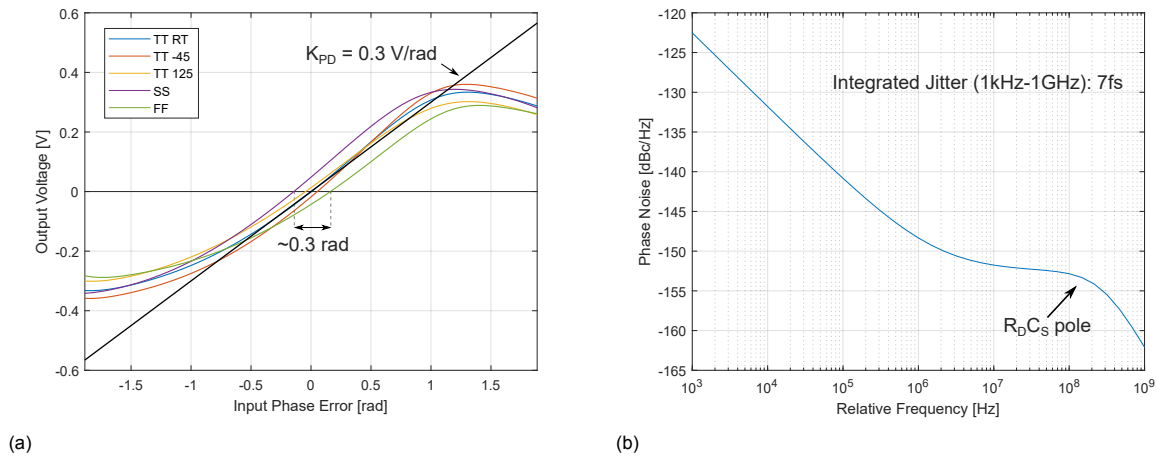

(a)

(b)

Figure 4.10: Post-layout simulation of a) K_{PD} over temperature and corners and b) the phase detector phase noise spectrum referred to 2.5GHz.

Figure 4.11: Layout of the phase detector.

4.2. V/I

The phase detector's differential output voltage is converted to a current with the use of a V/I stage. The resulting current then feeds into the loop filter. The V/I has to handle the phase detector's voltage swing and provide enough transconductance ($K_{V/I}$) to meet the bandwidth specification. The main design parameters are the gain, power consumption and bandwidth. Furthermore, the A_{CM-DM} and A_{CM} should also be sufficiently low.

4.2.1. V/I Design

Fig. 4.12 shows the folded cascode structure used to implement the OTA. It has a tunable degeneration resistance to allow calibration of the gain in testing. The folded cascode structure has been proven at cryogenic temperatures [5] and the transconductance can be described by

$$G_m = \frac{g_m}{1 + g_m R_S},$$
(4.15)

where g_m is the transconductance of the input pair and R_S the value of the degeneration resistance. The power consumption of the V/I is described by $2(1 + g_m R_S)I_{SS}$, where I_{SS} is the tail biasing current. Due to the high phase detector gain, the phase noise of the OTA will be highly suppressed, allowing for a low power consumption. Transistor-level schematic of the V/I is shown in Fig. 4.12.
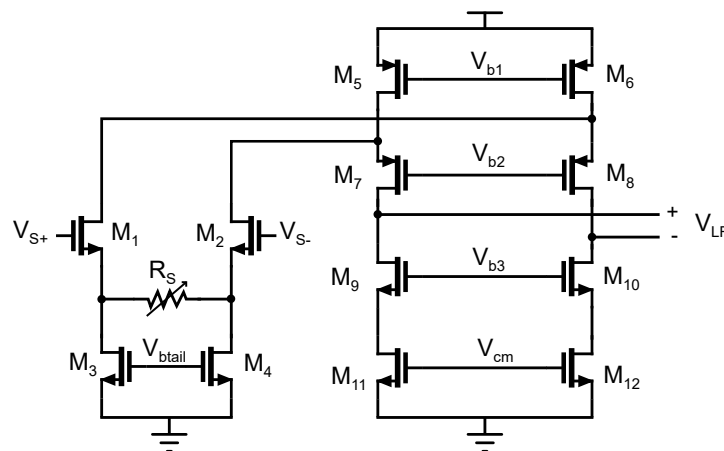

Figure 4.12: Transistor-level schematic of the V/I stage.

The biasing and common-mode feedback circuits are shown in Fig. 4.13, current biasing is used over constant gm-biasing as it has been proven at cryogenic temperatures [5].

Figure 4.13: Transistor-level schematics of the biasing and common-mode feedback of the V/I stage.

4.2.2. Post-Layout Simulation Results

The simulated post-layout $K_{V/I}$ is shown in Fig. 4.14a, in TT it is tunable from 100-140 μS with steps of roughly 4 μS. The gain varies 20% over PVT and can be compensated. The circuit draws 120 μA from a 1.1 V supply. The phase noise contribution of this block depends on the phase detector gain as well as its own. The simulated input-reffered voltage noise spectrum is plotted in Fig. 4.14b, we can see that it is indeed dependent on $K_{V/I}$. The A_{CM-DM} and A_{CM} of the V/I are -80 dB and -25 dB at 2.5 GHz, which suppresses the common-mode ripple of V_S below the data-dependent differential ripple.

Figure 4.14: Post-layout simulation of a) $K_{V/I}$ over temperature and corners and b) the input-referred voltage noise spectrum.

To verify the loop dynamics of the implemented OTA, we simulate the transfer function and check the phase margin. The non-idealities considered here are the DC output impedance, parasitic capacitance of the OTA and secondary loop capacitors. Fig. 4.15 shows the resulting Bode plot of the loop, the phase margin degrades to 69.4 degrees and the bandwidth also reduces slightly. However, both values are within acceptable ranges.
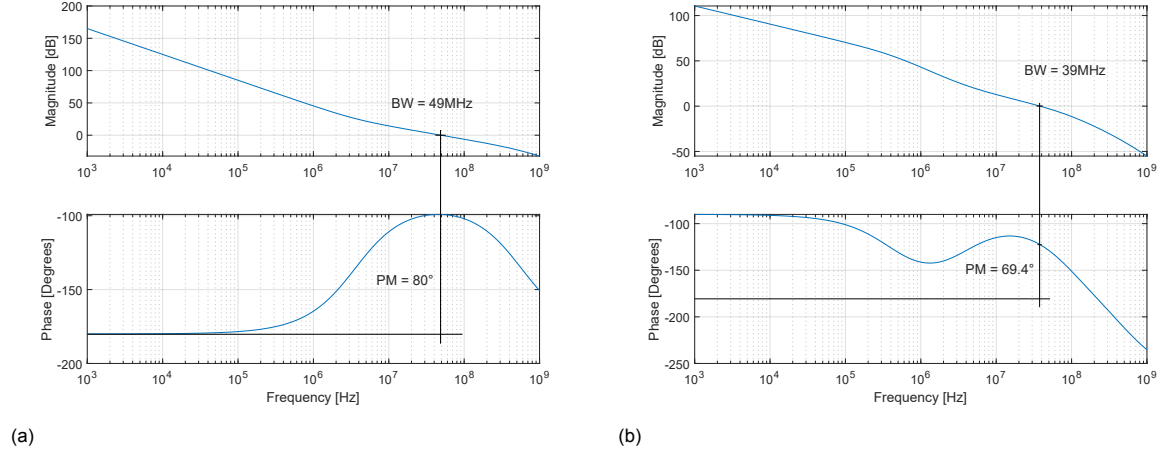
Figure 4.15: Bode plots of the loop with a) ideal $K_{V/I}$ and loop filter and b) post-layout extracted stage.
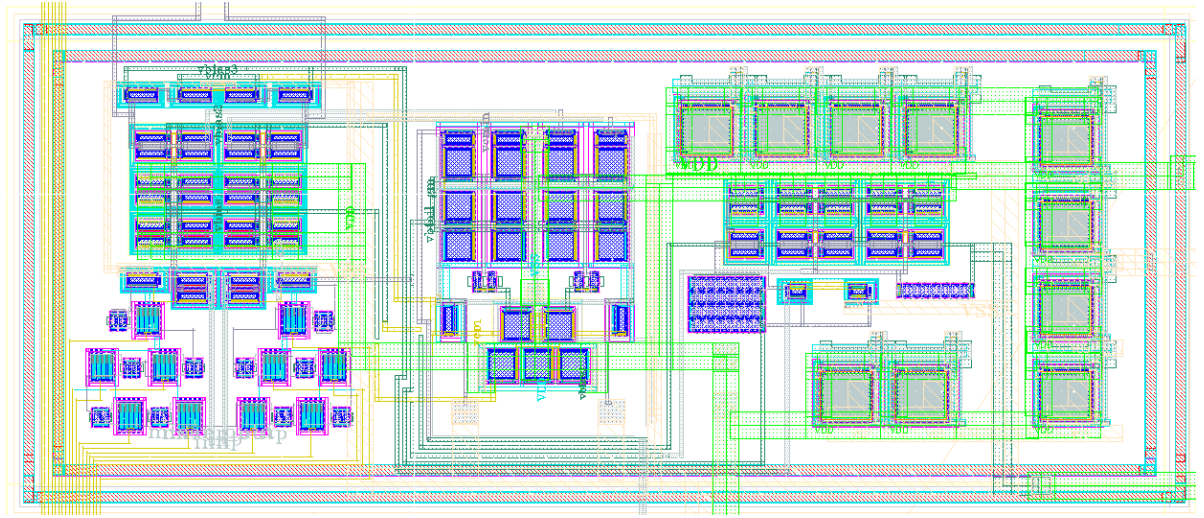

Figure 4.16: Layout of the V/I stage.

