POWER AND NOISE CONFIGURABLE PHASE-LOCKED LOOP USING MULTI-OSCILLATOR FEEDBACK ALIGNMENT

CHRISTOPHER WILLIAMS

A Thesis
in
The Department
Of
Electrical and Computer Engineering

Presented in Partial Fulfillment of the Requirements for the
Degree of Master of Applied Science (Electrical and Computer Engineering)
at

Concordia University
Montréal, Québec, Canada

August 2013

CONCORDIA UNIVERSITY
SCHOOL OF GRADUATE STUDIES

This is to certify that the thesis prepared

By: Christopher Williams

Entitled: "Power and Noise Configurable Phase-Locked Loop Using Multi-Oscillator Feedback Alignment"

and submitted in partial fulfillment of the requirements for the degree of

Master of Applied Science (Electrical and Computer Engineering)

Complies with the regulations of this University and meets the accepted standards with respect to originality and quality.

Signed by the final examining committee:

Dr. P. Valizadeh	Chair
Dr. P. Bianucci	Examiner, External To the Program
Dr. Y. R.Shayan	Examiner
Dr. G. Cowan	Supervisor
Dr. O. Liboiron-Ladouceur	Supervisor

Approved by: _____
Dr. W. E. Lynch, Chair
Department of Electrical and Computer Engineering

August2013

_____
Dr. Christopher Trueman
Dean, Faculty of Engineering and Computer Science

ABSTRACT

POWER AND NOISE CONFIGURABLE PHASE-LOCKED LOOP USING MULTI-OSCILLATOR FEEDBACK ALIGNMENT

Christopher Williams

On-the-fly data rate changes allow for the data rate to be lowered when peak speeds are not needed. A PLL is presented that contains a plurality of sub-VCOs, each able to be enabled or disabled. The goal of this technique is having the power dissipation proportional to the data rate, in order to obtain a fixed energy per transmitted bit. The proposed architecture accomplishes data rate changes by quickly reconfiguring itself and exploiting known power / jitter trade-offs in circuit design. The proposed architecture can be applied to either electrical or optical serial links that do not contain a forwarded clock.

By relaxing the jitter constraints at lower data rates, the receiver can enter a low-power mode enabling energy savings when maximum data rates are not required. A bank of sub-VCOs is introduced and can be brought up to speed and connected. An activation procedure and compensation methods have also been introduced in order to avoid arbitrary phases during start-up, which would lead to large phase excursions. Simulations show that by enabling the high-performance mode, data rates of 25 Gb/s are able to be obtained in a CDR setting. In the low power mode, the jitter increases by 1.5 times but the power reduces by 46%. In this mode, the architecture can support data rates of 12.5 Gb/s. Therefore, this system responds to the need of improving energy efficiency in receivers by allowing a dynamic reconfiguration of the circuit; varying power in response to jitter specifications.

ACKNOWLEDGEMENTS

I would like to thank my supervisor, Dr. Glenn Cowan, for his guidance and support during my undergraduate and Master's studies. His interest in electronics and fun-natured personality has made this a memorable experience. I would like to thank my other supervisor, Dr. Odile Liboiron-Ladouceur, for her help and guidance throughout this project as well.

I would also like to thank Mr. Ted Obuchowicz for his continued help with the CAD tools, as well as his support as an instructor and friend. His expertise was very useful at all points of the design.

I also want to extend my thanks to the members and "short-term members" of the VLSI group, with whom I shared many fun-filled experiences. I consider you all friends and wish you all the best.

Finally, I would like to thank my family and friends for encouraging me to continue my education and reach for my full potential. Edith Groulx-Robert deserves special thanks as she has been there every step of the way, and has made sacrifices in order for me to pursue this degree.

CONTENTS

List of Figures

List of Tables

List of Acronyms

BER	Bit Error Rate
CDR	Clock and Data Recovery
CML	Current Mode Logic
DLL	Delay-Locked Loop
I/O	Input / Output
MOSFET	Metal-Oxide Semiconductor Field Effect Transistor
PLL	Phase-Locked Loop
UI	Unit Interval
VCO	Voltage-Controlled Oscillator
XOR	Exclusive OR (logic gate)

CHAPTER 1

Introduction

1.1 Research Motivation

The world is now centered on the flow of information. What information is available to a person and how fast they can obtain it is the focus of many research topics. Just as a chain is only as strong as its weakest link, a high speed microchip is only as fast as its slowest I/O.

To recover serial data received without a forwarded clock, a clock and data recovery (CDR) scheme is required. This is true for serial data over both optical and electrical links. As CDR implementations at data rates above 25 Gb/s are now common place, research is focused on minimizing the power dissipation and chip area in integrated designs. However, it is a fact that more energy must be dissipated in order to properly receive data at higher rates. In the default approach for some implementations of variable data rates, fixed power dissipation is obtained by sending filler bits to create larger bit periods, and thus lower data rates. This however still requires the same number of bits to be sent as for higher data rates, except most of the data sent are copies of transmitted bits in integer multiples to lengthen the bit period. Nevertheless, a user may not always require the maximum data rate. Figure 1-1(a) illustrates a hypothetical situation and presents the required data rate versus time. Figure 1-1(b) compares the fixed power and proportional power schemes and shows that the latter follows the data rate. This is

referred to as energy proportionality or fixed energy per transmitted bit. The area between the fixed and proportional power curves is wasted energy that can be saved.

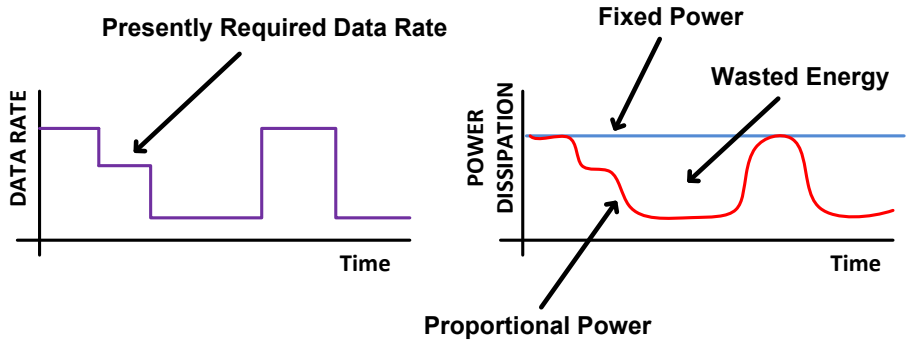

Figure 1-1: Example of power savings using on-the-fly data rate changing.

This power scaling can be realised through on-the-fly data rate changes that keep the energy per transmitted bit constant. On-the-fly data rate changes require a system capable of reconfiguring itself quickly enough to insure that power savings outweigh the overhead circuitry and system penalty due to the latency of switching to different data rates. The system must also be transparent enough not to disturb the dynamics of any active control loops. To scale the power with the data rate, it is necessary to carefully consider where this power savings will come from.

For example, *Toifl, et al.* has implemented a wide-band phase-locked loop (PLL) in order to suppress voltage-controlled oscillator (VCO) phase noise from a fixed-frequency, multi-phase compact ring VCO [1]. The reported energy of 1.8 pJ/bit is one of the lowest power CDRs at the time of writing. However, the power used by the VCO represents a significant part of the overall receiver power.

Other works proposed methods whereby CDRs are capable of multi-rate operation using fixed clock rates. As an example, *Rodoni, et al*. presented a CDR based on a delay-locked

loop (DLL), capable of discarding samples from an Alexander phase detector [2]. This meant that the system was able to be transformed into a quarter, half or full rate receiver by selecting which samples to discard. Despite this flexibility, the power did not scale with the data rate. Furthermore, three times as much power as in [1] was dissipated, mainly due to the power hungry blocks such as the phase rotators and phase interpolators [1]. Because the power reduction scheme consisted of deactivating some circuit blocks when working at lower data rates instead of targeting the VCO which represented a larger fraction of the system power, the power did not proportionally scale with the data rate. In addition, *Rodoni, et al*. does not discuss the speed at which the data rate can be changed [2].

In order to create a low power CDR and extend it to multiple data rates, the techniques in *Rodoni, et al*. can be applied to the architecture in *Toifl, et al.* [1]. A modified version of the CDR in *Toifl, et al.*is shown in Figure 1-2. This CDR operates as a $1/8^{th}$ rate architecture [1]. By disabling clock signals and capturing fewer data/edge samples per PLL clock period, the input data rate can be reduced but with modest power reduction. However, since the power dissipation of the PLL will remain constant, it will limit power reduction at lower data rates.

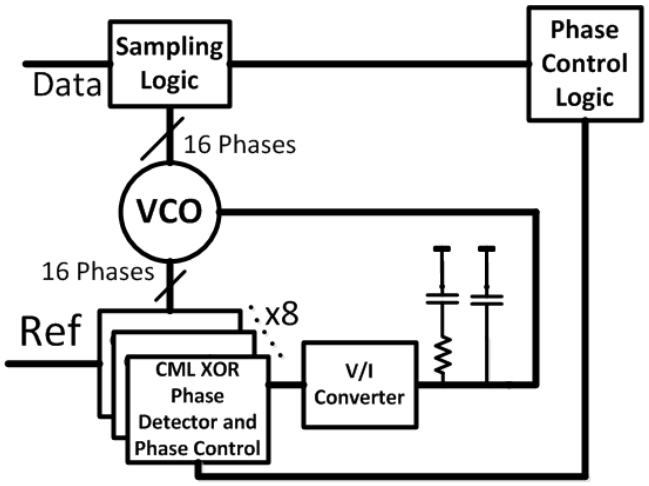

Figure 1-2: Modified CDR based on *Toifl, et al*. [1].

In *Mansurl, et al.*, highly scalable circuits were implemented with aggressive voltage supply scaling from 0.6V to 1.08V for use in memory and I/O circuits [3]. The output drivers used variable current mode and voltage mode circuits to enable noise rejection and allow power performance to be traded off depending on the required data rate. This variable data rate scheme led to a non-linear power efficiency of 2.6 pJ/bit at16 Gb/s down to 0.8 pJ/bit at 4 Gb/s, and 1 pJ/bit at 2 Gb/s. The clocks used injection locking techniques and ran at $1/4^{th}$ the data rate. This work demonstrated that there are industry interests in the area of data rate scaling, which is a necessary obstacle to overcome in order to reduce power density on chips and allow the trend of device and I/O scaling to continue.

1.2 Related Work

The power dissipation of a PLL is set by the maximum tolerable jitter when the CDR is operating at its highest data rate. If the input data rate is lowered, assuming a constant eye opening relative to the unit interval (UI), the PLL's timing jitter can increase without

degrading bit-error rate. This is illustrated in Figure 1-3, where the darkened grey squares represents bits of data.

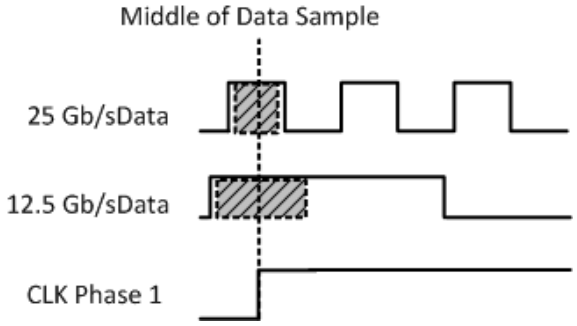

Figure 1-3: Unit interval versus data rate.

In order for proper clocking to occur, the clock must latch the data during this grey area, preferably towards the middle of the grey square. Jitter is the superposition of noise in time on top of the clock signal, which causes timing fluctuations in the crossings of the eye diagram. In the presence of higher jitter, data latching becomes less likely to occur without errors. However, at the lower data rate this type of clock jitter can be acceptable. This suggests that the jitter requirement of the fixed-frequency PLL can be relaxed when the data rate is lower.

In *Hajimiri, et al*., it was found that there exists a noise / power trade-off in circuit design [4]. It was found that if the requirements of a system called for tight jitter specifications during design time, then more energy would need to be dissipated in order to meet these conditions. If however, low power was more of an interest, then by reducing the effective size of transistors the power dissipation would reduce, but the noise would then increase.

With *M. Behbahani and G. E. R. Cowan*, the concepts in *Hajimiri, et al*. were applied to a variable ring oscillator [4,5]. By using a bank of uniformly weighted sub-VCOs,*M. Behbahani and G. E. R. Cowan*were able to connect up these sub-VCOs to increase the effective size of the transistors in the ring post fabrication [5]. This was done using switches to control which sub-VCOs would be interacting with the active nodes. By closing the switches, the effective size of the transistors in the ring increases and will follow the theory presented by*Hajimiri, et al*. [4].Figure 1-4 shows the schematic of this work.

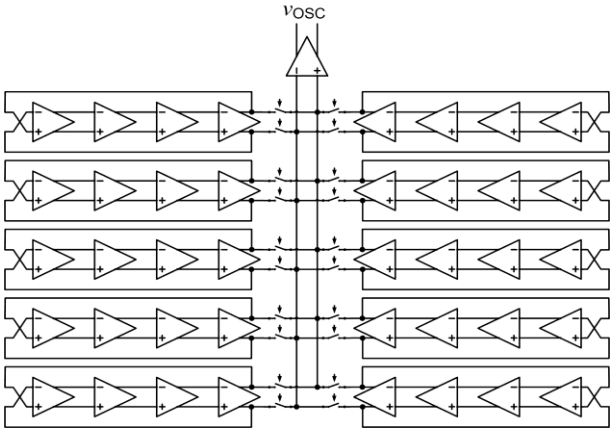

Figure 1-4: Sub-VCOs for static noise/power tuning [5].

Using this approach, *M. Behbahani and G. E. R. Cowan*were able to statically tune the power or noise performance of the oscillator to achieve a specified phase noise window with minimum power [5].This proved to be successful; however, it was not designed for dynamic connection of rings as the phases of deactivated rings would start up in an arbitrary state. If this design were to be extended to on-the-fly data rate changes, a method would need to be implemented that would avoid arbitrary phase start-up of the sub-VCOs, and thus minimize the impact on the system during reconfiguration.

Other research has focused on burst-mode type CDR architectures. Burst-mode CDRs present an alternative method to varying the data rate. In some burst mode architectures, data can be queued up and sent in packets at the highest data rate, and then the link can be shut down. Alternatively, *S. Kobayashi and M. Hashimoto* presented a burst-mode CDR that can vary the data rate quickly using a bit rate discrimination circuit [6]. The nature of the burst-mode architecture allows for quick start-up, and *S. Kobayashi and M. Hashimoto* reports a capability of varying the rates and recovering data on the second transmitted bit correctly [6]. Reported issues with this method stem from the fact that the gated oscillator will inherit all noise associated with the input data stream. Most of the burst-mode work reviewed was not very competitive in terms of power dissipation, and so this architecture was not chosen to be investigated further.

1.3 Thesis Contributions

This research work presents a novel architecture for providing a clock recovery system for use in on-the-fly data rate changes. The main contributions of this paper are:

- Novel circuit architecture for connecting multiple sub-VCOs together.
- A methodology for bringing a newly enabled sub-VCO up to speed, and readying it before connection to other already running sub-VCOs.
- Identification of the problems with this architecture, as well as innovative compensation methods to alleviate them.
- Implementation of the architecture in an integrated circuit using TSMC 65nm technology.

- Provisional U.S. patent filed by the Concordia office of research and Gestion Valeo.
- Presentation and publication of this work at IEEE Mid-West Symposium on Circuits and Systems (MWSCAS) 2013 in Columbus, Ohio, USA.

CHAPTER 2

Background Theory and Fundamentals

This chapter describes the general theory pertinent to the presented work. An overview of serial data transmission will be followed by general discussions on clock and data recovery (CDR) architectures. A clock generation circuit consisting of a phase-locked loop (PLL) will then be discussed, along with relevant information on each of the blocks contained within. Figure 2-1 shows a simplified CDR, containing a PLL using an external reference clock.

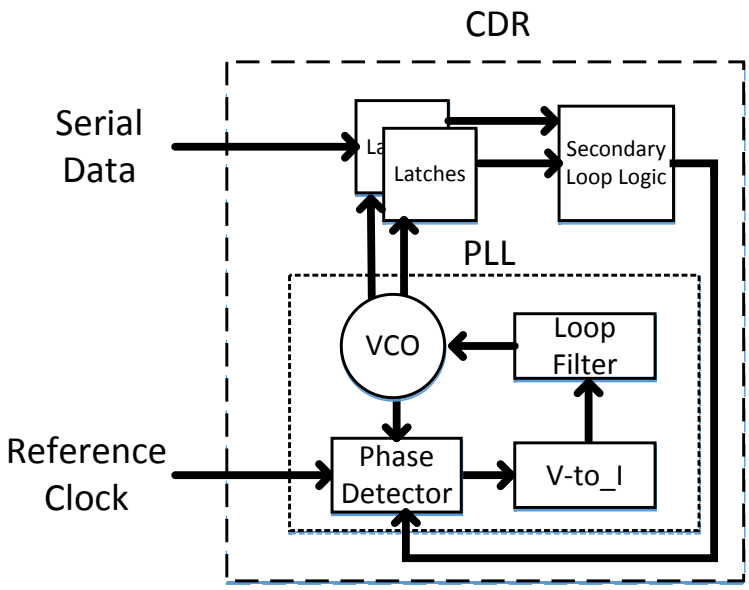

Figure 2-1: CDR structure encompassing a PLL.

This simplified model can be used as a reference for this chapter, as a macroscopic view of the entire system will lead to a microscopic explanation of each component in the system.

2.1 Random Data and Characterization

Receiving serial data consists of receiving one random bit of information after another, as each bit received (without accounting for maximum run lengths) has equal probability of being a 1 or 0. Normally, charge and discharge times are governed by the current entering or exiting a capacitive node, and thus are proportional to the product of the resistance and capacitance, referred to as the time constant. After five time constants, the voltage of the node asymptotically reaches a final state. If however, the transmitted data rate is too high for the voltage at the node to change before the next stage can resolve the intended value, then errors can occur. Another problem, not discussed in detail in this work, is the fact that previous data will have an effect on the present bit. This is due to the effect of the previous charge/discharge cycles on the present one, leading to Inter-Symbol Interference (ISI). ISI is the signal overlap of one bit to the adjacent bits.

Information concerning the quality of the received data can be found in an eye diagram. The eye diagram provides a visual representation of the ISI and noise in the random bits received. The diagram is created by folding all received bits into an interval two to three bits wide [7]. An example of a real eye-diagram is shown in Figure 2-2 [1]. The eye diagram gives insight into several important characteristics such as eye opening (top-bottom opening) and timing jitter (thickness of line between each eye).The UI is the spacing between *L Open* and *R Open*.

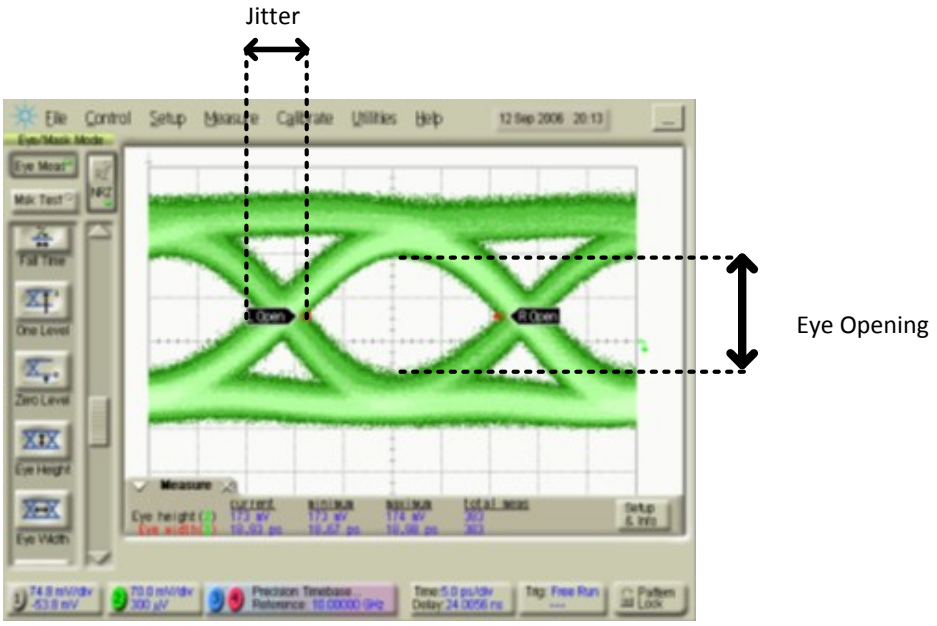

Figure 2-2: Example of eye diagram from*Toifl, et al*. [1].

2.2 CDR Architectures

2.2.1 Architectures with and without a Reference Clock

In order for a CDR to correctly receive data, it must be able to regenerate a clock if one is not sent. This implies having a type of oscillator in the receiver circuit, whose frequency can be controlled. As the approximate frequency of the transmitted data is known, then one method to accomplish clock alignment is to use an external reference clock on the receiving end. The reference clock is a locally generated clock for the receiver, and is assumed to be as noise free as possible. This then allows the locally controlled oscillator to obtain a frequency lock with the reference clock. A delay-locked loop (DLL) is a delay line that consists of a chain of inverters, and propagates the reference signal along the line but at delayed intervals. This allows for a particular delay to be selected and locked to the reference clock's phase. Figure 2-3shows an example of

an implementation using a reference signal and a DLL. This method allows for one switch to close at a time, in order to select the phase that will clock the data. This example does not have a phase interpolator, which limits the granularity to 1 inverter delay. Phase interpolators allow the weighted combination of two phases, which allow for an output phase equal to one of the two inputs or any combination of the two. For an alternative to CDRs requiring reference clocks, there exists a reference-less design. This is shown in Figure 2-4, and phase locks by using the incoming data transitions.

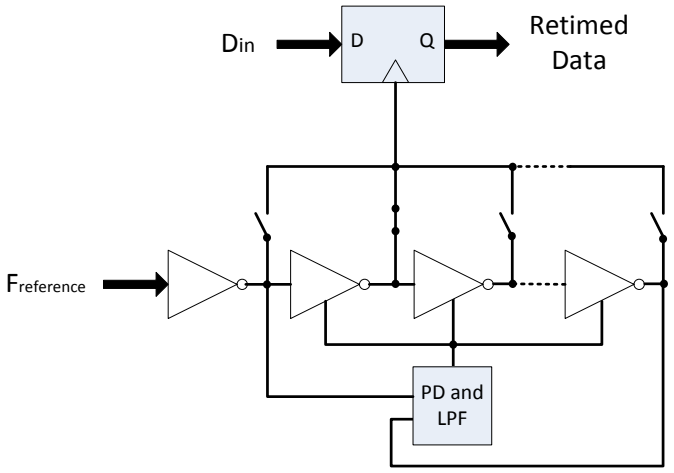

Figure 2-3: DLL based CDR with external clock reference.

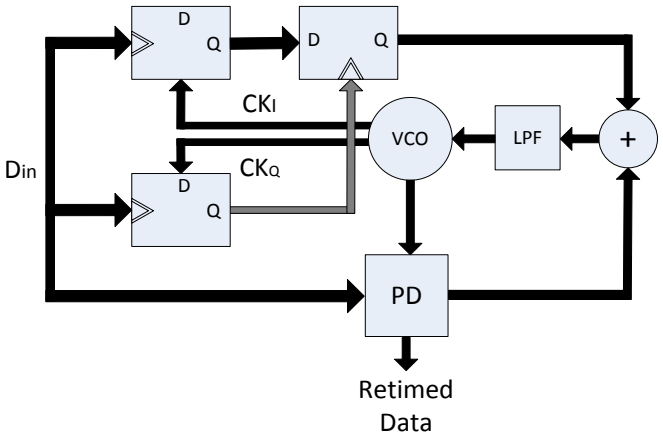

Figure 2-4: CDR without external clock reference [7].

The use of a reference clock in a CDR makes the design of the system simpler, as circuits can be designed to work with the periodic reference signal available. The availability of a reference will improve phase noise, as reference-less CDRs have long run lengths to contend with and thus have more of a chance to drift. However, when using a reference clock there are obstacles as well. If the reference is integrated into the chip, then there will be a need for more area to be reserved for it. If the reference is external, such as a crystal oscillator, then this will add to the cost of manufacturing. The reference should also be as clean as possible, for reasons that will be discussed further in this paper. This adds constraints to the reference clock, and will most likely increase the power dissipation of the system.

2.2.2 Sub-Rate Clocking

The availability of multiple phases of an oscillator gives rise to other CDR designs, such as sub-rate architectures. Sub-rate operation allows a high frequency serial stream to be broken down into several parallel streams, each at a fraction of the high frequency. This relaxes the specifications of blocks since the data rate of each parallel stream is lower. Such terms as 1/4th and 1/8th rate are terms used for sub-rate specifications, and refer to the fractional data rate at which the parallel links are operating at compared to the input serial data rate. These concepts will be expanded upon further in the following sections.

2.2.3 Alexander-Style Phase Detection

In order to receive data correctly, clocks must provide a reference for the system to use. This reference can be used for two separate purposes, for edge sampling or data

sampling. In edge sampling, an Alexander-style detection scheme can be used, the workings of which is illustrated in Figure 2-5below.

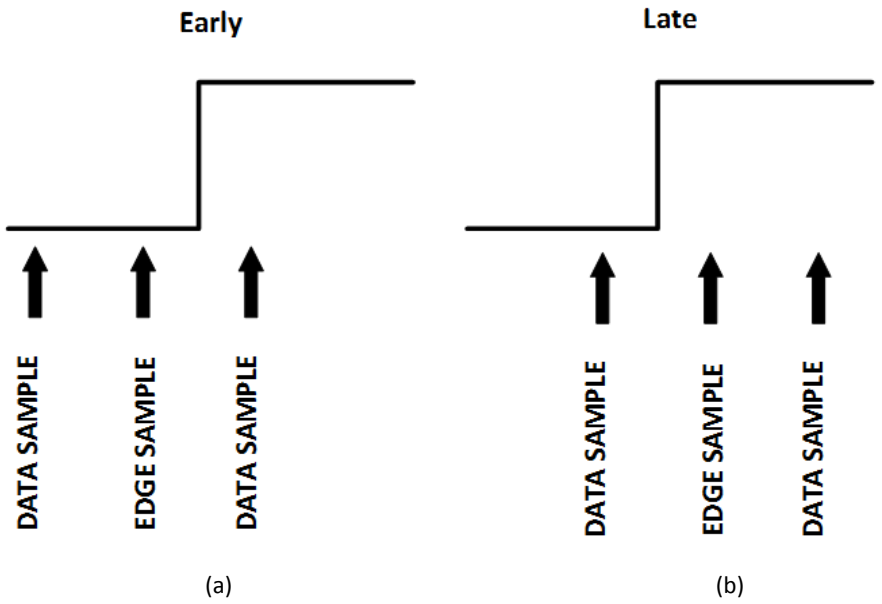

Figure 2-5: Alexander phase detector sampling data (a) early and (b) late.

In order to lock onto the correct sampling point for receiving data, the edge information is continuously updated, as seen in Figure 2-5 (a), and the clock sampling point is moved until the edge is surpassed in (b). Once this happens, the phase decrements until surpassing once again, and then repeats the process. In this way, during the lock condition, the phase will continuously alternate between ahead of and behind the data transitions. Once the edge information has locked, then data sampling is available for information extraction. This samples the voltage of the data signal and determines whether it is logic high or logic low. As ISI is usually present in data, more complex schemes may need to be used in order for correct transmission.

2.3 PLL and Loop Dynamics

A PLL is a control system which uses negative feedback in order to lock a VCO contained within the PLL to a reference input. This locking refers to both the output and input having constant phase differences. Output frequency can differ from the input frequency if a divider is added to the feedback path. In this case, the output frequency will be a multiple of the input.

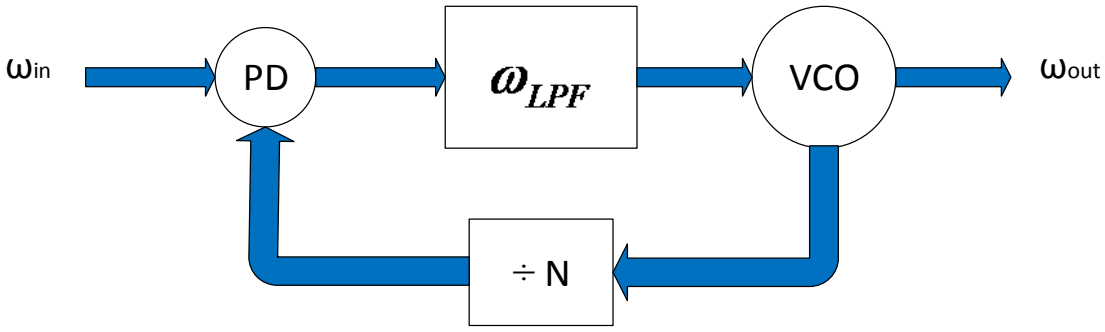

Figure 2-6: PLL block diagram.

A PLL block diagram is shown in Figure 2-6 where PD is the combination of a phase detector and charge pump, ω_{LPF} is the low-pass loop filter, VCO is the voltage-controlled oscillator and *N* is the division ratio. The function of the phase detector is to relay information about the relationship between the two input signals, which in this case, is the reference and the VCO signal divided by a value N. Several different varieties of phase detectors exist, each with their own strengths and weaknesses.

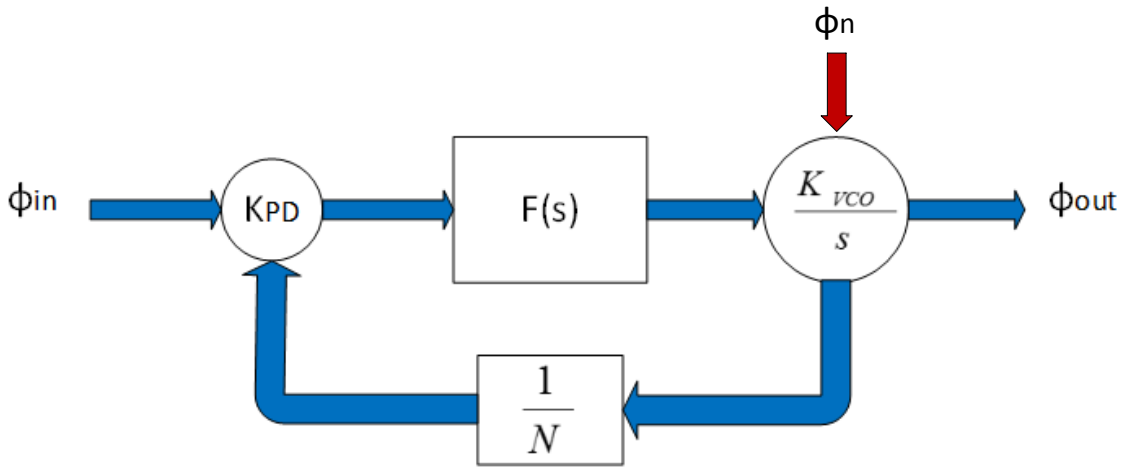

Figure 2-7: PLL as a control system.

As a PLL is concerned with the phase difference of two signals, Figure 2-7 is made to relate the phases of the input and output in a control loop containing *s-domain* parameters. By deriving the closed loop transfer function, it can be found that

$$H(s) = \frac{\phi_{OUT}}{\phi_{IN}} = \frac{\frac{K_{VCO}K_{PD}F(s)}{s}}{1 + \frac{K_{VCO}K_{PD}F(s)}{Ns}} = \frac{K_{VCO}K_{PD}F(s)}{s + \frac{K_{VCO}K_{PD}F(s)}{N}} \tag{2-1}$$

where s = jω. Substituting and taking the limits as $s \to \infty$, it becomes evident that (2-1)is that of a low pass filter. As the transfer function is comparable to a low pass filter, this leads us to the conclusion that any noise coming from the reference will be modulated onto the VCO output, up to the loop bandwidth. Anything higher in frequency will be rejected. This, however, is also true for phase correcting information that is fed back from the output to the input. If a high frequency error in the output is present, then this is filtered out by loop and thus the PLL cannot correct for this. By looking at an injected disturbance, ϕ_n, and the result this has on the output, the following noise transfer function is found to be:

$$N_{TF} = \frac{\phi_{OUT}}{\phi_n} = \frac{1}{1 + \frac{K_{PD}K_{VCO}F(s)}{Ns}} = \frac{s}{s + \frac{K_{PD}K_{VCO}F(s)}{N}} \tag{2-2}$$

This leads to a transfer function that is similar to a high pass filter. Therefore any high frequency noise generated in the oscillator is directly passed to the output, leaving the loop unable to correct for these errors. Thus, noise within the loop bandwidth is suppressed.

2.3.1 Phase Detector

The phase detector is the first block that will be discussed. The purpose of the phase detector is to compare the inputs and to provide an output that expresses the difference of the two. The phase detector can come in two types, linear or non-linear. The linear type provides information on both the sign and magnitude of the phase error of the two input signals. An example of this would be the *Hogge* detector. The output voltage is averaged to obtain magnitude information, and is shown in Figure 2-8 (a). The non-linear type will only provide sign information on the phase error. Examples of this would be the *Alexander* and *bang-bang* phase detectors. The ideal output of this type of detector is shown in Figure 2-8 (b) [7].

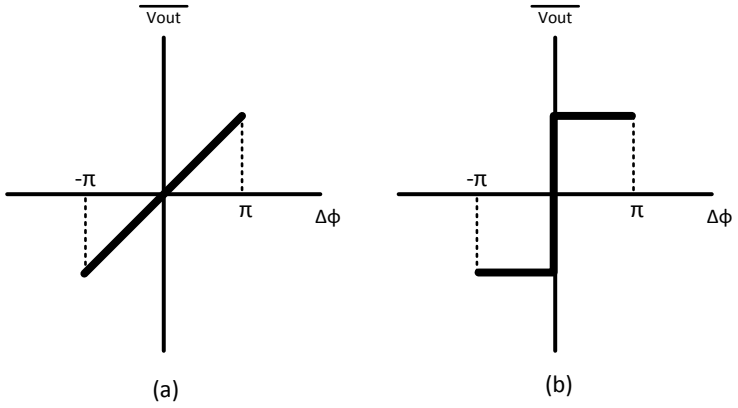

Figure 2-8: Phase detector average voltage output (a) linear (b) non-linear type.

2.3.2 Loop Filter

The loop filter and its design is a vital part of a properly functioning PLL. It is responsible for establishing the loop bandwidth of the system. The filter can receive either a voltage or current input, depending on the output of the phase detector, and transform it into a voltage for the VCO.A resistor can be added in series to this capacitor to add a zero to the PLL filter transfer function, shown below in (2-3).

$$Z(s) = \frac{sRC + 1}{sC} \tag{2-3}$$

As each unit of charge is given or taken from the loop filter, the resistor will respond to this current flow. Since the resistance develops a voltage across the terminals of the added resistor due to Ohm's Law, the consequence of this is a ripple on the loop filter voltage, also referred to as the control voltage. In order to reduce this ripple, a second capacitor in parallel can be added, shown in Figure 2-9 as C2.

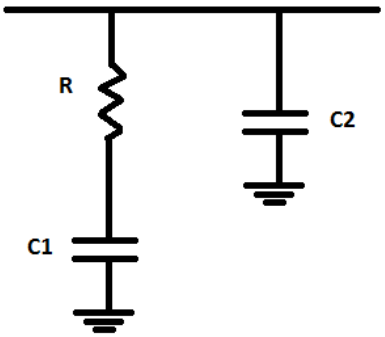

Figure 2-9: Second order loop filter for current output phase detectors.

The overall transfer function of the loop filter is presented in (2-4), with the result being in Ohms, as the output is a voltage and the input is a current. The extra capacitor adds a pole in the transfer function and thus must be chosen carefully.

$$Z_T(s) = \frac{V_{OUT}}{I_{IN}} = \frac{1 + sRC_1}{s^2 RC_1C_2 + s(C_1 + C_2)} = \frac{s\left(\frac{1}{C_2}\right) + \frac{1}{RC_1C_2}}{s^2 + s\left(\frac{C_1 + C_2}{RC_1C_2}\right)} \tag{2-4}$$

Figure 2-10 is a MATLAB generated plot that shows the effect of adding the parallel capacitor.