4.3. Input Buffer

In a complete wireline receiver, the input of the CDR originates from an equalizer which compensates for channel loss. This work does not focus on the channel response and equalization and instead assumes an input swing of $300\,\mathrm{mV_{pp}}$ provided at the input of the chip. The buffer has to terminate the incoming RF data line, amplify the signal and drive the differential phase detector.

4.3.1. Input Buffer Design

First, the incoming data line is low-side terminated at $50\,\Omega$ using a unsalicided resistor with 3-bit tuning to compensate for process variation. A single-ended input is used over differential due to the phase matching that would be required for the cables during testing. The signal now has to be amplified to full-scale, this is done by using an inverter with tunable drive strength for PMOS and NMOS to bias at the highest gain over PVT. A typical alternative would be the use of a self-biased inverter with capacitive coupling. However, an RC bias-tee will result in droop, which in turn translates into jitter due to the randomness of the data. The longest droop duration is 31 bits, 3.1 ns. Assuming a rise-fall time of 30 ps and a maximum allowed Δt of 150 fs, the maximum allowed droop is $0.15/30 = 5\,\mathrm{mV}$. The required corner frequency of the bias-tee to achieve this has to be

$$0.995 = \exp{-\frac{mT_b}{\tau}} \tag{4.16}$$

$$\tau \approx 618\,ns \tag{4.17}$$

$$f_c = \frac{1}{2\pi \cdot 618\,ns} \approx 260\,kHz \tag{4.18}$$

Achieving such a low corner frequency requires a large capacitor and resistance. Consequently, no bias-tee is applied here. The first amplifier stage does not fully drive the data to rail-to-rail and so two extra inverter stages were added. These are able to drive the line from the input to the center of the chip.

To generate the complementary NRZ signal, a transmission gate in combination with an inverter is used, as shown in Fig. 4.17. This relatively simple circuit has a large bandwidth required by the high frequency data. The drive strength of the transmission gate output is limited so a second stage of inverters is used to improve the slew rate. The phases are aligned using a cross-coupled pair followed by a final driver to the phase detector.

The driver should contribute negligible jitter to the system and the mismatch must not impact performance. To estimate the tolerable delay mismatch, the system was simulated with ideal VerilogA blocks aside from the phase detector. The complementary random data is given as input with a fixed mismatch ranging from 0ps to 10ps. The peak-to-peak jitter on the recovered clock is then compared relative to the ideal case as shown in Table 4.2. For a duty cycle error of 2ps, the increase in data-dependent peak-to-peak jitter is only 3%, which is subsequently set as the target.

Table 4.2: Impact of duty cycle mismatch on the peak-to-peak jitter.

$D_{in}\ \overline{D_{in}}$ Mismatch [ps]	Relative Jp-p w/ ideal input	Relative Jp-p w/ input jitter
0	1	1
2	1.03	1
4	1.13	1
6	1.31	1.04
8	1.48	1.09
10	1.68	1.12

With 15 ps rise and fall times, we can find the required area of the devices:

$$Slope = \frac{\Delta V}{\Delta t} = \frac{1.1}{15\,ps} = 73\,mV/ps \tag{4.19}$$

$$\Delta t = \frac{\Delta V_{th}}{Slope} = \frac{4\,mV}{\sqrt{WL}}\frac{1}{73\,mV/ps}. \tag{4.20}$$

Even for relatively small devices of $(3\,\mu m/40\,nm)_{W/L}$, the mismatch is limited.
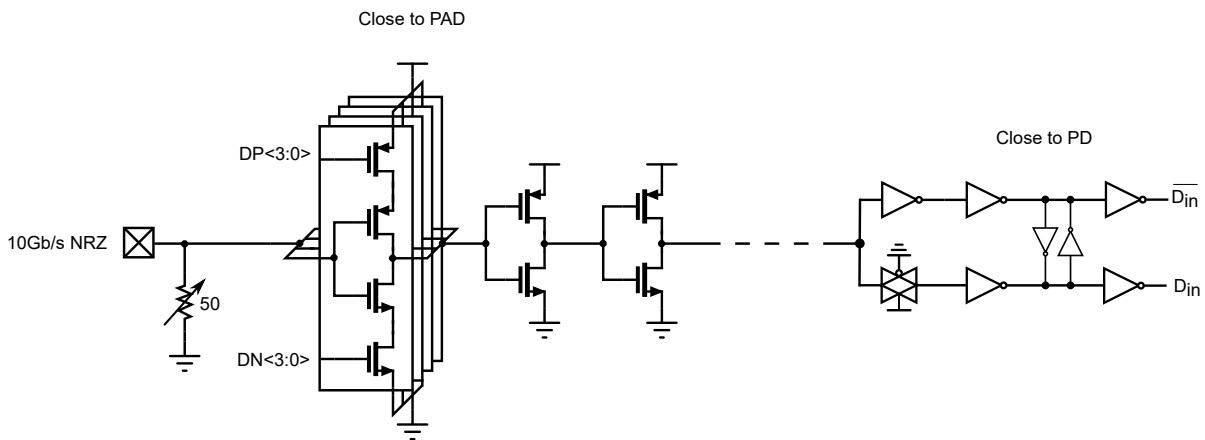

Figure 4.17: Schematic of the complete input buffer.

4.3.2. Post-Layout Simulation Results

The layout of the blocks that are close to the input PAD is shown in Fig. 4.18. The single-to-differential buffer is close to the PD and visible in the bottom-left of Fig. 4.11. The single-to-differential buffer draws 200 μW from a 1.1 V supply with 10 Gb/s NRZ data in TT. The duty cycle mismatch is shown in Table. 4.3 and only slightly exceeds target of 2 ps. The phase noise originating from the input buffer itself is plotted in Fig. 4.19, the loop low-pass filters the transfer, with an integration bound of 100 MHz, the integrated jitter is only 12 fs$_{\mathrm{rms}}$.

Table 4.3: Extracted Monte-Carlo duty cycle mismatch.

| Mean Duty Cycle Error [ps]	σ [ps]	Max Error (3σ) [ps]
0.619	0.522	2.19

Figure 4.18: Layout of the input termination and driver.

Figure 4.19: Phase noise spectrum of the input buffer.

4.4. Retimer

The incoming data must be retimed using the recovered clock. This can be achieved with the use of a D flip-flop. Since it operates at 10 GHz, the design has to be compact to limit its power consumption. Additionally, the setup and hold times should be low to limit degradation of h and the jitter tolerance.

A popular low-power flip-flop type is based on true single-phase clocking (TSPC), it however, fails to operate at frequencies of over 8 GHz. Hence, a complementary version introduced in [9] is used. The schematic is shown in Fig. 4.20a.
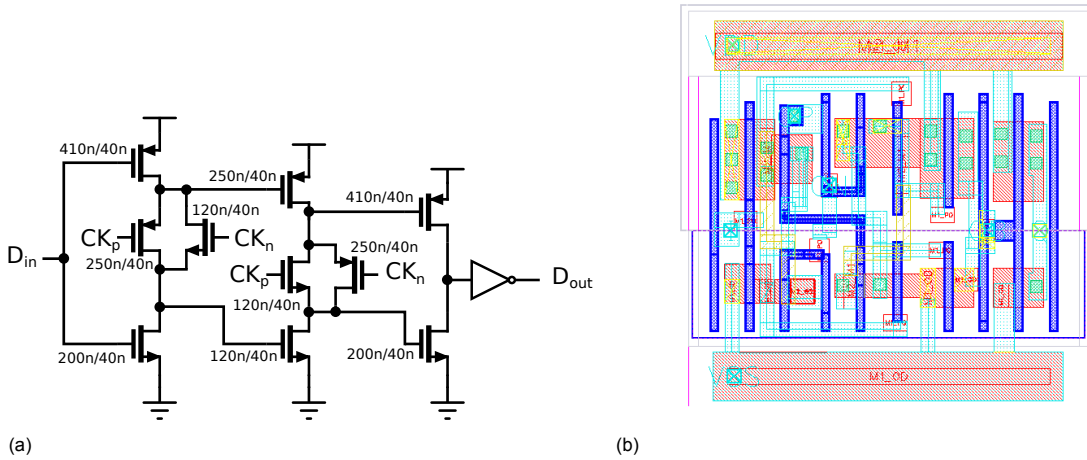

(a)

(b)

Figure 4.20: Retiming flip-flop a) transistor-level schematic and b) layout.

4.4.1. Post-Layout Simulation Results

The parasitics are extracted from the layout in Fig. 4.20b and post-layout simulation is performed to estimate the power consumption as well as setup and hold times and operating frequency. The power consumption is minimal at $60\,\mu$W while operating at 10GHz. To find the optimal moment for retiming, we simulate the values of T_{su} and T_{ho} as defined in Eq. 3.16. The output is no longer correct for a T_{su} below 32 ps and a T_{ho} below 7 ps. Consequently, the optimum distance between the clock and data edge is 62 ps, which would result in a maximum h.

The maximum operating frequency of the flip-flop can be evaluated by connecting $\overline{Q}$ and D to create a divide-by-2 circuit. Fig. 4.21 shows the resulting input/output frequency characteristics. The flipflop operates up to 13.5 GHz with extracted parasitics, fulfilling its required functionality.


Figure 4.21: Input frequency vs output frequency when using the flipflop as divide-by-2.