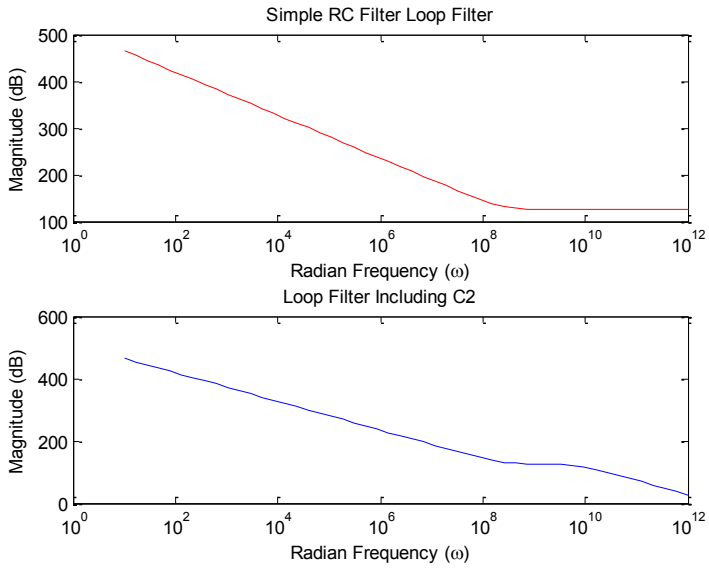

Figure 2-10: Bode plot of first order versus second order loop filter.

The loop filter is calculated to achieve a cut off frequency of 1/10th the reference frequency. This is due to the fact that the equations derived for the PLL are for a continuous-time system, whereas the sampling nature of the PLL makes it a discrete-time system. However, as shown in *J. P. Hein and J. W. Scott*, if the switching nature of the PLL is kept below 1/10th of the reference, then the PLL can be modeled as a continuous-time system with fairly high accuracy for Laplace domain analysis [8]. It is possible to violate this ratio, but this would require a more complex model that takes into account the switching nature of the loop filter voltage.

2.3.3 Phase Noise

Phase noise is a frequency domain representation of fluctuations in the oscillator's phase. Phase noise is superimposed on top of an input signal, V_o, giving rise to a mathematical definition of

$$V_{OUT} = V_O \cos[\omega_O t + \phi_n(t)] \tag{2-5}$$

where V_{out} is the output signal, ω_o is the oscillation frequency of the VCO and ϕ_n is a zero-average random phase component [7].A phase noise plot can be generated, and gives an overall picture of noise found at different offsets from the center frequency, or oscillation frequency. An ideal oscillator will produce tones at the frequency of oscillation, shown in Figure 2-11 (a). However, when phase noise is superimposed upon this frequency spectrum, side skirts appear ($\Delta\omega$) around the ideal tone (ω_o), as shown in Figure 2-11 (b).

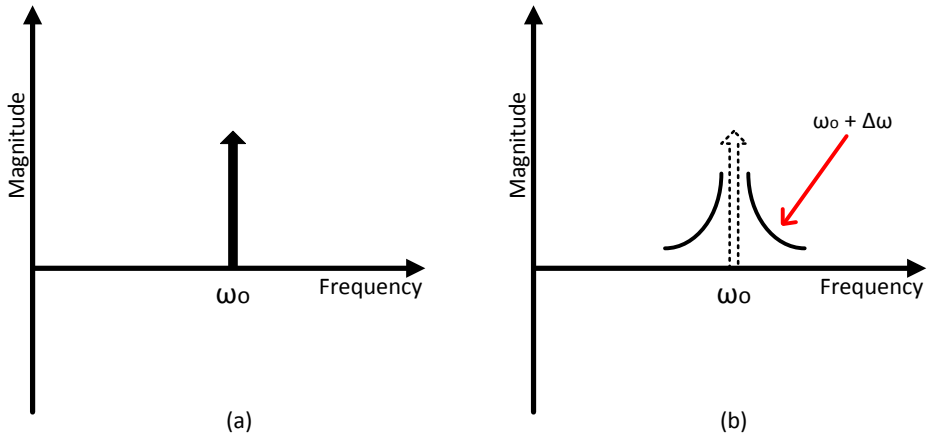

Figure 2-11: Frequency spectrum of (a) ideal signal (b) signal subjected to phase noise.

In order to obtain jitter information from phase noise, the time domain representation of phase noise, the following transformation must be applied [16]:

$$Jitter_{RMS} = \frac{2 \cdot \sqrt{10^{\frac{\int_{f1}^{f2} L(f)df}{10}}}}{2\pi f_o} \tag{2-6}$$

where f_1 and f_2 are the limits of integration of the phase noise spectrum, $L(f)$ is phase noise and f_o is the frequency of oscillation. Jitter reveals information about the zero-crossings of a waveform. If the zero-crossings of an output signal are compared to the zero-crossings of an ideal signal, the deviation from the ideal signal will be referred to as absolute jitter. If each period of the signal is compared to the proceeding period, then this will result in a cycle-to-cycle jitter measurement, where no reference signal is required [7]. Jitter can appear as two types, random or deterministic. Deterministic jitter is predictable, and can be reproduced under like-conditions. Random noise however is unpredictable, and is caused by random events.

2.3.4 Bandwidth Trade-Off

The loop bandwidth should be set to 1/10 of the reference frequency. This was shown in *J. P. Hein and J. W. Scott* to be a good practice due to the nature of the PLL [8]. Since there is no divider present in this design, the loop bandwidth could to be much higher than in other PLLs that contain dividers. Since the loop bandwidth is greater, more phase correction information is available to adjust the VCO due to the low-pass nature of the system allowing higher frequencies pass. On the other hand, this larger bandwidth also opens the door to a greater amount of low frequency noise from the reference to modulate the control line, and appear on the VCO output as noise.

2.3.5 Loop Behaviour

The PLL has the ability to lock its oscillator to a reference frequency provided that the VCO has a proper tuning range and the loop parameters have been correctly calculated. The PLL uses the phase detector in order to provide information on whether the VCO frequency is faster or slower than the input reference signal. If the VCO's frequency is higher than the reference frequency, then it is equivalent to the VCO accumulating phase quicker. The phase detector is then forced to slow down the VCO until the opposite happens and the reference is ahead of the VCO. When this happens, the phase detector can then make small increments and decrements in the VCO speed in order to keep it in phase with the reference frequency. This is analogous, as a first order approximation, to two cars trying to drive next to one another. The driver can look over in short intervals to the other car in order to obtain information on the placement of their car relative to the other car. If the driver is behind, the car can be sped up. If the driver is ahead, they can slow down.

In order for the control system to remain stable, some important parameters should be discussed. The grouping of the poles and zeros of the PLL results in a frequency plot showing the response of the system at a range of frequencies. Figure 2-12 shows two plots, the top one illustrates the magnitude of a system over a frequency range, where the crossover frequency is the frequency at which the magnitude is equal to 0 dB [15]. The bottom image is the phase of the system. The gain margin, GM, and phase margin, PM, are labelled in the figure. These variables are important as either of them can cause oscillation of the system, which would render the control loop unstable.

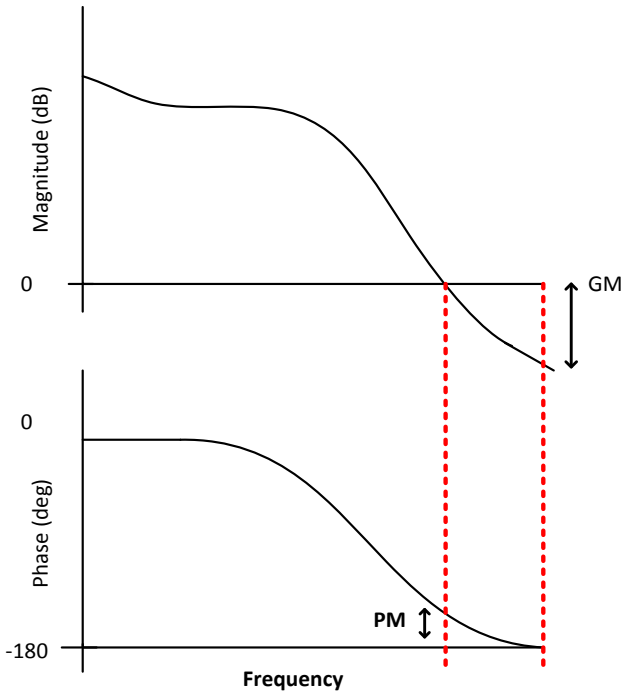

Figure 2-12: Gain (GM) and phase (PM) margins.

System stability implies that a bounded input will result in a bounded output. Using Barkhausen criteria, it must be ensured that if the phase difference between the input and output is 180^{o}, then the gain is not unity and vice-versa for system stability. The gain margin is then considered as the magnitude of gain present at the frequency where the phase difference is 180^{o}. The phase margin is the measure of excess phase, or how much less than 180^{o} the phase is at the frequency where the gain drops to unity, or 0 dB.

If there is one pole at the origin in the transfer function, then this corresponds to a type I PLL. The type I PLL gives zero steady state error for a phase step. If there are two poles at the origin, then the PLL becomes a type II system. This indicates that it can have zero steady state error for phase ramps. As the input to a PLL is phase ramp, type II PLLs are needed for steady state regular operation of a functional PLL.

2.4 Voltage-Controlled Oscillator (VCO) Theory

2.4.1 Operation

Voltage controlled oscillators can come in a variety of types, the two main ones being LC-VCOs and ring VCOs. The LC-VCO consists of a resonant circuit built of inductors and capacitors that oscillates at a frequency determined by the component values. A ring VCO establishes an oscillation frequency dependant on the time a signal takes to propagate through the circuit and continuously update each stage of the ring. The LC-VCO is usually known to achieve better noise performance, but consumes large area due to the need for inductors. Ring VCOs can be made more compact and implemented using CMOS gates. This can be an attractive feature as CMOS processes are used for digital circuits that normally are used in large quantities on a chip, and thus can take advantage of well-established device models.

Ring oscillators, at their base, can be implemented two ways to sustain oscillation. The first way, referred to as the single-ended approach, is illustrated in Figure 2-13.

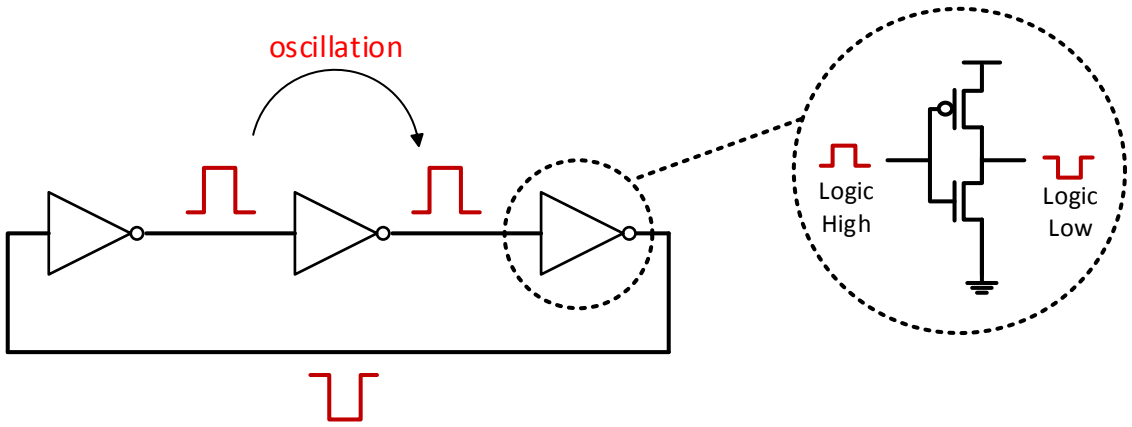

Figure 2-13: Ring oscillator.

With an even number of inverter stages, the loop "latches-up" and does not oscillate. If an odd-number of stages exists, the ring will start to oscillate. The frequency of oscillation will be determined by the delay of each stage, both the rising and falling portions, and are related by the following [9]:

$$f = \frac{1}{2Nt_d} \tag{2-7}$$

where N is the number of stages and t_d is the delay of the inverter. A second implementation also possible is called a differential ring oscillator. It consists of two single-ended rings, tied together with cross-coupled inverters. The cross-coupled pairs ensure that the two phases are 180 degrees apart. With a differential implementation, an even number of stages is possible, so long as another 180 degree phase shift is present elsewhere in the loop. This can be accomplished by crossing the forward path at one point in the ring, as presented in Figure 2-14.

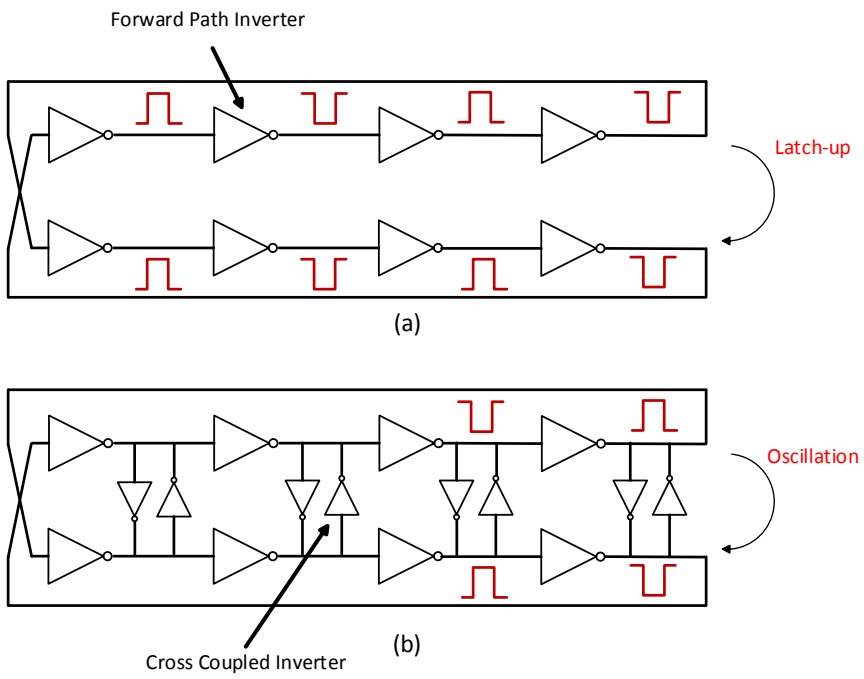

Figure 2-14: Differential ring oscillator (a) latched-up (b) oscillating.

Cross-coupled inverters do require careful implementation, as there exists a lower and an upper bound on the possible inverter strengths for correct operation. Consider the following: if the cross-coupled inverters are too weak compared to the forward path inverters, then it will be effectively turned into an 8-stage, single-ended ring. This will cause the ring to latch-up. If however the cross-coupled pair is made too strong compared to the forward inverters, then the cross-coupled inverters will prevent oscillation from occurring. This is due to the forward path being unable to overcome the drive strength of the cross-coupled pair.

2.4.1 Multi-Phase Generation

One important aspect of this project relies on its use of multiphase clock generation. This is similar to normal VCOs, but instead of only one phase brought out, each phase in

the ring oscillator is brought out. This provides multiple output clocks, each at the same frequency, but delayed by identical amounts.

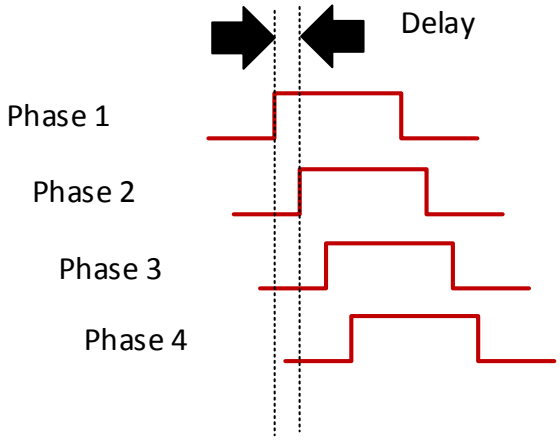

Figure 2-15: Multi-phase clock output example.

The image above illustrates four phases out of the possible sixteen available in this project, and is not drawn to scale. The delay, to a first order approximation, is equivalent to the delay of the inverter separating two stages of a ring oscillator.

2.5 Scaling

Impedance scaling is a common practice in filter design, which results in a scaling of noise, power and cut off frequencies. In filters, without changing the cut off frequency, resistors can be divided by K and capacitances can be multiplied by K in order to achieve better noise performance but with higher power dissipation. Given three transistors of width W, together they will behave as a larger transistor of width $3W$ in terms of power dissipation and noise performance.

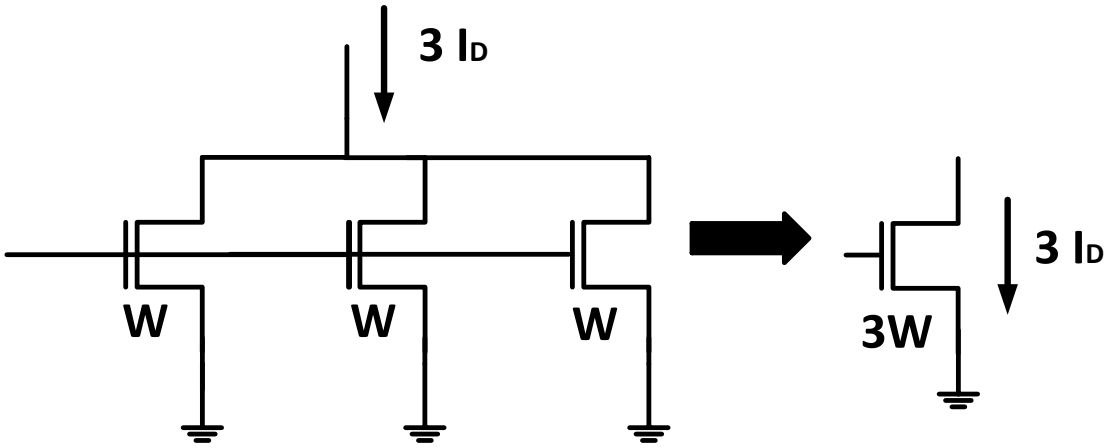

Figure 2-16: Impedance scaling.