4.5. Delayline

The delayline should fulfill two main functions, most importantly it has to provide the required delay to center the clock for retiming. By considering both the PD and retimer performance, the delay should be 30 ps. Furthermore, it has to sharpen the clock flanks to improve the setup and hold times of the retimer. Both of these tasks should be accomplished under PVT variations. Additional constraints are

power consumption and jitter.

PVT makes an open-loop solution unfeasible; a normal buffer's delay can vary more than 8 ps degrading h by over 0.15 and requiring 1.5x more bandwidth to reach the same jitter tolerance. To overcome this variance, the delay is made controllable. The targeted accuracy is 2 ps, which would require less than 10% extra bandwidth. An additional constraint here is the limited swing of the VCO itself, it produces a sine wave at mid-rail with ~900 mV$_{pp}$ swing. This results in a lower slew rate making the design more challenging.

There are two main strategies to implement such a delayline in CMOS: current starving and capacitive loading. Current starving limits the current available to an inverter; a difference in current will translate to a difference in the time required to charge the load capacitor. Capacitive loading on the other hand controls the load capacitance seen at the output; a difference in capacitance while having the same drive strength translates to a difference in rise time. Both strategies can be implemented analogously or digitally [41].

4.5.1. DCDL Design

In this design, a digital capacitive tuning method is chosen. The current starved method requires additional PMOS and NMOS devices in series with those of the inverter, effectively limiting the slew rate near the supply rails due to the devices dropping into the triode region. This compromises the second objective of the line; improving the setup time. Capacitive tuning does not have this drawback. Analog capacitive tuning can be implemented using varactors, however, the relative tuning range of $\Delta C/C$ is limited to less than 50%, meaning that the achievable delay range is significantly reduced. The digital alternative is switched capacitors which can achieve a far higher $\Delta C/C$. Consequently, we arrive at the proposed delayline with its transistor-level schematic shown in Fig. 4.22.

It consists of two complementary inverters with the switched capacitors connected to the center nodes. The first inverter in combination with the capacitors control the delay while the second inverter adds additional static delay to reach the total delay requirement. The switched capacitors are implemented similar to those in a DCO with unary coding. Unary was chosen over binary due to the low number of required steps and their inherent monotonicity. M_{b1} and M_{b2} prevent the internal nodes from floating and are minimum size. M_{SW} is the actual switch. The propagation time of an inverter assuming an ideal step at the input is

$$t_p = 0.69RC, \tag{4.21}$$

where R is the average 'on' resistance and C the output capacitance. The limited slew rate of the VCO output degrades this time by

$$t_{p,actual} = \sqrt{t_p^2 + (t_r/2)^2}, \tag{4.22}$$

where t_r is the rise time of the input signal. Considering the tradeoffs between delay, power consumption and jitter, we arrive at the values listed in Fig. 4.22.

4.5.2. Post-Layout Simulation Results

The layout parasitics were extracted and post-layout simulation was performed to estimate the delay, power consumption and jitter. Fig. 4.23a shows the delay range in process corners TT, SS, and FF as well as temperature. The tunable range is roughly 8 ps from 24.5-33 ps. The performance is heavily limited by the VCO swing and slew rate. Power consumption at the lowest setting is 900 μW and 1.3 mW at the highest. Due to the high power consumption, the jitter of the line is low at 15 fs$_{rms}$, as shown in Fig. 4.23b.

4.6. Delay-Loop

In order to set the correct delayline code, the delay must be measured and an error signal is needed in a feedback loop. Since the clock frequency is 10GHz, the RF part should be as compact as possible
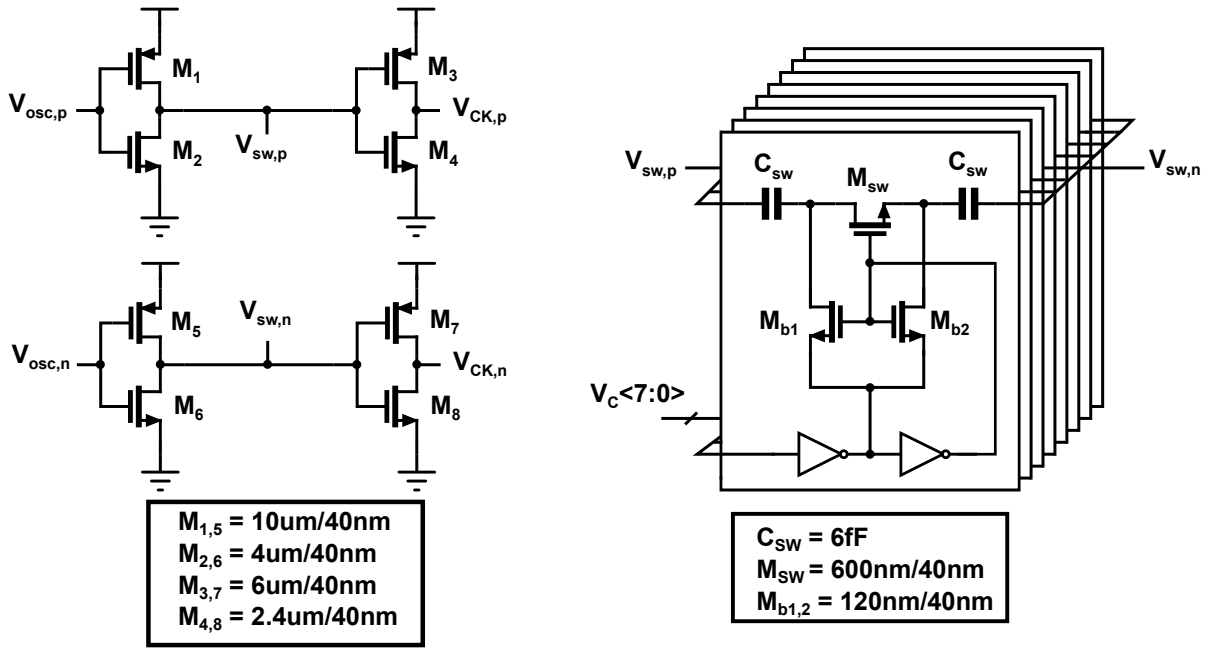
Figure 4.22: Circuit diagram of the DCDL.

to minimize its power consumption. The loop itself does not have to operate at high frequencies as it compensates for PVT. This Section covers the phase detector used to detect the error, as well as the subsequent base-band processing before entering the digital block. The complete circuit diagram is shown in Fig.4.24.

4.6.1. Phase Detector

Phase detection is done similarly to the main loop phase detector, only here, the complementary switching is replaced by a single switch. In this case, the analysis using T_P holds and it is equal to $T_{VCO}/2$. During the sampling pulse, the delayline's output phase error relative to the center of the VCO pulse is determined. The phase detector uses PMOS devices instead of resistors to have a locking point more towards 25 ps rather than 35 ps. It is a single-balanced setup, while a double-balanced setup might seem more appropriate, it has a locking point that is much further removed from the target at >35ps. The locking point of the phase detector serves as the reference delay and therefore should be robust over PVT. This phase detector suffers less from PVT than the main one due to the more ideal tail, allowing smaller devices to be used which reduces the output capacitance of the delayline. The gain of the detector is estimated by $\approx 2G_M A_{VCO} R_{on}/\pi$, where R_{on} is the on-resistance of the PMOS devices replacing the resistors. The gain is 0.25 V/rad to suppress the noise of the subsequent components and have a useful error voltage.

4.6.2. Baseband Blocks

The phase detector essentially functions as a mixer and its output consists of two components, the baseband error voltage and the $2f_{VCO}$ signal. Only the baseband voltage is of interest, therefore, a low-pass filter is first used to attenuate the 20 GHz component to below 1 mV, considerably below the minimum baseband voltage of 12.5 mV. Now, the high frequency component is removed but the the common-mode voltage of the signal is still near V_{DD}, making it difficult to work with. A source follower is employed to bring the signal to mid-rail, subsequently a 5-transistor OTA amplifies the error signal and converts it to single ended. Two inverters then bring it to rail-to-rail and interface with the digital block. The OTA is biased using a reference current generated from an on-chip resistor. Due to the increased threshold voltage at cryogenic temperatures, the locking point of the loop can change, to help compensate for this, tunable source degeneration resistors are added to the OTA, allowing tunability of the locking point.

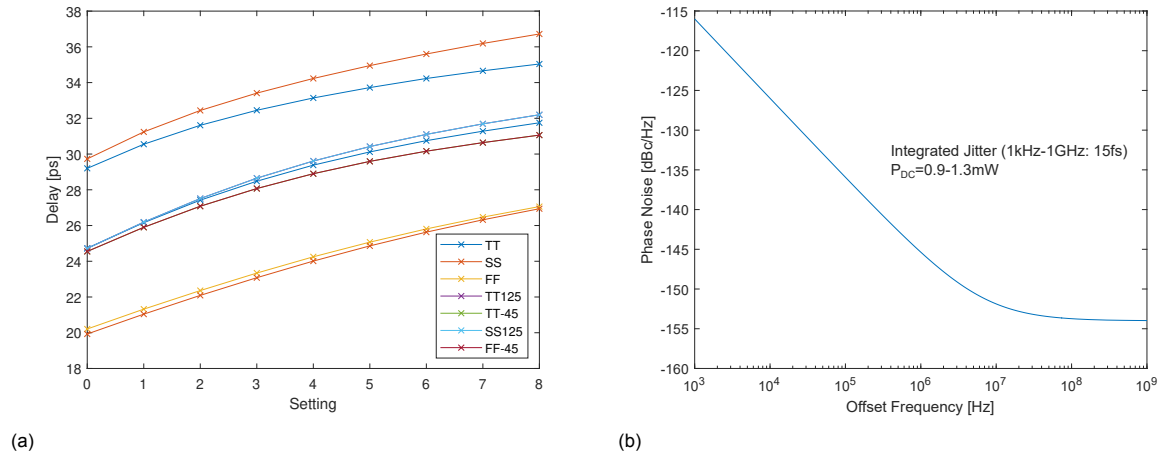(a)   (b)

Figure 4.23: Post-layout simulation of a) the delay range over temperature and corners and b) the phase noise spectrum in the SS-corner.


Figure 4.24: Circuit diagram of the delayloop analog frontend.

4.6.3. Post-Layout Simulation Results

The layout of the delayloop is shown in Fig. 4.26, the delayline is on the right, with the phase detector and subsequent blocks in order from right to left. The phase detector consumes <100 μW at 10 GHz and the 3σ mismatch is ±1.25 ps around a locking point of 29.5 ps. The spread over corners is also ±1.25 ps as shown in Fig. 4.25a. The phase detector output voltage for different delay settings is plotted in Fig. 4.25b. The minimum step voltage is 5 mV between the maximum and second to maximum setting. Large devices were used in the baseband blocks to limit the mismatch and eliminate the need for additional calibration as much as possible. The simulated $\pm3\sigma$ variation after the phase detector is <2 mV, enough to distinguish the individual settings. However as a consequence of the large devices, the bandwidth is low at just over 1 MHz.

Figure 4.25: Post-layout simulation of a) phase detector gain and locking point spread and b) the phase detector output voltage for all delay settings.
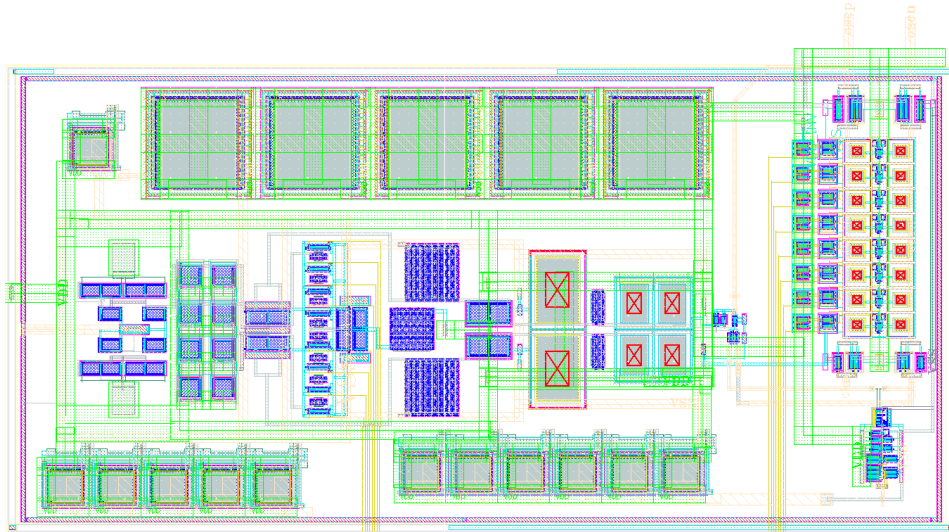

Figure 4.26: Layout of the delayline and delay-loop analog frontend.

4.7. Data Driver

After retiming is performed, the data is driven out of the chip for testing. While on-chip BER tests exist, their measurement capability is limited and it would add additional complexity to the system. Since the driver is purely for testing, it should not degrade the performance of the rest of the system. The timing margin should therefore not be affected by the driver. The driver should also have 50 Ω on-chip matching to properly drive the cable for testing. Disturbance on both supply and ground rails due to the driver's current consumption must also be considered.

4.7.1. Data Driver Design

High-speed wireline drivers come in two main flavors, voltage-mode and current-mode, also known as source-series terminated (SST) and CML drivers. The SST driver core consists of an inverter with a series resistance at its output which, in combination with some calibration slices, matches the line to 50 Ω. The main advantage of this driver over CML is its power efficiency and higher voltage swing [42]. However, neither of these two factors are important here. On the contrary, its downside of high supply rail disturbance goes directly against the desired functionality.

The advantage of the CML driver is low disturbance on the supply and ease of matching the lines

[43]. The driver is therefore implemented as a CML-stage terminated differentially at 50 Ω. To optimize the speed, a pre-driver stage is added [43]. Generation of the differential retimed data is achieved by using the same circuit that was used for the input driver, only here it was scaled up to further reduce mismatch. The tail current is biased using an external current source. The targeted driver output swing is 400 $\mathrm{mV_{pp}}$, high enough to ensure the external BERT will function.

Fig. 4.27 shows the transistor-level schematic of both the driver and pre-driver stage along with its sizing values.
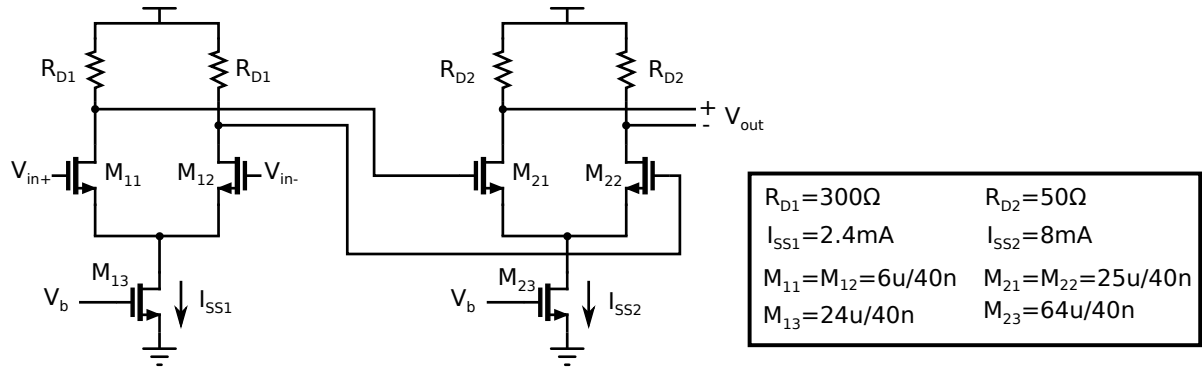

R_{D1}=300Ω	R_{D2}=50Ω
I_{SS1}=2.4mA	I_{SS2}=8mA
M_{11}=M_{12}=6u/40n	M_{21}=M_{22}=25u/40n
M_{13}=24u/40n	M_{23}=64u/40n

Figure 4.27: Circuit diagram of 10Gb/s CML output driver for the retimed data.

| |
|---|
| I_{ref}=1.5mA |
| M_{b1}=12u/40n |
| M_{b2}=36u/40n |
| M_{b3}=12u/40n |

Figure 4.28: Circuit diagram of the CML biasing.

4.7.2. Post-Layout Simulation Results

The layout of the driver is shown in Fig. 4.29. The simulated eye diagram are shown in Fig. 4.30 for TT corner. Monte Carlo simulation shows a three sigma variation of $\pm$1.5ps, negligibly affecting the jitter tolerance. The buffer draws 12 mA from a 1.1 V supply. The power supply ripple introduced by the driver is less than 1 mV and consequently does not affect operation of the delayloop.
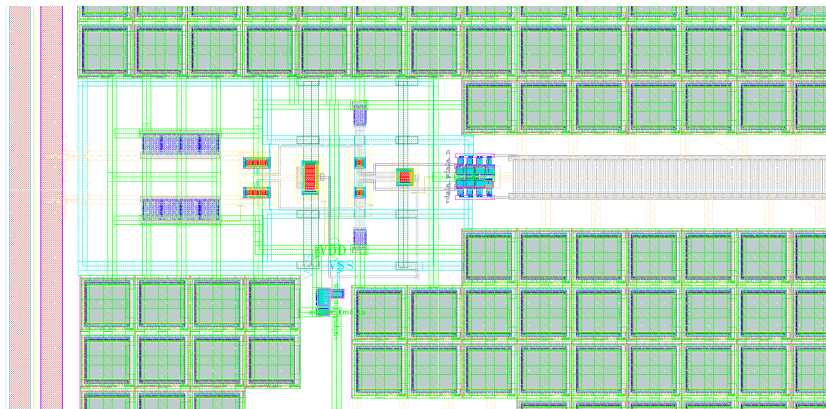

Figure 4.29: Layout of the 10Gb/s test driver.

Figure 4.30: Post-layout simulated eye diagram of the 10Gb/s driver output.

4.8. Other Blocks and Top Layout

This section covers minor blocks such as the loop filter and the modifications done on the VCO. It also shows the top-level floorplan and layout of the design.

4.8.1. Loop Filter

The loop filter consists of the fixed differential capacitor C, two resistors R and a second capacitor C_1 as shown in Fig. 2.20. The resistors are unsalicided poly and tunable between 11 kΩ and 22 kΩ with steps of 1.5 kΩ to compensate the process variation in measurement. The differential capacitor C is 2 pF and the secondary capacitors, C_1, were kept at 30 fF to prevent additional degradation of the loop stability.

The layout of the loop filter is shown in Fig. 4.31. Its input originates from the V/I stage directly below and flows upwards towards the VCO. The resistance and coupling capacitance of the interconnect between the V/I and VCO does not introduce noticeable loop delay, i.e., for 170 um length 0.3 um width M4, the loop delay is « 1ns.