When a transistor is larger, the current increases by three times along with the power dissipation. However, as a consequence of this, the mean squared noise current will be reduced by three times. This power / noise trade-off was shown in *Hajimiri, et al.* to be able to be extended to VCOs as well [4]. This work presented the following equation:

$$L\{\Delta f\} \approx \frac{8}{3\eta} \cdot \frac{kT}{P} \cdot \frac{V_{DD}}{V_{CHAR}} \cdot \frac{f_o^2}{\Delta f^2} \tag{2-8}$$

where $L\{\Delta f\}$ is the phase noise at a frequency offset (Δf) from the center frequency (f_o), V_{CHAR} is the characteristic voltage of a device, η is a proportionality constant, V_{DD} is the supply voltage and P is power. Keeping all variables constant but increasing the number of VCOs dissipating power from one to n, the following relationship is made:

$$L\{\Delta f\} \propto 10\log\left(\frac{1}{n}\right) \tag{2-9}$$

This means that if n VCOs are placed in parallel, the phase noise will reduce proportionally as stated above due to an increased amount of power dissipated.

2.6 Excursion Plots

In order to evaluate different methods of implementation and to plot the response to a disturbance caused by a VCO reconfiguration, an output plot was developed for visual comparisons. Excursion plots are graphs that describe the extent of deviation (in seconds) to which the output locked phase from the oscillator is from the input reference clock. In order to obtain these plots, the rising edges of the oscillator are examined for where they cross the $V_{DD}/2$ mark, or half of the supply voltage. Then, the same is done for the input reference clock. The difference of the two lines are then found and normalized to account for the offset, which in this case is due to the two signals being in quadrature. The usefulness of the plots is seen when a new sub-VCO is activated. In Figure 2-17 (a) a hypothetical case is shown where a comparison must be made between an ideal reference clock and the VCO output. In (b) the crossing information of the two signals presented in (a) are graphed. This shows how the VCO non-ideal waveform differs from the ideal reference clock waveform. In (c) the excursion plot is generated by taking the difference of the two lines found in (b).

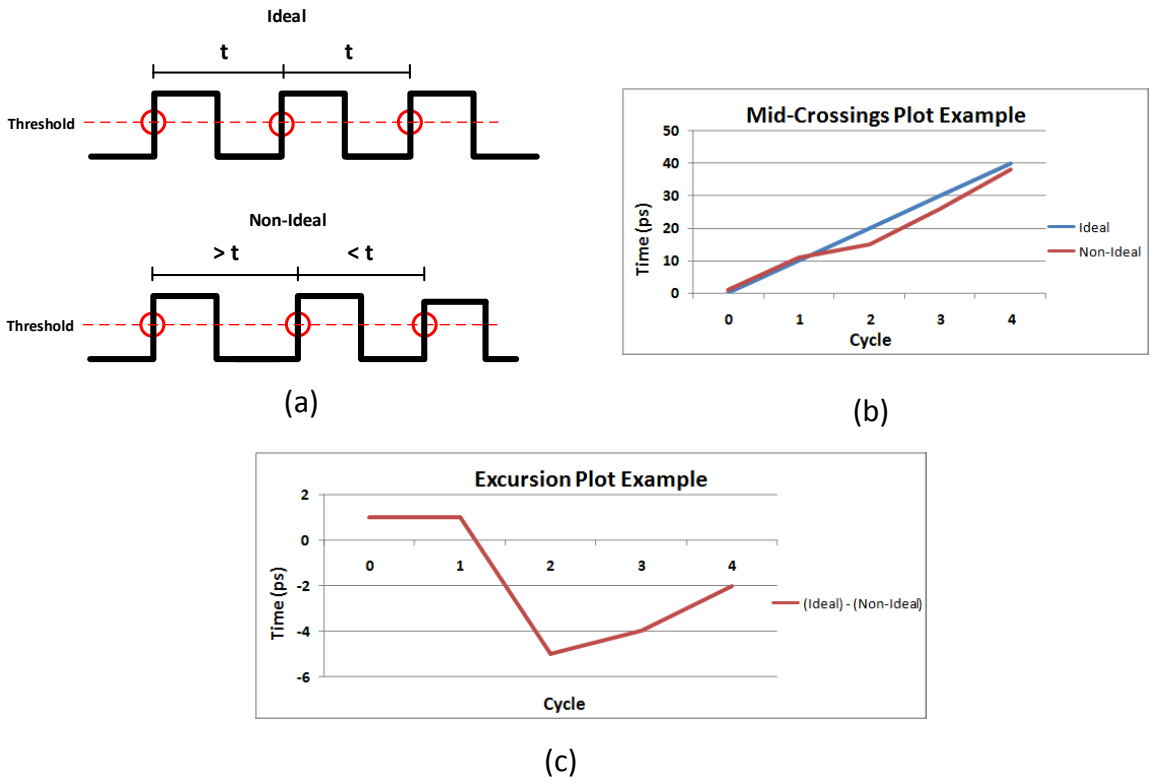

Figure 2-17: (a) Ideal clock versus one with jitter (b) theoretical mid-point crossings of reference and output (c) excursion plot generated from mid-point crossings.

If this excursion was related to a deviation from the center of a UI, needed for ideal data sampling, this deviation could mean incorrect sampling of data bits. For example, if 12.5 Gb/s data was being sampled, 16ps would be 20%or 0.2 UI. Figure 2-18 is taken from *Toifl, et al*. and is an example of a jitter tolerance diagram that was used in this project as a template for what can be expected of a real CDR [1].

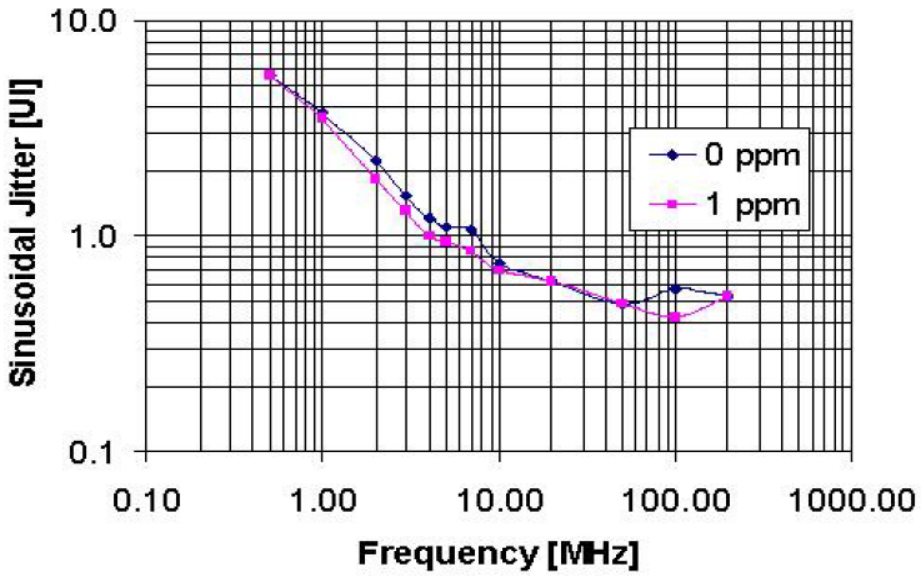

Figure 2-18: Jitter specifications of actual CDR in *Toifl, et al.*[1].

Taking the jitter tolerance found at 100 MHz and extending it further along the X-axis, the value of the jitter tolerance for the CDR at higher frequencies is found. As can be deciphered from this diagram, at higher frequencies, their CDR was able to tolerate 0.4 UI. This signifies that the CDR will incorrectly sample if nothing is done to keep the excursions within the jitter tolerances of the system.

CHAPTER 3

Design

In this chapter, details of the design will be presented. Insight into these design choices and reasoning behind designs will also be conveyed. The goal of the PLL was to meet the specification of the CDR in order to receive data at 25 Gb/s. This was to be done using a 1/4th or 1/8th sub-rate clocking scheme. The application of this PLL is for use in data centers, in order to respond to the need for power savings. The design was simulated using Cadence simulation tool using the TSMC 65nm technology.

3.1 Proposed Design

The proposed design of the supporting system for the presented VCO will now be introduced. Shown below in Figure 3-1 is the overall block diagram of the sub-VCOs in a PLL.

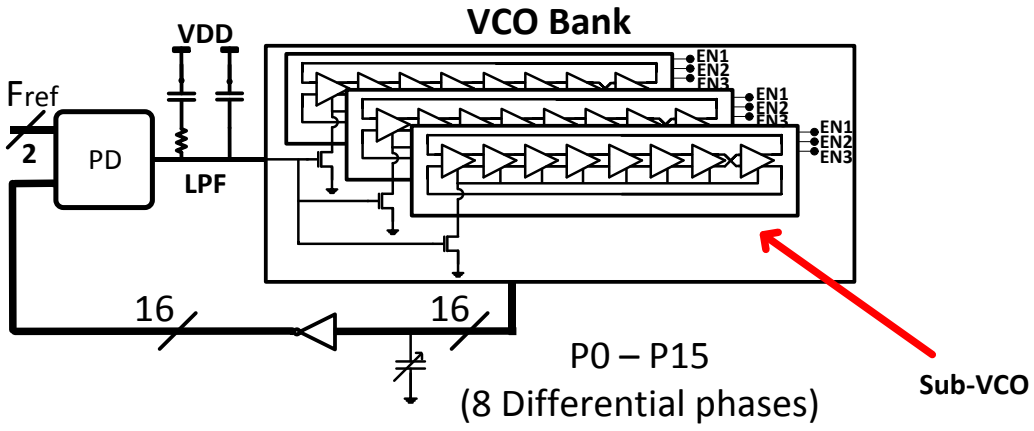

Figure 3-1: Proposed sub-VCO architecture in a PLL system.

Where *PD* is the phase detector, *LPF* is the low pass filter and the VCO bank is the combination of all three sub-VCOs. F_{ref} is the differential reference frequency of 3.125 GHz and VDD is the supply voltage of 1V DC. As no divider is used in the feedback path, the loop bandwidth is set to 300 MHz. In this proposed design, one sub-VCO in the VCO bank will always be running. If one wished to increase the data rate, the jitter tolerances would have to be tightened. In order to improve the jitter, another sub-VCO would be activated.

At the base, each of the 3 identical sub-VCOs is a current-starved 8-stage differential ring oscillator using static CMOS inverters. The nodes of the ring are then isolated from the rest of the circuit through the use of buffers. The result is 16 phase nodes per sub-VCO, as shown in Figure 3-2, labelled as V_A, V_B, etc. The corresponding phase node of each sub-VCO are then connected together in order to create a *Common Phase Node.* This node is shared between each of the sub-VCOs and one is created per phase, resulting in 16 *Common Phase Nodes*.

The current method of data rate changing is not yet known. As this is a proof of concept, we assume that there will need to be an automatic detection scheme in place that can detect and reconfigure the system quickly or that the system controller will signal data-rate changes to the PLL.

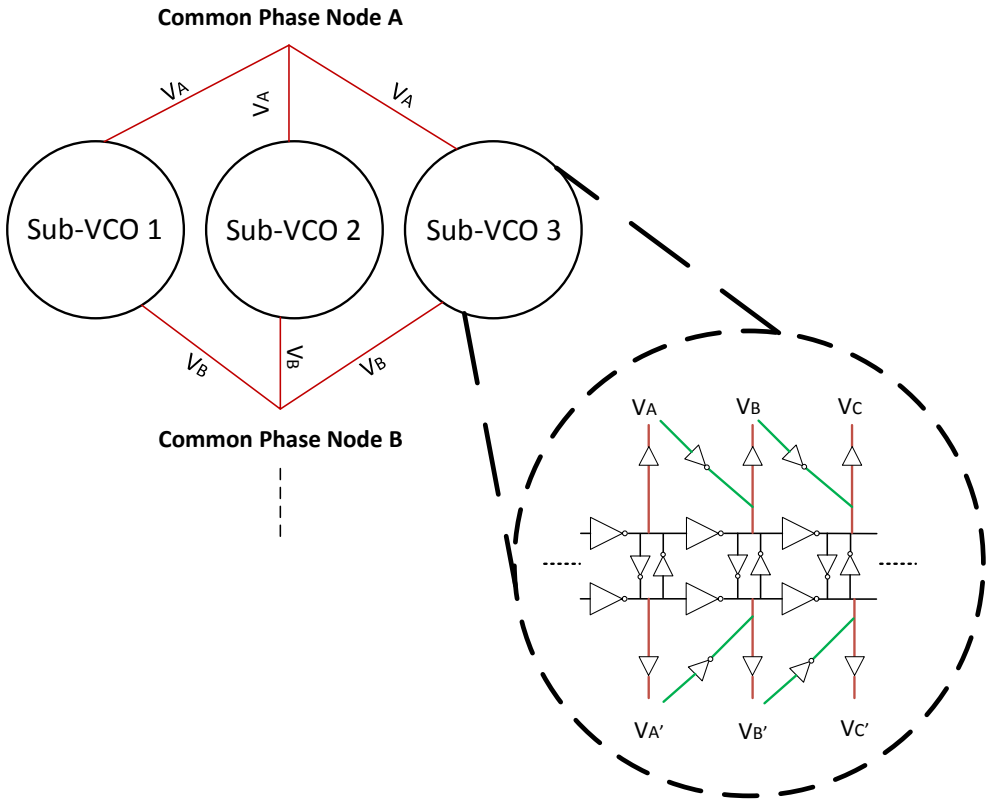

Figure 3-2: Proposed VCO bank architecture.

3.2 Voltage-Controlled Oscillator (VCO)

3.2.1 Basic Considerations

The frequency of operation and the phase noise of a VCO can be determined by not only the strength of each stage, but also the capacitive loading of each of the nodes in the ring. For this reason, it is important to select the ring strength carefully.

In differential inverter-based ring oscillators, each inverter has an ability to either source or sink current, the amount of which is determined by the size of the device. This sizing however contributes to node capacitance as well. This node capacitance will ultimately lead to a larger load that must be charged and discharged each cycle. For this reason,

importance must be given to this inverter sizing since it can directly affect not only power dissipation, but noise as well.

Consider the noise generated by a low pass filter using noise modeling analysis. Analysis shows an interesting conclusion, where the noise of the circuit is determined by the capacitor and not the resistance. This is due to the capacitor creating a filter where signals, including noise, above cut-off frequency are shorted to ground and eliminated. This leads to a trade-off, where a larger node capacitance will reduce the noise of the ring, but will also increase the power needed to charge the node each cycle as well as increase the per stage delay and thus lower the oscillation frequency.

3.2.2 Feed Forward VCO Architecture

Another VCO architecture exists that allows for higher frequencies to be achieved at reduced power, called a *feed-forward*(FF) architecture. This is illustrated below in Figure 3-3, where a signal is allowed to bypass the regular inverter path and is fed to an advanced stage of the ring. The arrow on the inverter is to signify a variable strength design where the drive strength is able to be changed, and thus increase or decrease the overall frequency of oscillation.

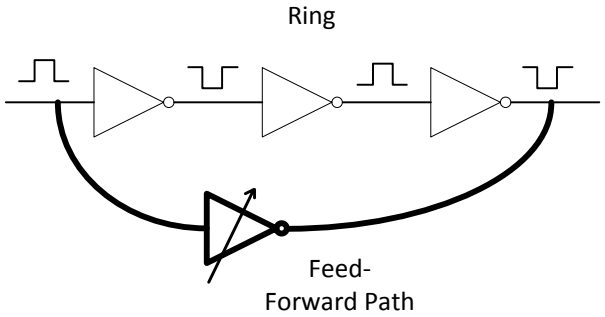

Figure 3-3: Ring oscillator using the feed-forward architecture.

This will allow for the node further down the path to begin charging before it would normally. This architecture was looked at in detail as well, as there was interest in seeing the advantages of this design. The following data in Figure 3-4 and Figure 3-5has been gathered to show the usefulness of the architecture. The *normal* oscillator was a 4-stage differential ring oscillator, and the *FF* was a feed-forward architecture imposed on top of the 4-stage differential oscillator. The different points represent different inverter strengths. There was evidence that the feed-forward architecture was able to hit higher frequencies with lower power and smaller devices.

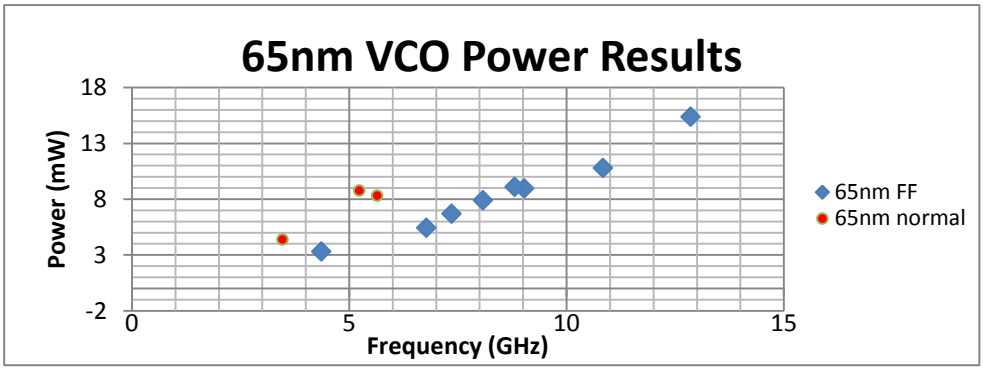

Figure 3-4: Power versus frequency results of feed-forward architecture.