Figure 4.31: Layout of the loop filter with tunable resistance.

4.8.2. VCO

The VCO used in this work is designed by Jiang Gong. It has an extremely low phase noise of -109 dBc/Hz at 1 MHz offset exceeding the specification calculated in Chapter 3 of -90 dBc/Hz by nearly 20 dB. Two changes have been made to the existing design: the varactor was scaled up and M6 dummies were added. The varactors are scaled up to increase the oscillator gain to help achieve the 40 MHz bandwidth requirement. The resulting K_{VCO} at both the maximum and minimum oscillation frequency is shown in Fig. 4.32. It varies from 100 MHz/V to 150 MHz/V over the tuning range.

Fabrication requirements demand that a reasonable M6 density is present in the inductor region. This was added in the form of square dummies and EM simulations were performed to see their impact on the Q-factor of the inductor. The dummies are too fine to directly simulate the actual layout, therefore, larger squares of M6 with the same total area density were used instead. The TSMCN40 layerstack was imported to ADS and a Momentum RF simulation extracted the S-parameters of the inductor. An AC simulation in Cadence shows the resonant peak of the LC-tank with ideal C. The the Q-factor is

(a)   (b)

Figure 4.32: Oscillation frequency versus varactor tuning voltage plotted around a) 9.65GHz b) 11.3GHz.

then found by

$$Q = \frac{f_0}{2f_{-3dB}} \tag{4.23}$$

where f_0 is the oscillation frequency and f_{-3dB} the cutoff frequency. The resulting Q-factors are shown in Table 4.4, the Q-factor degradation for the 20% fill is roughly 20%. As a result, the total tank Q-factor decrease will be less than 20%. The worst case degradation of the oscillator FoM described by Eq. 4.24 is $-10\log_{10}(1/0.8^2) = -1.94dB$. Since the VCO phase noise is already far below the specification this degradation was accepted.

$$FoM = -10\log_{10}(\frac{500KT}{Q_t^2 \cdot \alpha_I \cdot \alpha_V} \cdot (1 + \gamma)) \tag{4.24}$$

Table 4.4: Inductor Q-factor with different dummy densities.

	f_0 [GHz]	f_{-3dB} [GHz]	Q
M6 0%	14.0	0.277	25
M6 15%	13.92	0.330	21
M6 20%	13.86	0.347	20

4.8.3. Floor Plan and Grounding

The top-level layout is shown in Fig. 4.33. The core circuit including I/O drivers is $0.20\,\mathrm{mm}^2$, with the total chip area dominated by decoupling capacitors. The higher level of spectral content in random data in combination with the low desired jitter, requires careful management of the supply domains of the chip. The single-ended input signal introduces significant ground-bounce and, due to its single-ended nature, is especially vulnerable to supply disturbance. The supply of the input block together with the phase detector is therefore isolated from the rest of the system using a 'hard' cut, meaning that the ground and supply are only connected outside the chip. Signals crossing from this block to other blocks are preferably differential so that a high frequency return path is present, reducing ground bounce. This is the case for both the oscillator input and phase detector output. Simulation shows that the single-ended data signal to the retimer introduces negligible ground bounce and was therefore kept single-ended. The core clock recovery loop is fully differential and as a result does not suffer as much from the supply ripples. The test buffer for the recovered clock is also isolated using a 'hard' cut to reduce the coupling of the switching transients originating from the data to the single-ended clock output.

The ground of the delayline, retimer and data driver are shared to eliminate ground bounce originating from the return path of the retimed data. In order to further shield the blocks, they have been placed in deep N-wells to lower the switching transients to reduce coupling through the substrate. The VCO itself is placed directly on the substrate. To still provide shielding, a clean AC ground ring is added around the VCO, consisting of OD-M7 connected to an external AC ground via an additional pad. This aims to absorb high-frequency currents present in the substrate, subsequently resulting in lower coupling to adjacent blocks.
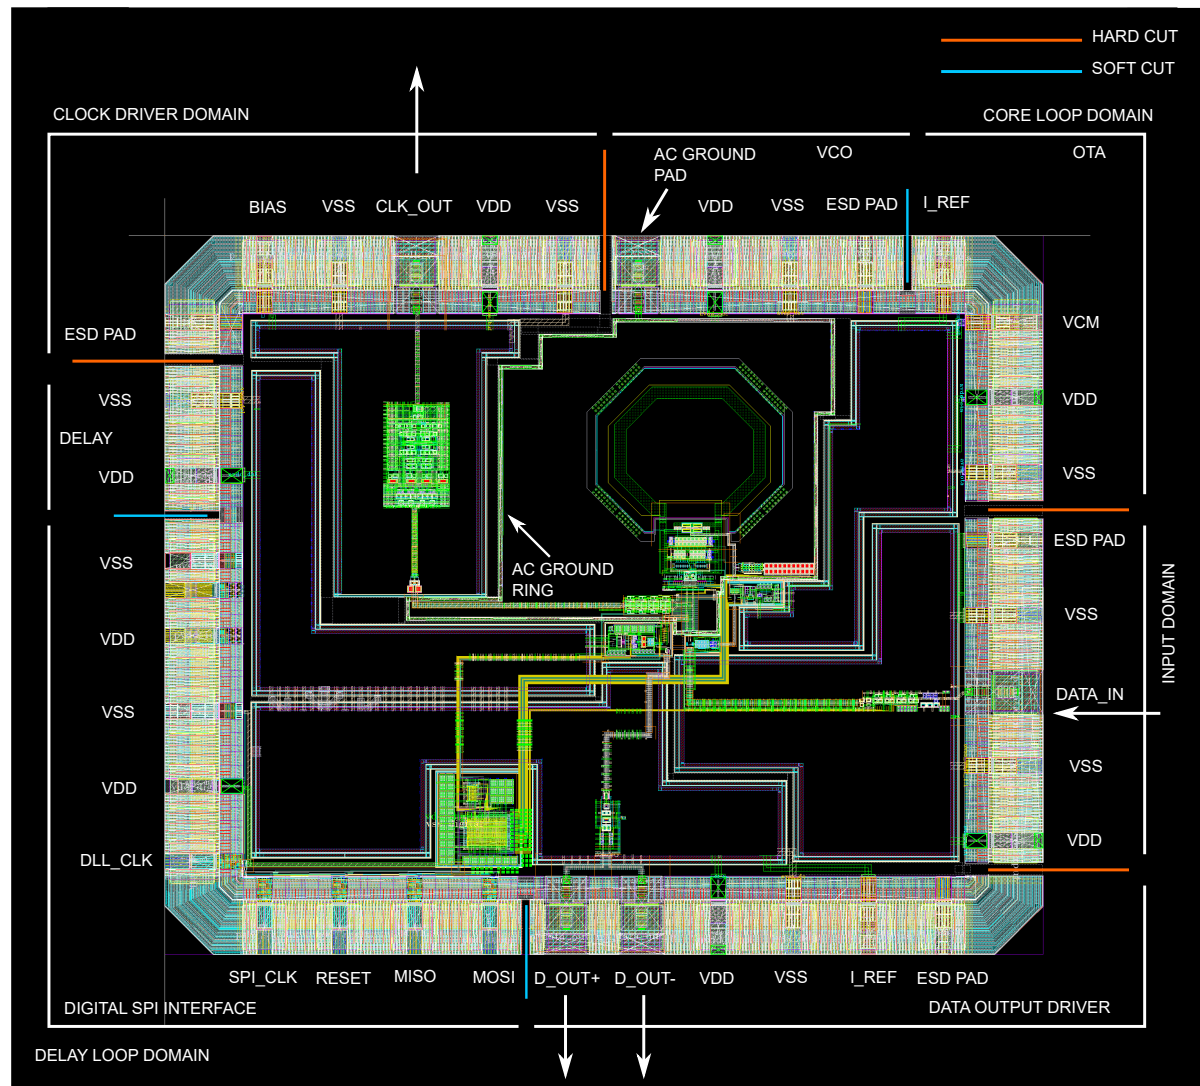

Figure 4.33: Top-level layout of the proposed CDR.

Figure 4.34: Top-level with supply domains and pads of the proposed CDR.

5

RTL Design and Simulation

This chapter will discuss the digital part of the design, entailing the digital part of the delayloop. It is clocked by a low frequency external source (<25 MHz).

5.1. RTL Design

The incoming error signal from the analog front-end is a static bit indicating that the phase is either lagging or leading. Based on this bit, the delay setting must be decreased or increased accordingly. A block diagram of the digital architecture is shown in Fig. 5.1. First, the error is accumulated and depending on its sign, the block will increase its output by '1' or decrease by '1'. A decoder then translates the binary output to an 8-bit thermometer code to format it for the delayline. A MUX is placed in front of the output to allow selection between the loop output and manual settings for testing.
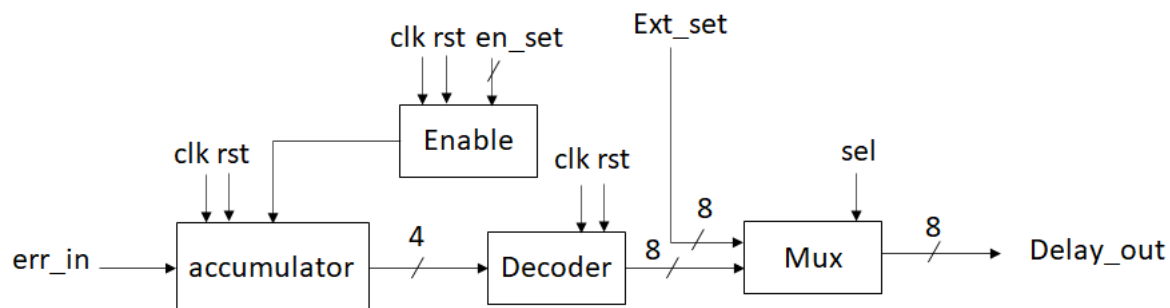

Figure 5.1: Block diagram of the delay loop digital part.

In order to ensure that no ambiguity will occur due to the analog and digital blocks being asynchronous, the accumulator is enabled at a lower rate. As shown in Fig. 5.2, with this block, the accumulator is only clocked at a fraction of the digital clock. The analog loop and digital can therefore run at such a low frequency that after each delay setting change, the analog front-end can settle before the next error bit is processed. Since there is no '0' error or locked state input signal, the system will continue to switch 1 bit around the PD's locking point. However, since a single bit toggle will only degrade the h by 1 ps or 0.01 UI it was accepted. The chance still remains that a setting will be sufficiently close to the locking point such that the analog block does not produce a correct error bit. In such a situation, the loop toggles around 2 bits. This can be resolved by adjusting the resistor setting in the analog block. The 'en_set', 'ext_set', 'rst' and 'sel' are set and controlled externally via serial-peripheral interface (SPI). The various tuning bits in the analog blocks are also set using registers controlled via SPI.
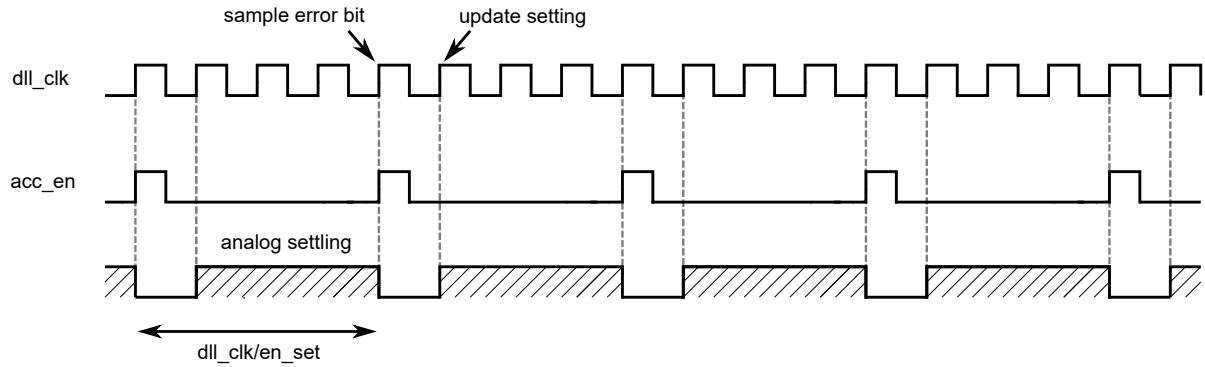
Figure 5.2: Timing of the digital loop.

5.2. Delay-Loop Post-Layout Simulation

The synthesized digital block and post-layout analog delay block are simulated to verify the functionality and performance of the loop. Fig. 4.24 shows the locking behavior in TT, we can observe that the system gets locked and toggles between a single delay setting. In this simulation, an offset in the resistor was used to prevent 2-bit toggling. The system also gets locks in the SS corner but is out of range of FF.

(a)   (b)
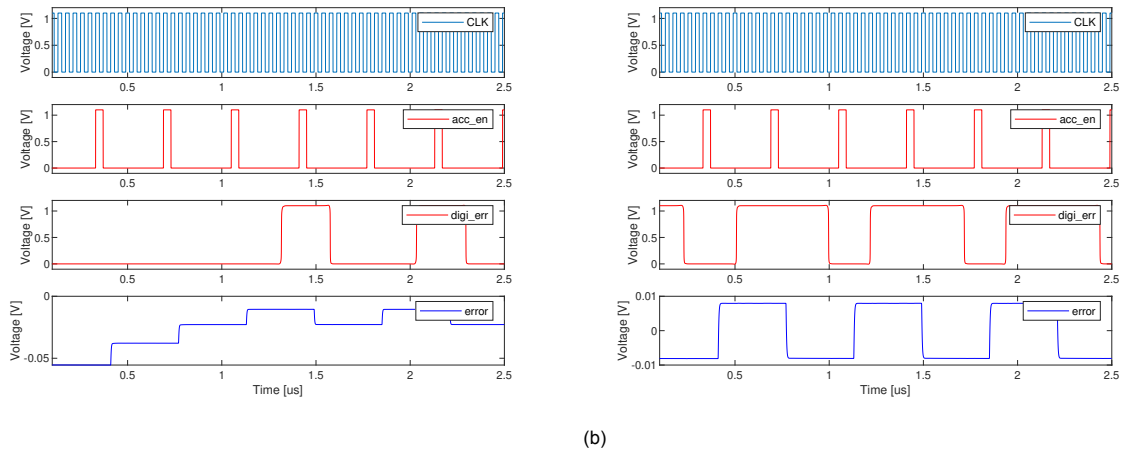
Figure 5.3: Waveforms in the delayloop in the a) TT corner, b) SS corner.

The simulated eye diagrams in the locked state for each corner are shown in Figs. 5.5 and 5.6a. To observe the overall spread, the three eye diagrams are layered together in Fig. 5.6. The spread over the corners is down from 10 ps to just over 4 ps. In the SS and FF corner this increases the h of the jitter tolerance by 0.1, a noticeable improvement. Power consumption of the digital block is negligible at less than 10 μW.

Figure 5.4: Waveforms in the delayloop in the FF corner.

(a)

(b)

Figure 5.5: Input and output eye diagrams of the delayloop in a) TT corner and b) SS corner.

(a)

(b)

Figure 5.6: Input and output eye diagrams of the delayloop in a) FF corner and b) combined corners.

6

System-Level Post-Layout Simulation Results

This chapter covers the top-level simulations done to verify the performance of the system and provides a comparison with the current state-of-the-art.

6.1. CDR Locking

The settling of the system is simulated for different initial phases as well as multiple corners. The input is a pseudo-random binary sequence (PRBS) with a $300\,\mathrm{mV_{pp}}$ swing. Fig. 6.1 shows the oscillator locking to 10 GHz and the test divider output mimicking the lock to 2.5 GHz. The settling is well-behaved and the difference in the initial phase causes the different initial trajectories. From the settling, we can estimate a bandwidth of roughly 45 MHz. Increasing the bandwidth by setting the loop parameters will show a more pronounced effect of the loop latency. The data-dependent jitter can also be observed from the figure, this will be further addressed in the next section.

(a)   (b)

Figure 6.1: Simulated initial VCO locking with a PRBS31 input at a) the VCO nodes b) the test divider output.

6.2. CDR Phase Noise

The phase noise spectrum of the recovered clock is simulated for the complete system. We obtain the spectrum by running a pseudo-steady state (PSS) simulation while providing a '0011' data pattern at the input. Due to limited simulation power, only the core loop consisting of the PD and its driver, the V/I, loop filter and VCO are extracted, the other blocks are schematics. Fig. 6.2 shows the phase noise spectrum with both the top-level simulated result as well as a composite that is calculated using

each block's individual contribution according to Eq. 3.20. A reasonable agreement is observed with a calculated bandwidth of ~44MHz. There is a minor peaking of ~1 dB in the simulation, likely due to additional parasitics in the layout that slightly decrease the phase margin. Due to the high bandwidth, the VCO contribution (yellow) is negligible. Combined with an extremely low in-band generation by the PD and V/I, the integrated rms jitter generation of the CDR is only 40 fs. However, this simulation effectively uses a clock as its input, not random data and therefore fails to capture the impact of this randomness on the entire system.


Figure 6.2: Phase noise spectrum of the recovered CDR clock with a 300mVpp '0011' input.

To obtain an estimate of the recovered clock jitter when operating with a random data input, a post-layout transient noise simulation is performed. The resulting phase noise spectrum is then calculated in MATLAB based on the zero crossings of the recovered clock and is shown in Fig. 6.3. This simulation was performed with minimum bandwidth settings on both the V/I gain as well as the loop filter resistance. Consequently, the resulting bandwidth is only around 30MHz and a peaking is observed of ~2 dB. To ensure that the steady-state and transient simulations match, the transient noise simulation was also run for the '0011' input. The simulated increase in phase noise is roughly 7dB for the plotted range, indicated by the dashed line. This difference is due to a range of effects, such as the data dependency present in the PD, the ISI in the input path, and switching noise in the supplies. The calculated absolute rms jitter of the transient noise simulation is 78 fs. At a higher bandwidth setting, the jitter will increase due to the higher in-band contribution. With a setting of 40 MHz, the simulated jitter is 84 fs_{rms}. Simulation of the full spectrum using transient noise is not possible due to excessive simulation times. Nevertheless, it can be observed that the phase noise of each block is well below the specification, and even the transient noise simulation shows a jitter below the target. To properly verify the performance, physical measurements should be performed.

Figure 6.3: Phase noise spectrum of the recovered CDR clock with a 300mVpp PRBS31 input.