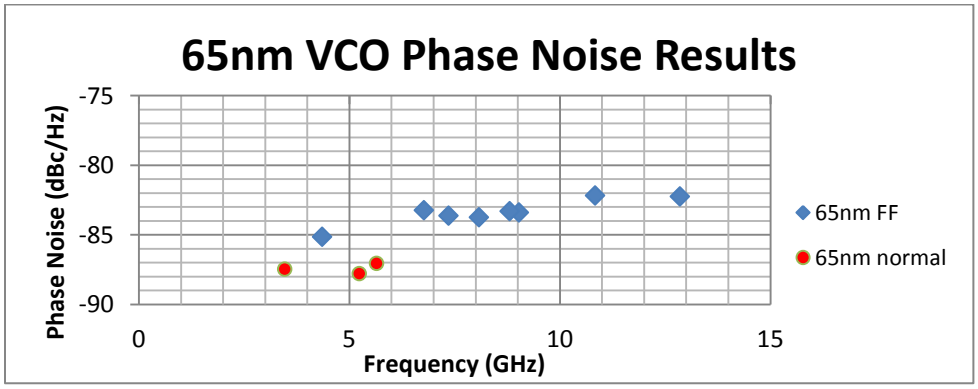

Figure 3-5: Phase noise versus frequency results of feed-forward architecture.

This design could achieve higher oscillation frequencies with lower power; however it does this with higher phase noise than a normal differential ring oscillator. The above plots enable us to clearly see the noise / power trade-off previously discussed. A figure of merit (FOM) was calculated for the above data and is shown to equal out and be approximately the same. (3-1) was used from *G. Zhou*[11].

$$FOM = 10 \cdot \log\left[\left(\frac{f_{osc}}{\Delta f}\right)^2 \cdot \frac{1}{Power}\right] - L(\Delta f) \tag{3-1}$$

Table 1 was assembled and shows the similar FOM results obtained. The drive strength rating of an inverter is a method to describe the current sourcing / sinking capabilities of the logic gate. The data in the table is logged using drive strengths of inverters found in the design kit's standard cell libraries (i.e., 0, 1, 2,etc.) and the feed-forward trials are coded using the following syntax: *forward inverter strength / cross-coupled strength / feed-forward strength (ex: 16/2/12)*. In the case of the normal ring trials, the syntax is as follows: *forward strength / cross-coupled strength (ex: 24/12).*

Table 1: FOM calculation for regular versus feed-forward VCO.

Type	Frequency (GHz)	Current (mA)	PN @ 1MHz (dBc/Hz)	FOM
Differential (24/12)	5.65	8.29	-87.06	152.91
Differential (24/16)	5.24	8.76	-87.78	152.74
FF 4-stage (16/2/12)	10.84	10.8	-82.17	152.53
FF 4-stage (16/2/8)	9.03	8.96	-83.39	152.98
FF 4-stage (16/2/6)	8.08	7.91	-83.75	152.91
FF 4-stage (24/2/16)	12.86	15.37	-82.24	152.55

3.2.3 Explored Methods for Sub-VCO Architecture

It was a starting point to build on the work of *Hajimiri, et al.* and *M. Behbahani and G. E. R. Cowan*, as the authors showed that this noise and power trade-off can be exploited [4,5]. It was the next step that was missing in order to take statically connected sub-VCOs and to implement them in a dynamic way. It was a concern from the beginning to keep all loading symmetrical to every stage, as in a multi-phase VCO each phase is used. In *M. Behbahani and G. E. R. Cowan*, one large node was created by attaching one phase from each sub-VCO to a common node through pass-transistor switches [5]. It was taken as the initial starting point of this project from which several connection schemes

were considered. In order to pick the best ideas, parameters such as loading, isolation, power dissipation and feasibility were taken into account.

3.2.4 Ring Length Selection

In CDRs with sub-rate architectures, the number of stages in the ring VCO will affect several different system parameters. In a CDR, data is received serially and, in this case, without a forwarded clock. The PLL must use the data to align itself in phase to the serial data stream. The data is then fed to multiple latches in parallel, each of which is clocked in by different phases of the VCO. This process is shown in Figure 3-6.

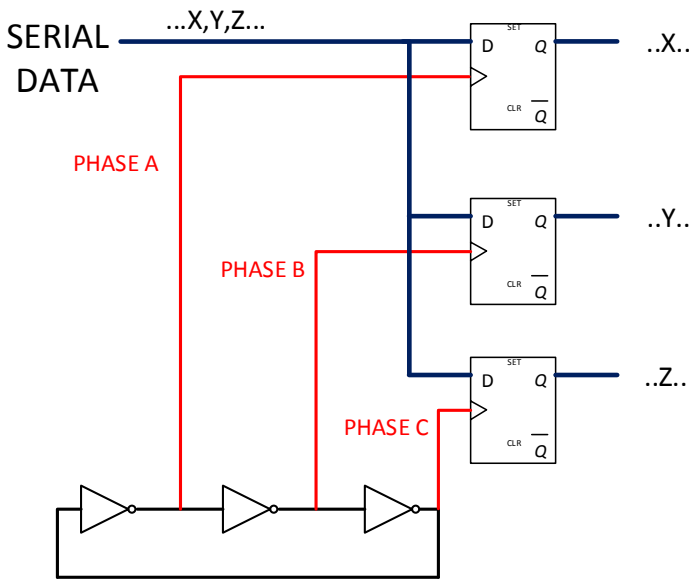

Figure 3-6: Sub-rate architecture example.

As a result, the number of latches increases as the number of phases increases. When this happens, the required frequency of the clock decreases. For this reason, an investigation was needed in order to find out the trade-off between a 4-stage ring oscillating at a frequency of 6.25 GHz versus a longer 8-stage ring oscillating at 3.125 GHz.

Phase noise was a concern between the two categories. Since the 8-stage ring would be running at half of the frequency compared to the 4-stage ring, it was important to see the noise trade-off. The 8-stage ring runs at a lower frequency than its 4-stage counterpart, and so the loop bandwidth will also be lower. This indicates that by changing the cut-off frequency to half of the original frequency, while keeping the same -20 dB/decade slope, it will result in a phase noise improvement of:

$$20 \cdot \log\left(\frac{f_2}{f_1}\right) = 20 \cdot \log(2) = 6dB$$

The integrated jitter was then found and the two scenarios were compared. It was observed that in order for both the integrated jitters to be equal, the 8-stage ring would need to have an additional 3dB of phase noise improvement. This means that the 8-stage ring would need to exhibit an improvement in phase noise performance of9 dB in order for the jitter to be equal for both cases.

It was decided to go with the longer 8-stage ring as it was believed to have less accumulated jitter effects, as well as it required a lower reference frequency which would simplify the testing process of the implemented chip. The 8-stage ring would also enable a $1/8^{th}$ sub-rate architecture that would allow for lower data rates.

3.3 Presented Method

After considering the information and results presented in the previous sections, the following design shown in Figure 3-7was chosen as the architecture of a sub-VCO.

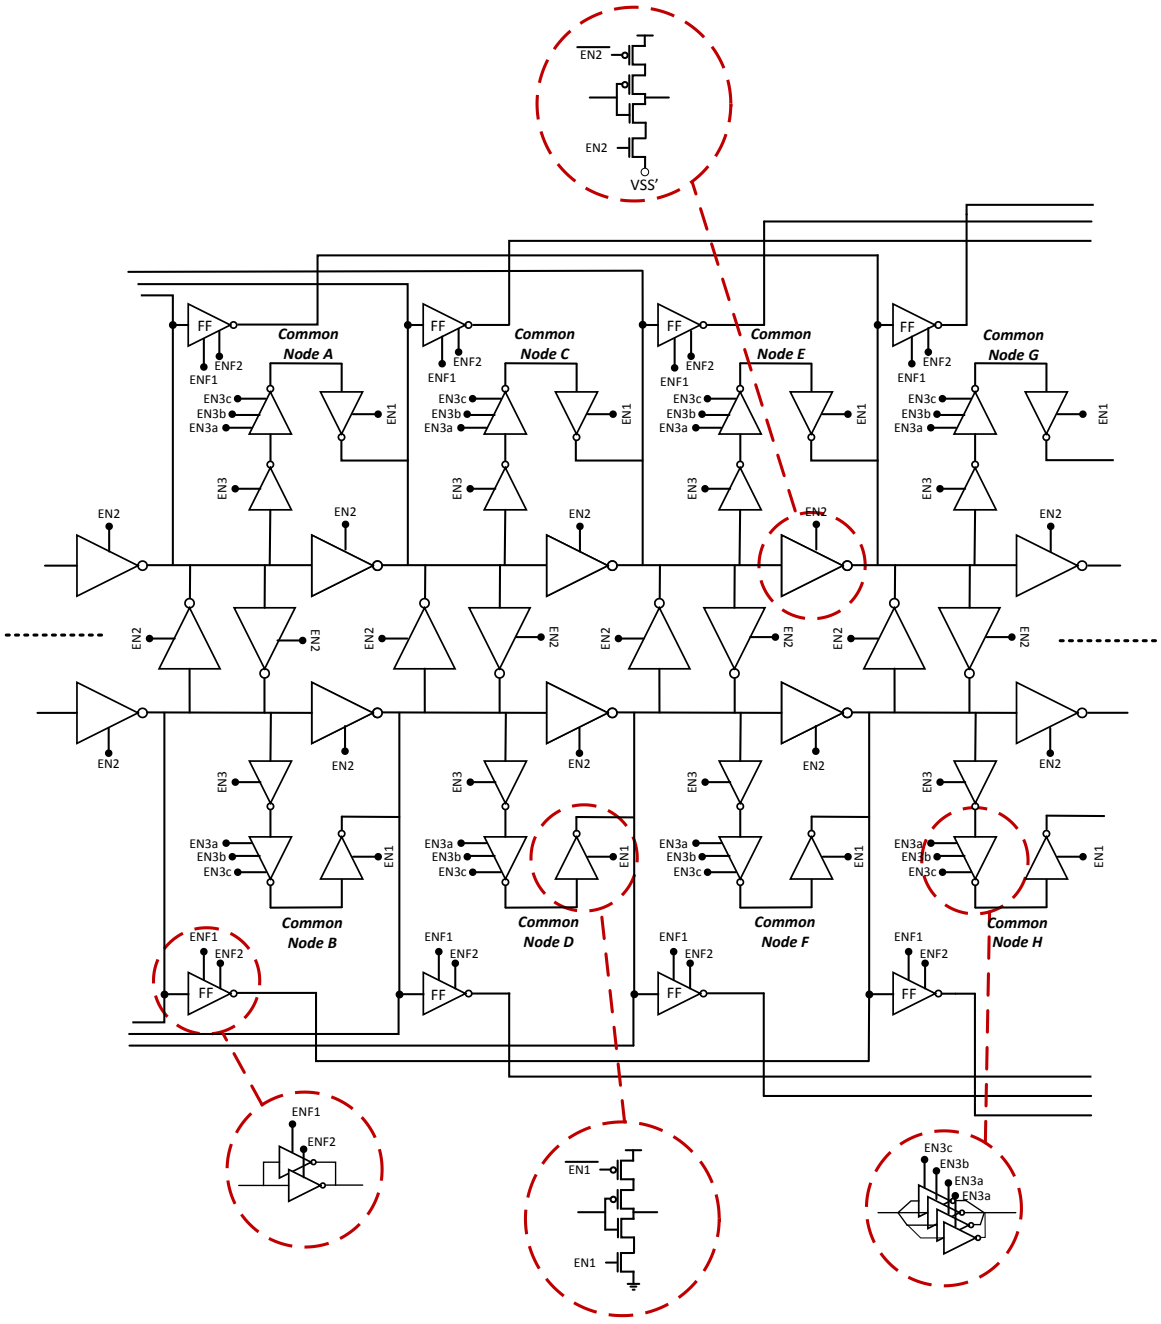

Figure 3-7: Overall schematic of proposed sub-VCO architecture.

The sub-VCO comprises an 8-stage ring, 4 stages of which are shown above, with an auxiliary path consisting of the buffered output of each phase of the ring. The inverters labelled "FF" make up the feed-forward loop that can be used to increase the speed of the VCO. The feed-forward loops are made up of two standard cell inverters, strength 0 and 1, and are able to be independently enabled. This gives control to increase or decrease the frequency of the ring post-fabrication. Each of the sixteen auxiliary paths forms a common node by connecting to the corresponding node in each of the other sub-VCOs. This will be discussed in detail in the following sections. The inverters in the main ring as well as the cross-coupled are current starved and thus have a pseudo ground, labelled *VSS'*. The other devices were not current starved in the hopes of keeping their delay and drive strength constant. In this work, the overall VCO is comprised of three current starved sub-VCOs. Each of the like phases is connected together, and so the common node has the effect of one VCO with transistors three times as big, as described in *M. Behbahani and G. E. R. Cowan*[5]. The problem of phase excursion in [5] is believed to occur when VCOs were attached while their phases were not aligned. To resolve this, a procedure was designed in which the rings were brought up to speed and aligned before the connection is made, thus reducing the chance of the output showing a significant error. This design relies on feedback as well as feed forward paths in order to properly work.

3.3.1 Activation Procedure of the Sub-VCOs

The developed procedure for bringing a sub-VCO up to speed with the goal of establishing connection with minimum excursion is described here in detail. This is necessary as without this method, unaligned phases may be paired together causing a

glitch in the output which may cause undesired system level effects or errors in the data clocking.

In Figure 3-8, inverters in the main path are the primary reason for oscillation. Cross-coupled inverters ensure the outputs of each stage are 180 degrees out of phase. The auxiliary path, created from inverters labelled *EN1* and *EN3*, is the mechanism that is intended to allow the connection of all the sub-VCOs together dynamically. The *ENx* pins are for the activation procedure, which allow for minimum phase excursions during VCO connection.

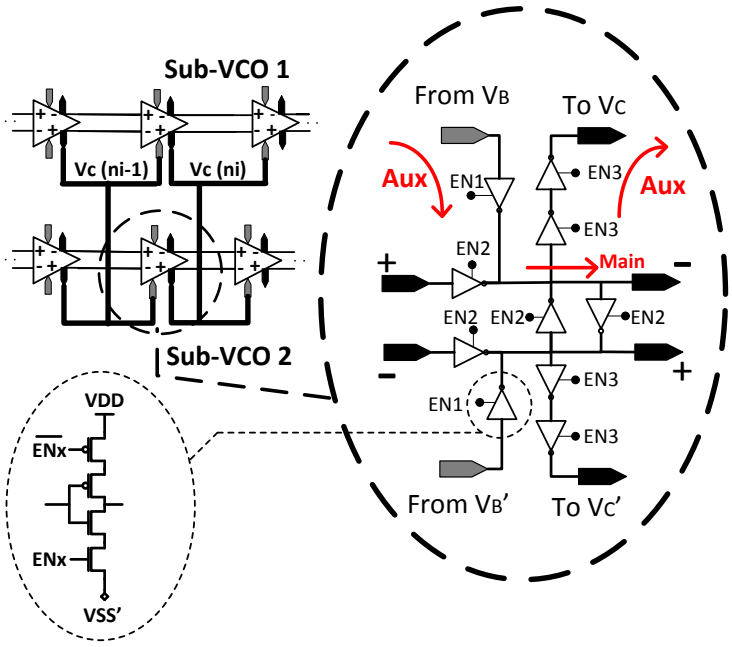

Figure 3-8: Power on sequence of sub-VCO.

Consider the following example where only one sub-VCO is enabled and the system is to operate at a higher data rate. This requires the jitter requirements to tighten up and thus another sub-VCO needs to be activated. At this point, each of the common nodes are oscillating, due only to the already active sub-VCO. The other sub-VCOs have no

interaction with this node at the moment as they are isolated through the buffer EN3. The first step of the procedure is to enable *EN1*, which will allow current to flow though the CMOS inverters. These first sets of inverters are for feedback alignment. As the name suggests, the enabled inverters allow a small amount of current to start charging up the nodes of each phase in the ring. When *EN2* is activated, the ring is started and the pre-charged nodes from step 1 are pulled apart due to the forward and cross-coupled devices. Shortly after, the ring should be oscillating closely to the already activated VCO, as the feedback inverters (from *EN1*) are providing an injection lock to the already enabled sub-VCO signal on each of the 16 common nodes. The final step is to activate *EN3*, which will connect the newly activated sub-VCO through inverters to other sub-VCOs already active and connected. At this point, the jitter will have improved and higher data rate functionality can commence.

3.4 Phase Excursion Compensation Methods

Consider the case when multiple sub-VCOs are connected together. In this scenario, two sub-VCOs share a node, shown previously as the common phase node. With the same node capacitance, but now with two sub-VCOs contributing current instead of one, the voltage at the common phase node will raise faster according to

$$i = C\frac{dV}{dt} \tag{3-2}$$

If the voltage rises faster, then the inverters in this path will experience smaller delay, as illustrated in Figure 3-9. If the path has smaller delay then, according to (2-7), the oscillation frequency will increase.

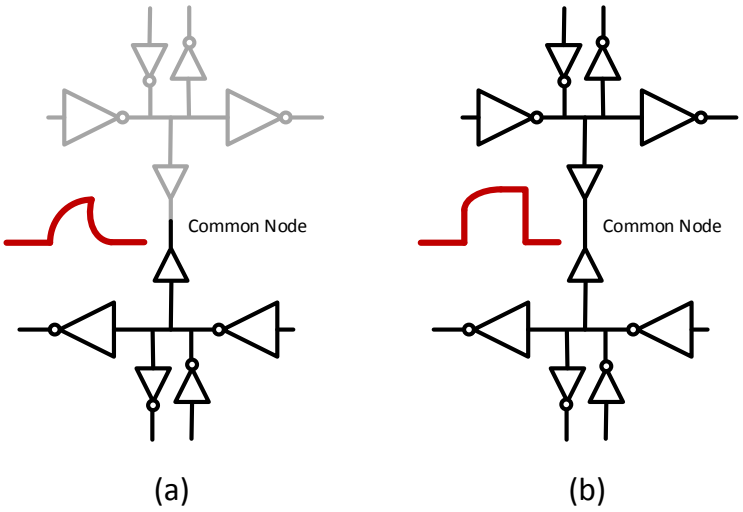

Figure 3-9: Common node rise time (a) with 1 sub-VCO active (b) with 2 sub-VCOs active.

This is the essence of the problem that was faced. If the delay of the auxiliary path reduces, then the ring will behave as it did in the feed forward ring structure, i.e., it will speed up. This may have a problematic effect on the loop as it will cause a large phase excursion, or in the worst case scenario, may cause the PLL to lose lock if the frequency change is too abrupt. Therefore, two techniques have been developed in order to counteract this problem, each of which is controlled by the presented delay path. As the delay path has tuning mechanisms, the activation of each compensation method will rely on experimental results that will require the user to fine tune these parameters. If needed for adequate compensation, this procedure shall be done on a per die basis, if the required measurement can carried out quickly enough,.