6.3. CDR Jitter Tolerance

The jitter tolerance itself cannot realistically be simulated. The system would have to run long enough to accurately detect BERs of 10e-12 for each frequency offset point. It can therefore only be measured with the physical chip. However, based on the post-layout simulation results, a reasonable estimate can be made. Using Eq. 3.15, along with all the extracted loop parameters, we plot the estimated jitter tolerance in Fig. 6.4. The nominal locking point in TT is 57.3 ps. The combined worst case 3σ deviations of all the block reduces this by 5 ps, to 52.3 ps. With a setup time of 32 ps, the value of h is $0.40\,\mathrm{UI_{pp}}$. The estimated jitter tolerance at 10 MHz offset is 1.1, fulfilling the target specification. There is a bump present due to high additional capacitance at the loop filter node due to the varactor. However, it does not violate the STM-256 mask.


Figure 6.4: Estimated jitter tolerance based on post-layout loop parameters.

6.4. Power Consumption

An overview of the power consumption is shown in Fig. 6.5. The total simulated power consumption is 3.89 mW which is lower than the targeted specification. The VCO consumes the most, at nearly

two-thirds of the power followed by the delay-line and -loop at 1.2 mW. The PD (excluding driver), V2I, retimer contribute 80 μW, 130 μW and 60 μW, respectively. Including the S2D would increase the total by 0.19 mW. The power consumption of the VCO can be reduced without degrading the total system's phase noise performance to lower the total power consumption further. The proposed phase detector itself consumes just 80 μW, a negligible amount of the total power. Note that the power consumption of the VCO is expected to increase during measurement as the reported value here utilizes a transient simulation with the inductor schematic provided by TSMC, which has a significantly higher Q factor than in reality.


Figure 6.5: Power consumption breakdown of the CDR core.

6.5. Performance Summary and Comparison Table

Table 6.1 lists the performance summary of the proposed PLL-based CDR and compares it with other reported state-of-the-art CDRs. The recovered clock jitter while operating at 10 Gb/s is 84 fs_{rms}, the lowest reported, with an estimated jitter tolerance of roughly 1.1 UI_{pp} at 10MHz offset. The bit efficiency is 0.39 pJ/bit. Note that the performance of this work is based on post-layout simulation not actual measurement.

Table 6.1: Comparison table with state-of-the-art CDRs.

	This Work	[33] JSSC'13	[9] JSSC'19	[16] JSSC'18	[7] JSSC'10
Data Rate [Gb/s]	10	25	20	25	25
Jitter Tolerance @ 10MHz [UIpp]	1.1	0.3	1	0.6	0.4
Rec. Clock Jitter [ps]	0.084	1.5	0.459	1.46	0.254
Power [mW]	3.89	5	3	46	99
Efficiency [pJ/bit]	0.39	0.2	0.15	1.8	3.96
Architecture	Type-II PLL	Type-II PLL	Type-I PLL	Type-II ADPLL	Type-II PLL
Technology [nm]	40	65	45	40	65
Supply [V]	1.1	1	1	1.15	1.2

7

Conclusion

7.1. Thesis Conclusion

Analysis, design and validation of a phase-locked loop (PLL)-based clock and data recovery (CDR) system for data recovery are presented in this thesis. The specifications are determined in Chapter 1. Chapter 2 then studied the basics of CDR along with state-of-the-art designs. Based on the specifications and prior art, a PLL-based CDR architecture is proposed. Chapter 3 focused on the system modeling to analyze its dynamics to fulfill the specifications. Chapter 4 then presented the transistor-level design of analog and RF blocks in the system. The RTL part of the design is covered in Chapter 5. Top-level performance is presented in Chapter 6 to verify functionality of the system.

A new phase detector is proposed for a PLL-based CDR operating at 10 Gb/s. The phase detector utilizes complementary switches to detect data transitions and convert the phase error into a voltage. A Type-II loop with a 40 MHz bandwidth is used to subsequently lock the VCO to the data. To achieve a more optimal timing margin between the clock and data, a digitally controlled delayline is employed, improving jitter tolerance in corners. The CDR, excluding testbuffers, consumes a total power of 3.89 mW and achieves a recovered clock jitter of $84\,\mathrm{fs_{rms}}$. The estimated jitter tolerance is $1.1\,\mathrm{UI_{pp}}$ at 10 MHz offset and fulfills the STM-256 mask requirement. The power, jitter and jitter tolerance specifications have all been met.

7.2. Future Work

7.2.1. Delay Loop

The performance of the delayloop in its current state can be greatly improved. The gain of the phase detector can be increased by >2x due to its limited phase error range. Yielding a larger error voltage. Furthermore, a hysteresis region can be used to remove the toggling that is currently present, for example, two phase detectors with a built-in offset can determine a locked state. Combining this with a lock detection block, the loop could be turned off after locking.

7.2.2. VCO Power Consumption

The VCO used in this work is not optimized for the loop bandwidth and subsequent phase noise specification. As such by optimizing the design, the bit efficiency can be significantly improved. One method would be to increase the oscillation frequency; a 1.5x increase in frequency will not result in 1.5x more power, subsequently increasing the efficiency.

7.2.3. Frequency Tracking

The presented design does not include a frequency tracking loop, limiting the acquisition range to roughly ± 50 MHz. To be able to cover the entire VCO range, a separate loop should be added. A classical frequency detector can suffice. Alternatively, the aliased frequency at the output of the phase detector can be used to detect the frequency difference between the VCO and incoming data.

7.2.4. PAM4

Currently, PAM4 CDR is an active research area and the proposed phase detector can be integrated in such a system. For example, the use of an edge selector allows the loop to only process 'full' data transitions, the resulting reduced transition density could be compensated by increasing R_D of the phase detector. Fig. 7.1 shows a block diagram of a PAM4 CDR incorporating the proposed phase detector.
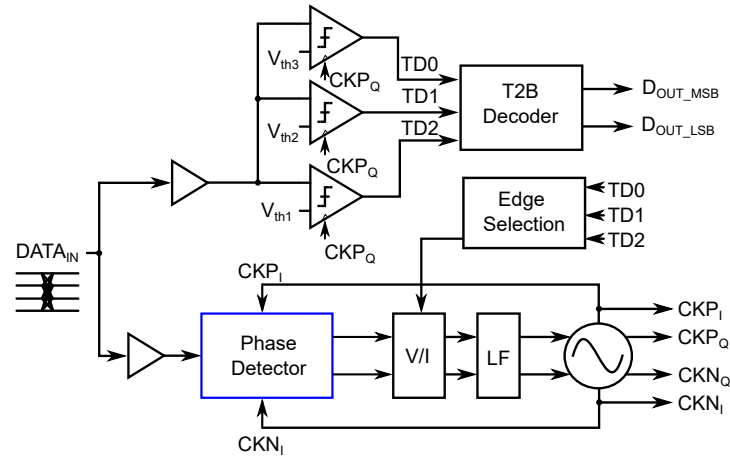

Figure 7.1: PAM4 CDR architecture with proposed phase detector.

Bibliography

[1] Behzad Razavi. *Design of Integrated Circuits for Optical Communications*. 1st ed. USA: McGraw-Hill, Inc., 2002. ISBN: 0072822589.

[2] Edoardo Charbon et al. "15.5 Cryo-CMOS circuits and systems for scalable quantum computing". In: *2017 IEEE International Solid-State Circuits Conference (ISSCC)*. 2017, pp. 264–265. DOI: `10.1109/ISSCC.2017.7870362`.

[3] Bishnu Patra et al. "Cryo-CMOS Circuits and Systems for Quantum Computing Applications". In: *IEEE Journal of Solid-State Circuits* 53.1 (2018), pp. 309–321. DOI: `10.1109/JSSC.2017.2737549`.

[4] Gerd Kiene et al. "13.4 A 1GS/s 6-to-8b 0.5mW/Qubit Cryo-CMOS SAR ADC for Quantum Computing in 40nm CMOS". In: *2021 IEEE International Solid- State Circuits Conference (ISSCC)*. Vol. 64. 2021, pp. 214–216. DOI: `10.1109/ISSCC42613.2021.9365927`.

[5] Jiang Gong et al. "A 2.7mW 45fsrms-Jitter Cryogenic Dynamic-Amplifier-Based PLL for Quantum Computing Applications". In: *2021 IEEE Custom Integrated Circuits Conference (CICC)*. 2021, pp. 1–2. DOI: `10.1109/CICC51472.2021.9431541`.

[6] ITU-T. *G.825 : The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)*. 2008.

[7] Ke-Chung Wu and Jri Lee. "A 2x25-Gb/s Receiver With 2:5 DMUX for 100-Gb/s Ethernet". In: *IEEE Journal of Solid-State Circuits* 45.11 (2010), pp. 2421–2432. DOI: `10.1109/JSSC.2010.2074291`.

[8] Behzad Razavi. *RF Microelectronics (2nd Edition) (Prentice Hall Communications Engineering and Emerging Technologies Series)*. 2nd. USA: Prentice Hall Press, 2011. ISBN: 0137134738.

[9] L. Kong, Y. Chang, and B. Razavi. "An Inductorless 20-Gb/s CDR With High Jitter Tolerance". In: *IEEE Journal of Solid-State Circuits* 54.10 (2019), pp. 2857–2866. DOI: `10.1109/JSSC.2019.2930899`.

[10] Verbeke, Marijn. "Low-power subsampling all-digital clock and data recovery techniques for multi-gigabit passive optical networks". eng. PhD thesis. Ghent University, 2018. ISBN: 978-94-6355-088-8.

[11] Jri Lee and Ke-Chung Wu. "A 20-Gb/s Full-Rate Linear Clock and Data Recovery Circuit With Automatic Frequency Acquisition". In: *IEEE Journal of Solid-State Circuits* 44.12 (2009), pp. 3590–3602. DOI: `10.1109/JSSC.2009.2031042`.