3.4.1 Variable Capacitor Technique

As an increased number of inverters start injecting current into the common node, the delay of the auxiliary path will decrease. This was previously found assuming that the

current into the node is increasing and the node capacitance is constant, then to satisfy(3-2), the rate of voltage change at the node increases. If however, one were to consider the case where the capacitance at the common node was not constant but variable, then the delay could be dynamically re-adjusted. This is one of the implemented solutions to the issues brought about by dynamically enabling the sub-VCOs.

A method using NMOS transistors as variable capacitors was used. Consider an NMOS transistor, where the drain and source of a MOSFET are connected together. With the body grounded, the gate will be considered as one terminal of a capacitor. If the drain and source are connected together at a high potential, then the gate capacitance of the MOSFET will have a capacitance equal to depletion region capacitance in series with the oxide capacitance. If the drain-source terminals are brought to a low potential, this removes the depletion region capacitance and results in a higher capacitance.

In this way, the drain-source connection can be either attached to VDD for a small capacitance value, or connected to ground in order to have a larger capacitance value. Table 2 shows the effect of the width, length and number of fingers on the low and high capacitance values. The capacitance values were found by applying a step response to the input of a resistor connected to the gate of a NMOS transistor. The corresponding time constant, the time at which the capacitor charges to 63% of the total applied voltage, was recorded and the capacitance values were extracted from these results.

Table 2:Results of changing variable capacitance parameters.

Width	Length	# of Fingers	Drain-Source @ GND	Drain-Source @ VDD
500nm	1um	5	39.52 fF	10 fF
		15	107 fF	19 fF
500nm	500nm	5	22.7 fF	8.5 fF
		10	40.2 fF	10.5 fF
		15	57 fF	12 fF

The implemented circuit is shown in Figure 3-10, where two variable capacitors are shown. These variable capacitors are implemented at each of the 16 common phase nodes.

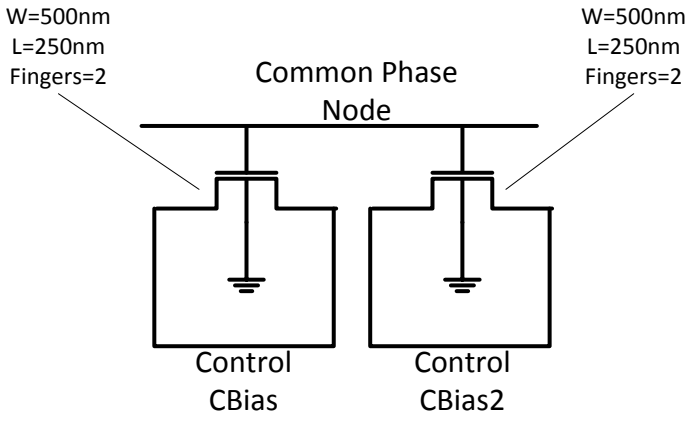

Figure 3-10: Variable capacitance schematic.

Each one of the two capacitors per node has a separate activation signal, and gets activated when the corresponding sub-VCO is enabled. If sub-VCO #2 is activated, then only one of the variable capacitors per common phase node should be enabled to support the sudden change. If sub-VCO #3 is also activated (and thus a total of 3 sub-VCOs

active), then both of the variable capacitors will be enabled. In this way, each extra unit of capacitance is enabled depending on how much of a sudden change to the node is expected.

3.4.2 Variable Buffer Strength Technique

In addition to the variable capacitance added to each node, the strength of the inverters in the buffer providing the current to the common node was made variable. For this to correctly function, the total amount of drive strength applied to each common node always has to be the same. Figure 3-11 (a) below shows the buffer of a sub-VCO connected to the common node delivering a total drive strength of three(strength 0 is half of the strength of 1) in a scenario where no other sub-VCOs are active.

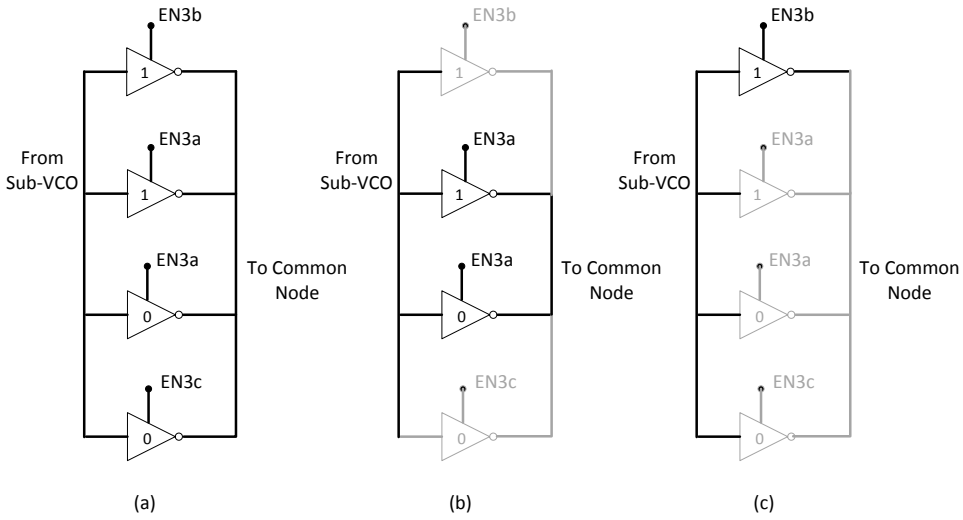

Figure 3-11: Configuration of a variable strength buffer connected to a sub-VCO when (a) 1 sub-VCO is active, (b) 2 sub-VCOs are active and (c) 3 sub-VCOs are active.

In turn, (b) shows the configuration of the same buffer when two sub-VCOs are active. To keep the same drive strength of 3, each buffer would have to drive the common

node with a strength of 1.5. If three sub-VCOs were active as in (c), the strength of each buffer driving the common node would have to be 1. In order to achieve these combinations, four inverters per buffer were needed.

3.4.3 Activation Timing

The process of activating the different stages of this design comes down to proper timing. If the ring is not given enough time to start oscillating and become in phase with the other active sub-VCOs, then the technique will not be useful. On the other hand, if too much time elapses, then the circuit will be of no use for the desired application of on-the-fly data rate switching. Some of the trials are shown in Table 3below, which demonstrate the effect of different activation timings. All times were made relative to the first column, EN1, which is the first step in activating a sub-VCO.

Table 3: Activation procedure time trials.

EN1	EN2	EN3	Excursion
0 ps	100 ps	300 ps	12 ps
0 ps	100 ps	500 ps	7.2 ps
0 ps	50 ps	600 ps	7 ps
0 ps	300 ps	700 ps	6.2 ps

It was found that the activation of EN1, EN2 and EN3 should occur at 0, 300 ps and 700 ps, respectively, relative to EN1. The effect of CBias signal, the signal used to turn on the variable capacitors, was then studied. It was found that best results occur if CBias is activated 200 ps *before* EN3. In order to achieve activation signals produced at the

optimal time intervals, a variable delay path was used. This delay path consists of a row of CMOS inverters, capacitors and multiplexers, and is shown in Figure 3-12.

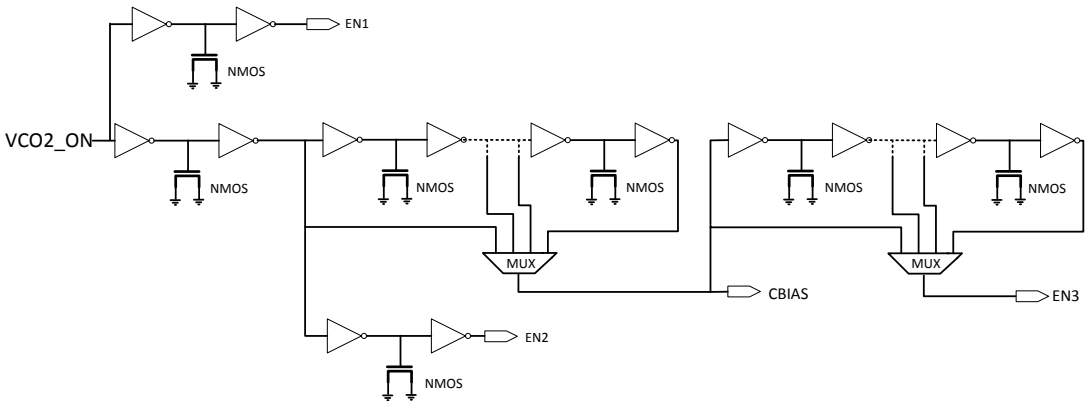

Figure 3-12: Delay path for implementing power-on procedure of sub-VCO.

This is a simplified schematic of the actual delay path; the actual one contains more inverters in order to achieve the needed time intervals of 300 ps for EN2, 500 ps for CBias and 700 ps for EN3. Two of these circuits exist, one for sub-VCO #2 and one for sub-VCO #3. As the following inverter stage will not be triggered until the previous stage surpasses a threshold, this charge time can be thought of as a delay time. The multiplexers allow for the selection of input that will propagate to the output of the block. This is a used as a post-fabrication tuning method of the delay paths. As shown, depending on which point is allowed to propagate through will determine the activation signal that causes triggering. The delay path is activated upon the user giving logic high to *VCOx_ON*, indicating that they wish a sub-VCO to become active.

In order for the variable buffers to be activated properly, another logic block had to be designed to accept the control signals from the delay path, as well as inputs from the user. This allowed for the activations or deactivations of the variable buffers to keep constant

buffer strength to the common node, as shown in Figure 3-11. Figure 3-13 is the circuit diagram for this operation.

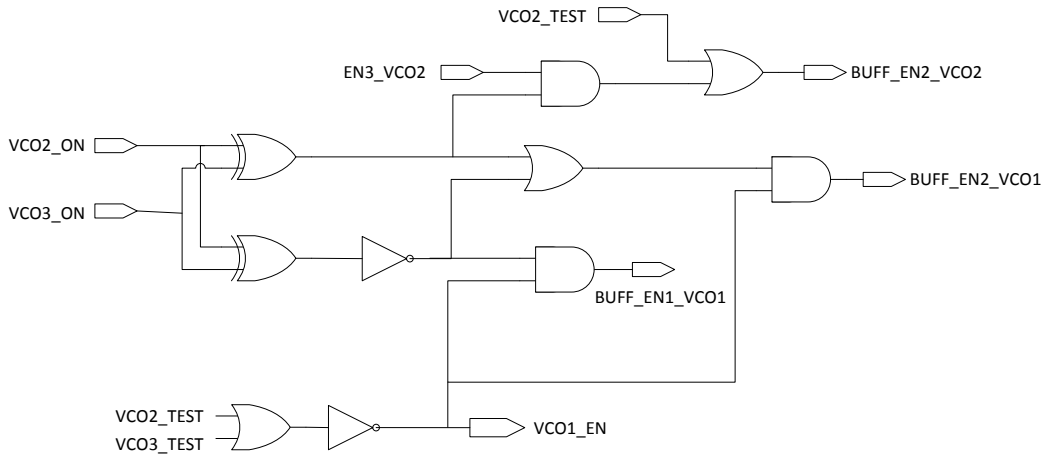

Figure 3-13: Logic for controlling sub-VCO bank.

Under normal circumstances, sub-VCO #1 would always be active. This block took into account the need for each sub-VCO to be activated independently, in order to tune and apply compensation methods post-fabrication to sub-VCOs #2 and #3. This was done using the *VCO2_TEST* and *VCO3_TEST* pins.

3.5 Phase-Locked Loop (PLL) Design

3.5.1 CML XOR Phase Detector

In *Toifl, et al.*, current-mode logic, or CML, was used to implement a phase detector [10].Figure 3-14 illustrates this design.

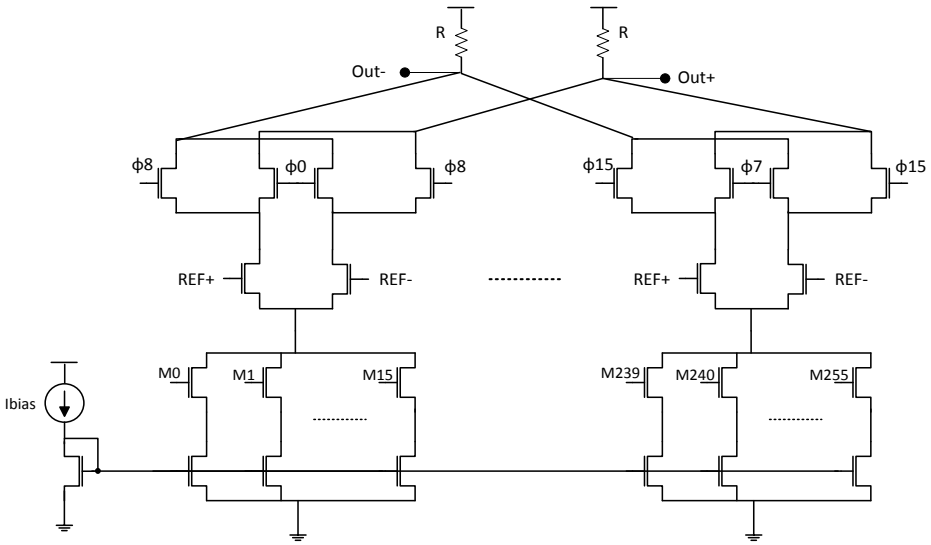

Figure 3-14: Schematic of CML phase detector.

Figure 3-14 shows a CML phase detector based on a16-way XOR design. Each of the XOR gates has the reference clock applied to one input. Each of the PLL's 16 differential output phases is applied to one XOR gate. Bias current to each XOR gate is then controlled by thermometer encoded current sources. The current will always be flowing, but the paths can change. In this respect, if one path has all transistors actively conducting current, then there will be a voltage drop across the resistor associated with this active path. The voltage drop can be easily expressed by Ohm's law, where the voltage potential across the resistor is equal to the multiplication of the resistance value by the current through that resistance. This Gilbert Cell multiplier circuit is interesting as there are two

output pulses per input cycle. This indicates that phase information is generated at a rate of twice the input frequency.

Transistors M0 - M255 provide a method of interpolating between phases. At any one time, 16 of the transistor current mirrors can be active. This can be 16 transistors conducting current through one Gilbert Cell or two neighboring cells, and thus enabling interpolation. In order to keep power low, a bias current of 320uA was chosen. This value was desirable due to the bias current circuit that is responsible for interpolation. With sixteen interpolation points, and thus sixteen current branches, this results in only 20uA required to flow through each branch for a total of 320uA. This small amount of current per branch keeps the size of the biasing transistors small.

3.5.2 Differential to Single-Ended Transconductance

The voltage-to-current conversion is a necessary step in the loop. As the output of the phase detectors is a voltage, and the loop filter works with current, an intermediary step is required. The design used was similar to that of the transconductor in *Toifl, et al*.[10].

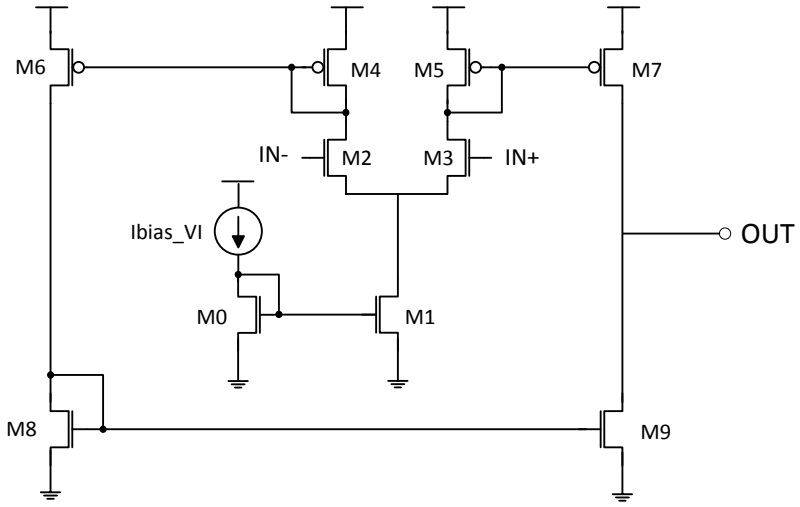

Figure 3-15: Schematic of transconductor.

According to *Toifl, et al*., the combination of this transconductor and phase detector gives gain

$$K_{PD} = \frac{2 \cdot I_{CP}}{2\pi} \tag{3-3}$$

where I_{CP} is the bias current of the transconductor [10]. The design employs a differential pair that will fully switch (ideally) and steer current into one of two PMOS diode-connected loads. This will then mirror the current out to either transistor M6 or M7. If the current is mirrored to M7, then current is sourced to the following stage, the loop filter. If the current is mirrored to M6, then the diode connected device M8 will act as another mirror and will result in M9 sinking current from the loop filter connected to the output. As the frequency of switching will be twice the reference frequency, due to the characteristics of the XOR phase detector, the gate capacitance will start to have an effect on the switching signal and thus may not allow devices M2 and M3 to switch fully.

3.5.3 Loop Filter

The loop filter is the next block to be designed. In order to simplify calculations, the parallel capacitor C2 will be ignored for the moment, to keep the PLL as a second-order system.

3.5.3.1 Finding VCO Gain

The VCO gain, or K_{vco}, is a variable needed for proper calculation of the loop filter. As supply voltage decreases with newer technologies, the VCO gain must increase in order to keep the same tuning range. This results in larger noise due to noise on the high gain node. In order to find the gain, the VCO was taken out of the loop and the control voltage was

swept. The output frequency was then examined to find the slope of the line. Figure 3-16 shows the tuning curve of one sub-VCO. As can be seen, the VCO gain does change as the voltage applied to the current starving devices increases. It was also found that as more sub-VCOs are added, the tuning curve will change slightly. In this case, the average of the sub-VCO gains was used. The region around the gate voltage of 0.5 V was used, in order to have both the NMOS and PMOS working optimally. This would put the VCO gain at around 3.55 GHz/V.

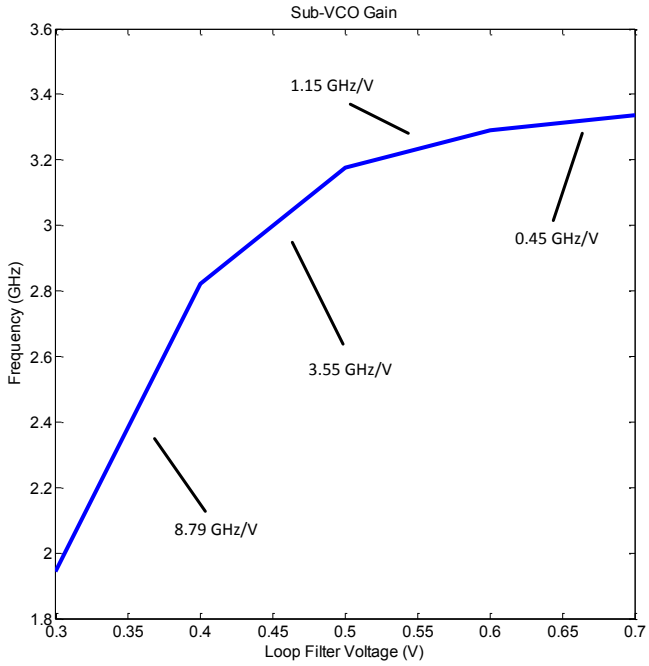

Figure 3-16: VCO gain (K_{vco}).

3.5.3.2 Determining Loop Filter Components

The necessary loop filter components will now be found. Starting from a standard second order transfer function:

$$H(s) = \frac{2s\xi\omega_n + \omega_n^2}{s^2 + 2s\xi\omega_n + \omega_n^2} \tag{3-4}$$

Applying feedback theory, the following transfer function of the second-order PLL is found to be:

$$H_{CL}(s) = \frac{A}{1 + A\beta} = \frac{sRK + K/C}{s^2 + sRK + K/C} \tag{3-5}$$

$$\frac{2s\xi\omega_n + \omega_n^2}{s^2 + 2s\xi\omega_n + \omega_n^2} = \frac{sRK + K/C}{s^2 + sRK + K/C} \tag{3-6}$$

where K = K_{PD} * K_{VCO}, *R* is the loop filter resistor and *C* is the loop filter series capacitor. If (2-3) and (3-3) are substituted, the following equations are obtained:

$$\omega_n = \sqrt{K/C} = \sqrt{\frac{K_{PD} \cdot K_{VCO}}{C}} = \sqrt{\frac{2 \cdot 2 \cdot I_{CP} \cdot K_{VCO}}{2\pi C}} \tag{3-7}$$

$$\xi = \frac{RK}{2 \cdot \omega_n} = \frac{R}{2}\sqrt{\frac{2 \cdot 2 \cdot I_{CP} \cdot K_{VCO} \cdot C}{2\pi}} \tag{3-8}$$

The variables found are now able to be placed into the equations and the loop filter components can be found. The work in *Toifl, et al.* found that the values of $\zeta = 1.6$ and $\omega_n/\omega_o = 0.04$ were optimal for their implementation, and they were used here as a starting point [10]. These equations provide values for the series resistor and capacitor. As a rule of thumb given in *Toifl, et al.*, the parallel capacitor C2, shown in Figure 2-9, should be less than $1/10^{th}$ the series capacitance for stability and greater than $1/50^{th}$ of the series

capacitance in order to keep deterministic jitter low [10]. The value of C2 was chosen to be approximately 1/40th of the series capacitor, C1.

3.5.4 VCO Starving and Compensation

The chosen method for frequency tuning involves current starving using NMOS transistors, illustrated in Figure 3-17. Each VCO is controlled by a starving device with the same control voltage. M0 controls the current I_A that is available for sub-VCO #1, M1 controls the current I_B available for sub-VCO #2, and so on. This was done in order to have all starving currents similar in each sub-VCO, with the goal of having free-running frequencies of the oscillators close together. In order to choose the size of these starving devices, it was first observed how much current is needed for the VCOs to oscillate at its maximum frequency. After, iterative simulations were done in order to find a transistor width able to sink the needed current, while taking into account the drain-source voltage with each change. The lengths of the devices were then also increased in order to increase the output resistance, r_o. By having high output resistance, the drain currents can become less sensitive to the drain-source voltages. These current-starving devices are operated at the edge of the saturation region.

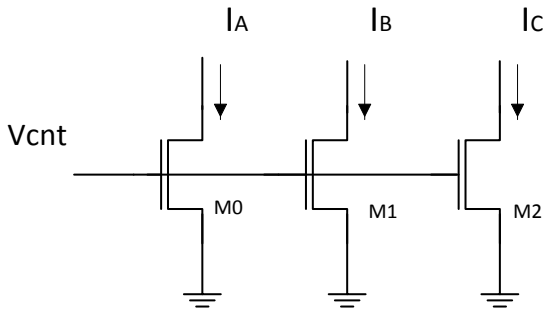

Figure 3-17: Current-starving method for sub-VCOs.

Compensation methods were also implemented in order to adjust oscillation frequencies post-fabrication to account for mismatch between sub-VCOs and variations due to global variation. Below in Figure 3-18 is an illustration of a control method for each oscillator. Transistors M5 through M8 provide fine tuning and activate a current flow through transistors M1 through M4. These are binary weighted transistors that have their gates controlled by the same control voltage as M0.In this way, the effective width of M0 is able to be increased or reduced depending on which combination of M5 – M8 switches are active.

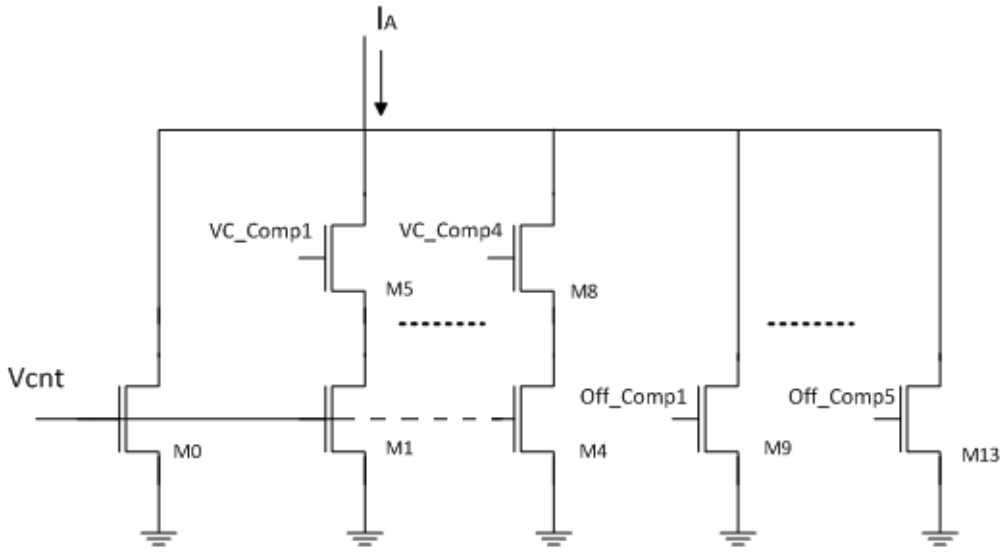

Figure 3-18: Frequency compensation method for sub-VCOs.

Coarse tuning is done using binary weighted transistors M9 through M13, where M0 is one of the current starving devices shown in Figure 3-17. These will increase the total current available for the ring oscillator to use, and thus increase the frequency of oscillation. However, as more of these coarse tuning devices are activated, we approach

the maximum oscillation frequency and thus the current starved devices have less of an effect on the overall current. Enabling M1 to M4 increases f_o, the frequency of oscillation, and contributes to K_{vco}, whereas M9 to M13 only increase f_o. It is expected that after fabrication, the optimum settings for M1-M4 and M9-M13 will be found experimentally.

3.5.5 Output Buffers

The output buffers take the signal from each of the common phase nodes and buffer them out to the phase detector. The signal that is passed to the phase detector is then further buffered out to be passed off to MUXs and power buffers to be sent off-chip. This secondary set of buffers is digitally controllable and can be deactivated in order to save power when at a lower data rate. The output buffers contain the implemented variable capacitors. Figure 3-19 shows the schematic diagram for one phase of the output buffers, which is the same for every other phase.

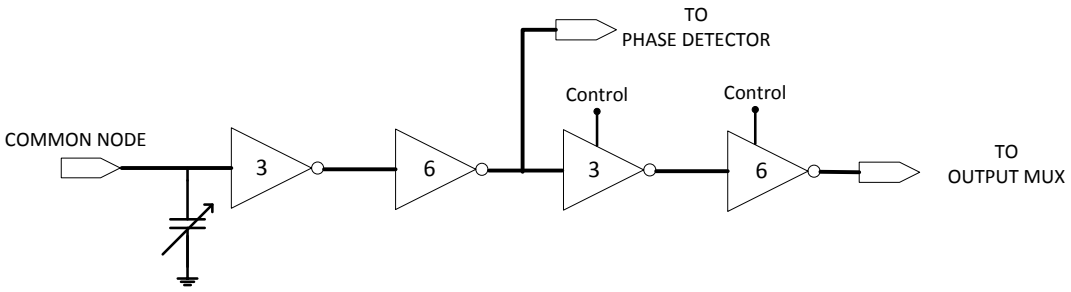

Figure 3-19 Phase buffer schematic

The numbers inside the inverters correspond to the inverter drive strength, as listed in the standard cell library. The inverters that are controllable drive a MUX that selects 1 of the 16 phases to be sent off-chip. This is done to save I/O pins.

3.5.6 Variable Loop Filter Resistance

In order to provide more freedom during the testing phase of the chip, a variable loop filter resistance was implemented. This is shown in Figure 3-20, where the transistors act as switches and bypass the unwanted resistors. The series resistance of the switches was found to be approximately 60 Ω, and this was taken into account when selecting the resistor values.

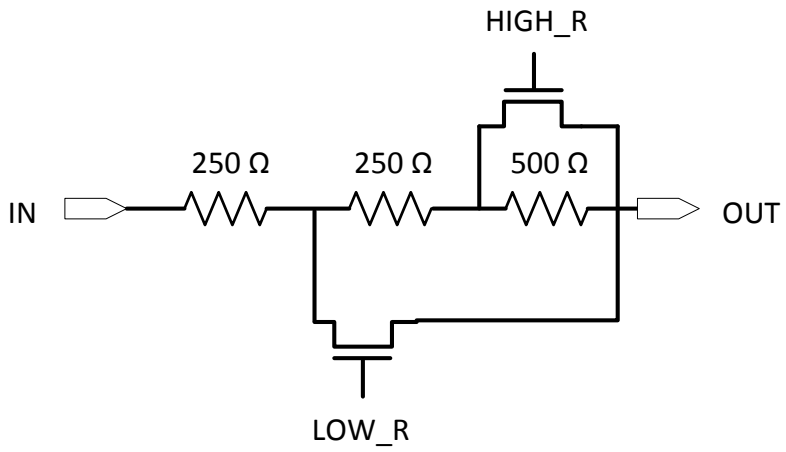

Figure 3-20: Variable loop filter resistance.

3.5.7 State Registers

In order to setup the system as well as tune parameters, a state register was required. This accepts a serial input and loads it into 64 flip-flops using an external clock. Once the data is entered, the data is latched into a secondary set of registers that provide logic signals to different blocks in the system and set circuit parameters. The secondary set of registers was used in order to prevent states from changing until all state bits were entered. Figure 3-21 shows a simplified circuit of this block.

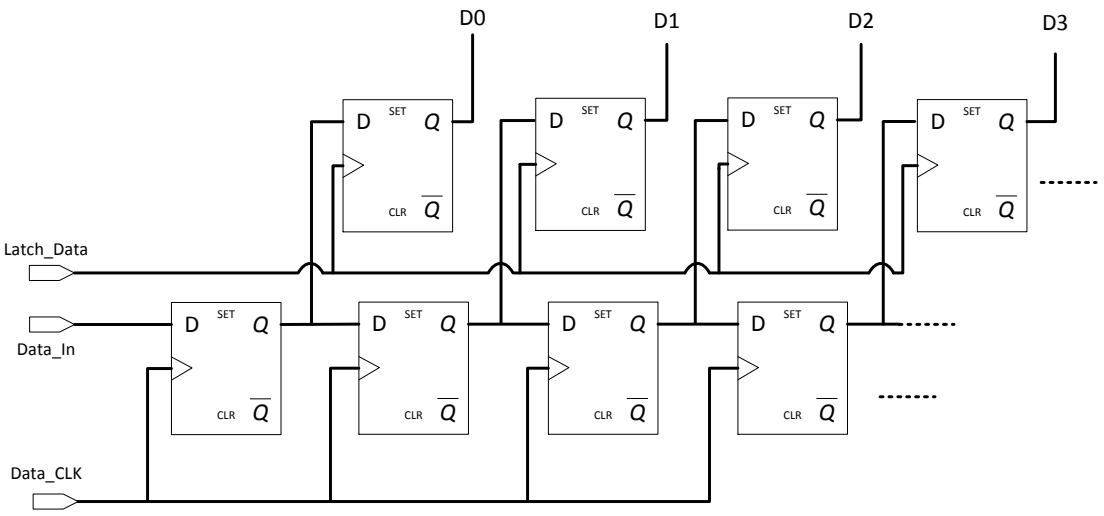

Figure 3-21: State bit registers.

The bits *D0 – D3* are state bits that are latched in when the *Latch_Data* signal goes high. This controls transistor gates that in turn reconfigure circuits and parameters.

3.5.8 Creating a Differential Input

Differential circuits provide common-mode rejection, and thus are appealing for analog circuit designers. However, during testing of a chip a differential clock is not always available. For this reason, a single-to-differential converter was placed on-chip. This block was made of standard cell CMOS inverters to avoid the need to use analog differential pairs with larger passive filtering components. Figure 3-22 below shows the schematic of the implemented circuit.

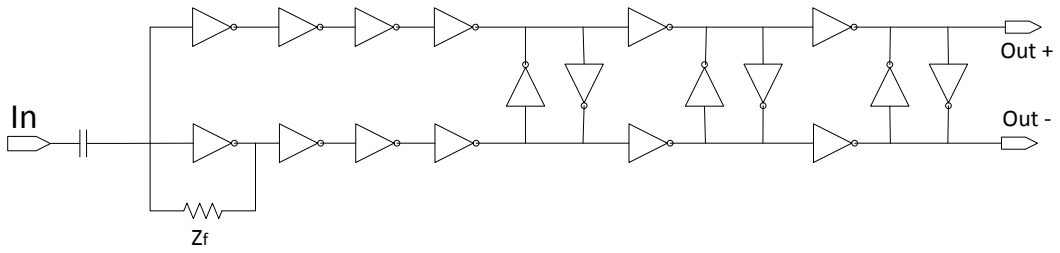

Figure 3-22: Single-ended to differential clock generator.

The capacitor at the input of the chain represents an off chip capacitor for AC coupling, possibly a Bias-T connector, due to the signal generators having a zero DC requirement at the output port. This results in the circuit requiring self-biasing.

The feedback resistance, *R*, sets Vin and Vout. The input resistance can be estimated using Miller's theorem.

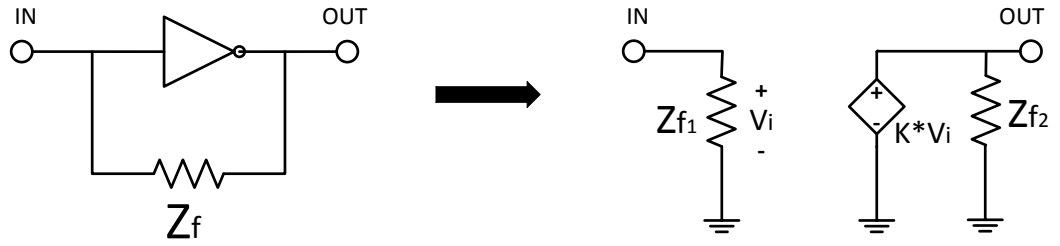

Figure 3-23: Miller effect on input stage [9].

Where Z_f is the feedback resistance, Z1 and Z2 are the Miller equivalent components, and K is the gain of the amplifier block. By applying a small signal to the inverter (without feedback) DC biased at VDD/2 and observing the output, it was found that the gain of the inverter is -3. It is of importance to match the impedance that the source will see when the chip is connected to the generator. This is to prevent reflections due to a load that is

unable to absorb the signal energy. The output impedance of the signal generator is 50 Ohms, and thus Z_1should equal to 50 Ohms for proper matching. The feedback resistance is then found using (3-9).

$$Zf_1 = 50 = \frac{Z_f}{1-K} = \frac{Z_f}{1-(-3)} \tag{3-9}$$

$$Z_f = 50 \times 4 = 200\Omega$$

3.5.9 CDR Secondary Loop

With the chosen design of the VCO, multiple phases are available to the user. As the phase detector allows locking of the loop to one phase, this opens the door to a secondary control loop. This allows different phases, or a combination of phases, to be the fed to the phase detector. First, let us explore the reason why one may need this functionality.

3.5.9.1 CDR Ideal Sampling Location

The ideal sampling location of incoming data can be deduced from an eye-diagram. As previously explained, the eye-diagram gives a large range of possible outcome characteristics depending on previous bits. With this information, it may be found that the ideal sampling location is not located in the geometric center of the eye. With an eight-stage differential ring, 16 phases are available for use. Out of these 16 phases, eight will be used for data sampling and eight will be used for edge sampling, as shown in Figure 2-5. This enables for the sampling location to change, by varying which phase of the VCO us used to phase lock to the reference.

3.5.9.2 Using Secondary Loop

Data is ideally sampled in the center of the eye. In order to do this, a secondary control loop, one that interacts with the PLL, will be used. This loop will dynamically cycle through the phases and will lock to the incoming reference signal. The design of the secondary loop is based on that of *Toifl, et al.*, where not only each individual phase can be used to lock to the reference, but the interpolation of any two adjacent phases as well [10]. This secondary loop is designed to work at a much slower rate than the primary loop.