[12] B. Razavi. "Challenges in the design high-speed clock and data recovery circuits". In: *IEEE Communications Magazine* 40.8 (2002), pp. 94–101. DOI: `10.1109/MCOM.2002.1024421`.

[13] Hui Wang and R. Nottenburg. "A 1 Gb/s CMOS clock and data recovery circuit". In: *1999 IEEE International Solid-State Circuits Conference. Digest of Technical Papers. ISSCC. First Edition (Cat. No.99CH36278)*. 1999, pp. 354–355. DOI: `10.1109/ISSCC.1999.759292`.

[14] Pavan Kumar Hanumolu et al. "A 1.6Gbps Digital Clock and Data Recovery Circuit". In: *IEEE Custom Integrated Circuits Conference 2006*. 2006, pp. 603–606. DOI: `10.1109/CICC.2006.320829`.

[15] Yong-Hun Kim et al. "A 10-Gb/s Reference-Less Baud-Rate CDR for Low Power Consumption With the Direct Feedback Method". In: *IEEE Transactions on Circuits and Systems II: Express Briefs* 65.11 (2018), pp. 1539–1543. DOI: `10.1109/TCSII.2017.2758923`.

[16] Marijn Verbeke et al. "A 1.8-pJ/b, 12.5–25-Gb/s Wide Range All-Digital Clock and Data Recovery Circuit". In: *IEEE Journal of Solid-State Circuits* 53.2 (2018), pp. 470–483. DOI: `10.1109/JSSC.2017.2755690`.

[17]   Behzad Razavi. "The Delay-Locked Loop [A Circuit for All Seasons]". In: *IEEE Solid-State Circuits Magazine* 10.3 (2018), pp. 9–15. DOI: `10.1109/MSSC.2018.2844615`.

[18]   Ming-ta Hsieh and Gerald E. Sobelman. "Architectures for multi-gigabit wire-linked clock and data recovery". In: *IEEE Circuits and Systems Magazine* 8.4 (2008), pp. 45–57. DOI: `10.1109/MCAS.2008.930152`.

[19]   T.H. Lee and J.F. Bulzacchelli. "A 155-MHz clock recovery delay- and phase-locked loop". In: *IEEE Journal of Solid-State Circuits* 27.12 (1992), pp. 1736–1746. DOI: `10.1109/4.173100`.

[20]   Woogeun Rhee et al. "A 10-Gb/s CMOS clock and data recovery circuit using a secondary delay-locked loop". In: *Proceedings of the IEEE 2003 Custom Integrated Circuits Conference, 2003.* 2003, pp. 81–84. DOI: `10.1109/CICC.2003.1249364`.

[21]   Christian Kromer et al. "A 25-Gb/s CDR in 90-nm CMOS for High-Density Interconnects". In: *IEEE Journal of Solid-State Circuits* 41.12 (2006), pp. 2921–2929. DOI: `10.1109/JSSC.2006.884389`.

[22]   Hiok-Tiaq Ng et al. "A second-order semidigital clock recovery circuit based on injection locking". In: *IEEE Journal of Solid-State Circuits* 38.12 (2003), pp. 2101–2110. DOI: `10.1109/JSSC.2003.818576`.

[23]   J. Savoj and B. Razavi. "A 10-Gb/s CMOS clock and data recovery circuit with a half-rate linear phase detector". In: *IEEE Journal of Solid-State Circuits* 36.5 (2001), pp. 761–768. DOI: `10.1109/4.918913`.

[24]   Wahid Rahman et al. "A 22.5-to-32-Gb/s 3.2-pJ/b Referenceless Baud-Rate Digital CDR With DFE and CTLE in 28-nm CMOS". In: *IEEE Journal of Solid-State Circuits* 52.12 (2017), pp. 3517–3531. DOI: `10.1109/JSSC.2017.2744661`.

[25]   Changzhi Yu et al. "A 6.5–12.5-Gb/s Half-Rate Single-Loop All-Digital Referenceless CDR in 28-nm CMOS". In: *IEEE Journal of Solid-State Circuits* 55.10 (2020), pp. 2831–2841. DOI: `10.1109/JSSC.2020.3005750`.

[26]   Hsiang-Hui Chang, Rong-Jyi Yang, and Shen-Iuan Liu. "Low jitter and multirate clock and data recovery circuit using a MSADLL for chip-to-chip interconnection". In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 51.12 (2004), pp. 2356–2364. DOI: `10.1109/TCSI.2004.838147`.

[27]   Dong Hoon Baek et al. "2.6 A 5.67mW 9Gb/s DLL-based reference-less CDR with pattern-dependent clock-embedded signaling for intra-panel interface". In: *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*. 2014, pp. 48–49. DOI: `10.1109/ISSCC.2014.6757332`.

[28]   Guanghua Shu et al. "A 4-to-10.5 Gb/s Continuous-Rate Digital Clock and Data Recovery With Automatic Frequency Acquisition". In: *IEEE Journal of Solid-State Circuits* 51.2 (2016), pp. 428–439. DOI: `10.1109/JSSC.2015.2497963`.

[29]   J. Kenney et al. "A 9.95 to 11.1Gb/s XFP transceiver in 0.13/spl mu/m CMOS". In: *2006 IEEE International Solid State Circuits Conference - Digest of Technical Papers*. 2006, pp. 864–873. DOI: `10.1109/ISSCC.2006.1696127`.

[30]   Kyongsu Lee and Jae-Yoon Sim. "A 0.8-to-6.5 Gb/s Continuous-Rate Reference-Less Digital CDR With Half-Rate Common-Mode Clock-Embedded Signaling". In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 63.4 (2016), pp. 482–493. DOI: `10.1109/TCSI.2016.2528480`.

[31]   C. R. Hogge. "A self correcting clock recovery circuit". In: *IEEE Transactions on Electron Devices* 32.12 (1985), pp. 2704–2706. DOI: `10.1109/T-ED.1985.22402`.

[32]   J. D. H. Alexander. "Clock recovery from random binary signals". In: *Electronics Letters* 11.22 (1975), pp. 541–542. DOI: `10.1049/el:19750415`.

[33]   Jun Won Jung and Behzad Razavi. "A 25-Gb/s 5-mW CMOS CDR/Deserializer". In: *IEEE Journal of Solid-State Circuits* 48.3 (2013), pp. 684–697. DOI: `10.1109/JSSC.2013.2237692`.

[34] Zhao Zhang and C. Patrick Yue. "A 12.5-Gb/s 4.8-mW Full-Rate CDR with Low-Power Sample-and-Hold Linear Phase Detector". In: *2018 IEEE International Conference on Integrated Circuits, Technologies and Applications (ICTA)*. 2018, pp. 96–97. DOI: `10.1109/CICTA.2018.8706047`.

[35] Zhao Zhang et al. "A 32-Gb/s 0.46-pJ/bit PAM4 CDR Using a Quarter-Rate Linear Phase Detector and a Self-Biased PLL-Based Multiphase Clock Generator". In: *IEEE Journal of Solid-State Circuits* 55.10 (2020), pp. 2734–2746. DOI: `10.1109/JSSC.2020.3005780`.

[36] Jri Lee, K.S. Kundert, and B. Razavi. "Analysis and modeling of bang-bang clock and data recovery circuits". In: *IEEE Journal of Solid-State Circuits* 39.9 (2004), pp. 1571–1580. DOI: `10.1109/JSSC.2004.831600`.

[37] Amir Amirkhany. "Basics of Clock and Data Recovery Circuits: Exploring High-Speed Serial Links". In: *IEEE Solid-State Circuits Magazine* 12.1 (2020), pp. 25–38. DOI: `10.1109/MSSC.2019.2939342`.

[38] Jiang Gong et al. "A Low-Jitter and Low-Spur Charge-Sampling PLL". In: *IEEE Journal of Solid-State Circuits* (2021), pp. 1–1. DOI: `10.1109/JSSC.2021.3105335`.

[39] Sung-Yong Cho et al. "A 2.5–5.6 GHz Subharmonically Injection-Locked All-Digital PLL With Dual-Edge Complementary Switched Injection". In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 65.9 (2018), pp. 2691–2702. DOI: `10.1109/TCSI.2018.2799195`.

[40] Gang Xu and Jiren Yuan. "Comparison of charge sampling and voltage sampling". In: *Proceedings of the 43rd IEEE Midwest Symposium on Circuits and Systems (Cat.No.CH37144)*. Vol. 1. 2000, 440–443 vol.1. DOI: `10.1109/MWSCAS.2000.951678`.

[41] Bilal I. Abdulrazzaq et al. "A review on high-resolution CMOS delay lines: towards sub-picosecond jitter performance". In: *SpringerPlus* 5 (2016). DOI: `10.1186/s40064-016-2090-z`.

[42] Marcel Kossel et al. "A T-Coil-Enhanced 8.5 Gb/s High-Swing SST Transmitter in 65 nm Bulk CMOS With -16 dB Return Loss Over 10 GHz Bandwidth". In: *IEEE Journal of Solid-State Circuits* 43.12 (2008), pp. 2905–2920. DOI: `10.1109/JSSC.2008.2006230`.

[43] P. Heydari and R. Mohanavelu. "Design of ultrahigh-speed low-voltage CMOS CML buffers and latches". In: *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 12.10 (2004), pp. 1081–1093. DOI: `10.1109/TVLSI.2004.833663`.