Through the use of a current-mirror structure, the weight each phase has on the interpolation can be varied between 0 and 16 units of strength (measured in current). Only 16 units will be active at any one time. Using a shift register, it is possible to shift through the phases in increments of one unit of strength. Using this method, it allows all 16 units to contribute to a single phase, or enable an interpolation of two adjacent phases to occur where the interpolation weight is dependent on the active units given to each phase .The shift register also allows moving up or down in phase. Using 16 units of strength per phase, with 8 phases, allows for increments as fine as three degrees of the total 360 degrees per step. Figure 3-24 is a simplified schematic, showing only 2 D-flip flops out of the actual 256.

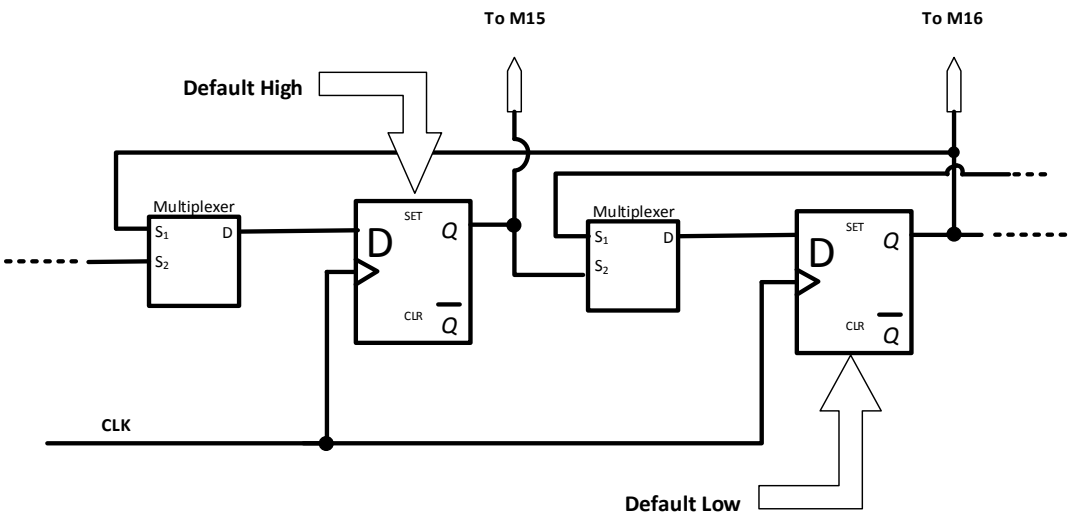

Figure 3-24 Simplified design of barrel register

The shift operation is accomplished using D-flip flops and multiplexers. The direction can be controlled by selecting the corresponding multiplexer input. The SET and CLEAR functions of the Dflip-flops are also used, as they serve as a method to reset the shift register to a known state. As a default, the first 16 registers are set, which enables M0 - M15 in Figure 3-14, and allows the ring to lock to phase 1. The other flip-flops controlling M16 – M255 are cleared. After this event, external control signals can move up or down amongst M0 – M255,shown in Figure 3-14, in order to lock to other phases. In order to move up or down, the multiplexer bit is set or cleared allowing it to propagate the output to the left (shift right), or to the output to the right (shift left).Upon a reset signal, the default state will be entered, where M0 – M15 are high and M16 - M255 are low. This enables for easier troubleshooting if questionable behavior is observed during testing.

3.6 Laboratory Testing

The sub-VCOs and the supporting PLL were laid out using TSCM 65nm technology, and were made to be part of a clock generation circuit for the receiver. The chip was designed to have multiple supporting blocks in order to cover a wide range of possible experiments involving the clock generation circuits. The design of the clock generation circuit that was implemented on the chip is presented in Figure 3-25. The reference clock is fed to two projects, the PLL and the injection-locked oscillator (ILO). The ILO, designed and implemented by Michael Di Perna, also supervised by Dr. Glenn Cowan, in a separate project that will be tested as well to find the merits of the design as compared to the PLL.

The 16-phases from the PLL and ILO are sent to a MUX that will give the ability to select which of the projects to test. There is then a MUX to select which of the 16 phases will be sent off-chip, through a large buffer that can drive the loads associated with both going off-chip and impedance matching. In parallel to the phase selection, the 16 phases will also control latches. This is considered a secondary experiment to test the bit-error rate (BER) in a lab setting, using a serial input data stream. After latching data, any 3 of the 16 samples, corresponding to a data, edge and adjacent data sample, can be sent off chip. This is envisioned as being part of a secondary loop architecture that will enable the use of an Alexander-phase detector in a larger CDR project.

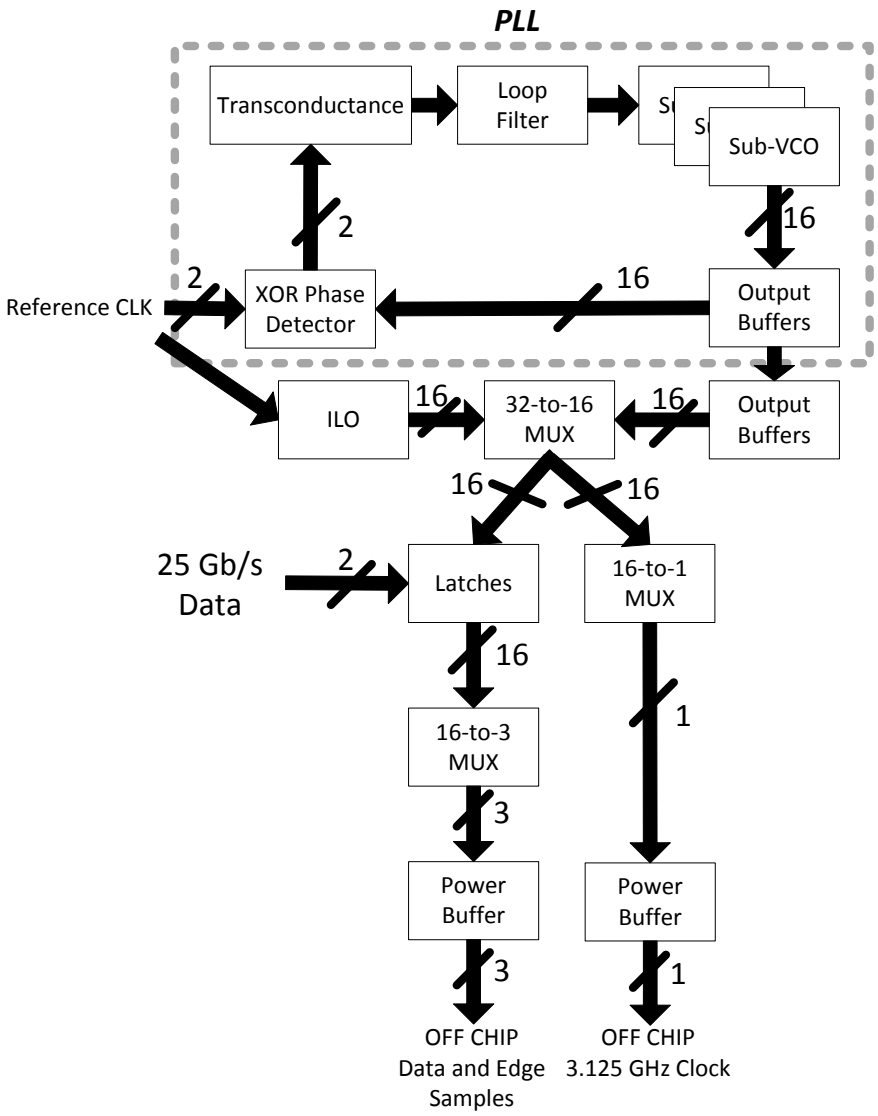

Figure 3-25: Block diagram of implemented test chip.

CHAPTER 4

Project Results

Simulation results will now be presented of the proposed sub-VCO design and PLL. As the fabricated chip will only return after the final draft of this thesis, testing results will not be available at the time of submission. The design was submitted for fabrication on the 22nd of July, 2013, for a scheduled fabrication run. It is expected to return January 2014.

4.1 Simulation of Blocks

4.1.1 Differential Clock Generator

The first block tested was the clock source of the system. The block uses an off-chip single-ended input and converts it to a differential output. The oscillating frequency of the input source is 3.125 GHz. In order to show proper functionality, Figure 4-1 below illustrates the input reference generator and the resulting differential output clock. This differential signal is then fed to the differential XOR phase detector.

In an ideal signal, the waveform would be high 50% of the time and low the other 50%. According to the data in Figure 4-2, one output has a high time of 150 ps, while the other has high time of 167 ps. Knowing that the reference clock is oscillating at 3.125 GHz, with a 50% duty cycle and a period of 320 ps, it can be concluded that the duty cycles are slightly skewed. Upon comparing the results of the system using (i) the designed block and (ii) ideal differential signals, it has been found that this slight variation in duty cycle does not affect the system performance.

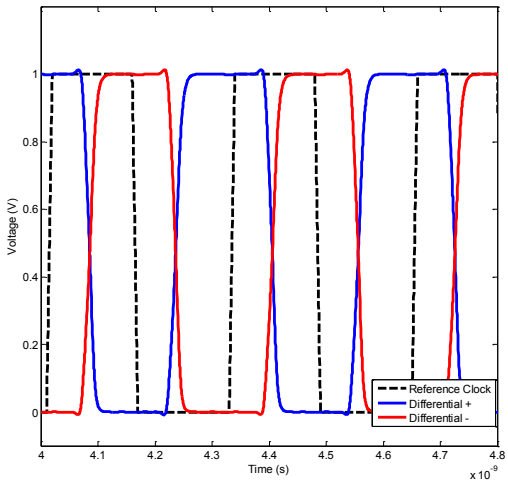

Figure 4-1: Differential clocks results from single-ended off-chip clock reference.

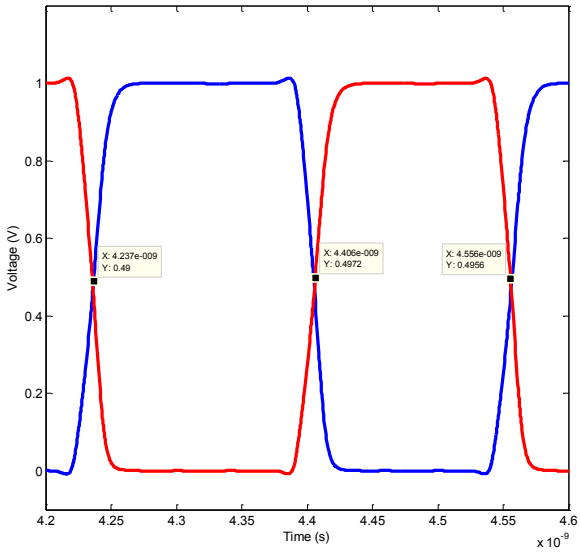

Figure 4-2: Differential clock duty cycle.

The graph in Figure 4-3 is a sweep of the input duty cycle and the corresponding output of one of the two differential outputs. This illustrates that the converter is dependent on the reference clock's duty cycle.

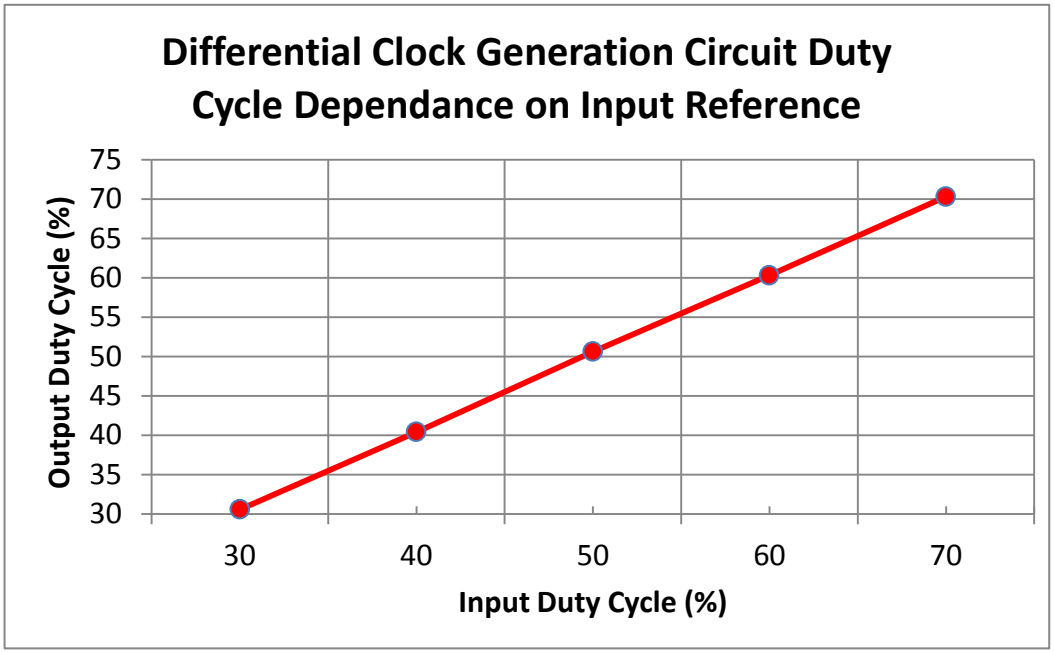

Figure 4-3: Output duty cycle dependence on reference's duty cycle.

4.2 Phase Detector

The results of the XOR CML phase detector are presented below. It is expected to follow the behaviour of an XOR logic gate with the truth table found in Table 4.

Table 4: XOR Truth Table.

IN_1	IN_2	OUT
0	0	**0**
0	1	**1**
1	0	**1**
1	1	**0**

The simulated output is shown in Figure 4-4, along with a visual indication to help identify the XOR operation. The figure is comparing one phase of the differential reference clock, against one half of a differential phase from the ring VCO bank. The results show that when both the reference clock and the VCO signal are equal (high or low logic level), then the output is pulled low. On the other hand, if both signals are different, the output is high.

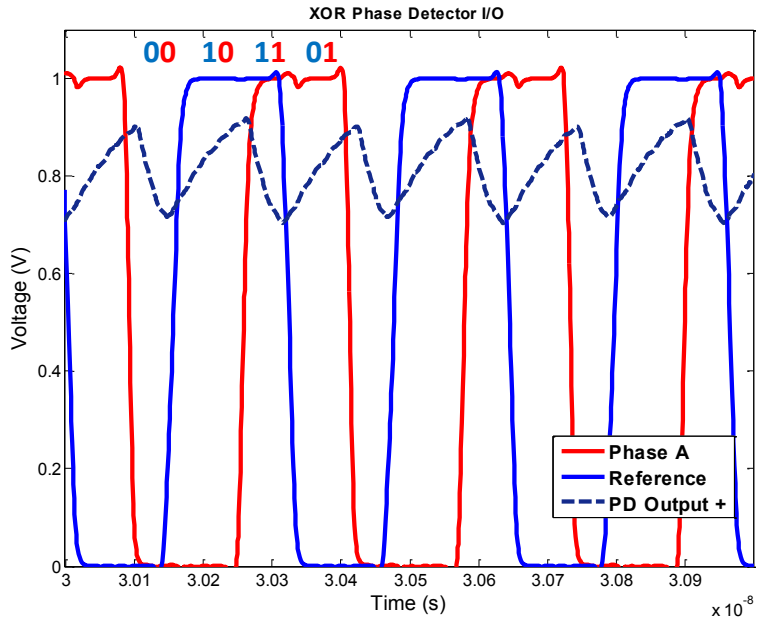

Figure 4-4: CML XOR phase detector output.

One of the major advantages of the XOR phase detector is that it has the ability to update two times per reference cycle. This can be seen in the results above where two output pulses exist in one period of the reference clock cycle. Not shown in this image is the other half of the phase detector differential signal, which mirrors this operation.

4.2.1 Transconductor

The transconductor accepts the differential output of the phase detector and outputs a current into the loop filter. Presented below is the current output of the detector under normal locked conditions. The bias current of this block is 200µA. This current output was recorded for a loop filter voltage (V_{LF}) of approximately 0.48 V. Referring back to Figure 3-15 where the loop filter voltage is an indication of the source-drain voltage for the NMOS, and is $1-V_{LF}$ of the source-drain voltage of the PMOS in the transconductor. These source-drain voltages will ultimately affect the current supplied by the transconductor. As the output resistance of the circuit is not infinite, it is observed that the output current is less than the bias current of 200 µA. Figure 4-5 shows the current output of the transconductor under phase locked conditions, with the inputs signals to the phase detector being in quadrature.

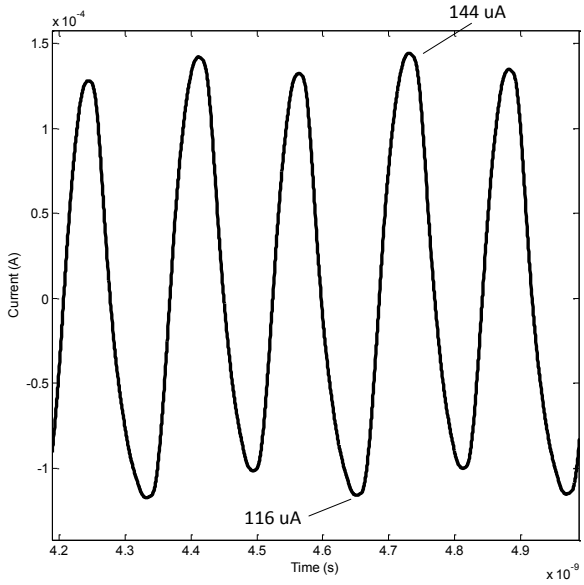

Figure 4-5: Current output of transconductor.

Next, the inputs to the phase detector were swept and the output current recorded. The results are shown in Figure 4-6, where the average transconductor current is plotted against the phase error. Ideally, the average current would be zero when the reference and the output phase are in quadrature, or 80 ps of the 320 ps period. However, due to non-ideal circuits, the current is slightly skewed. This causes the zero average to occur at a phase error around 65 ps. The maximum and minimum currents show the current for a phase error of 180 degrees (160 ps) and 360 degrees (320ps).

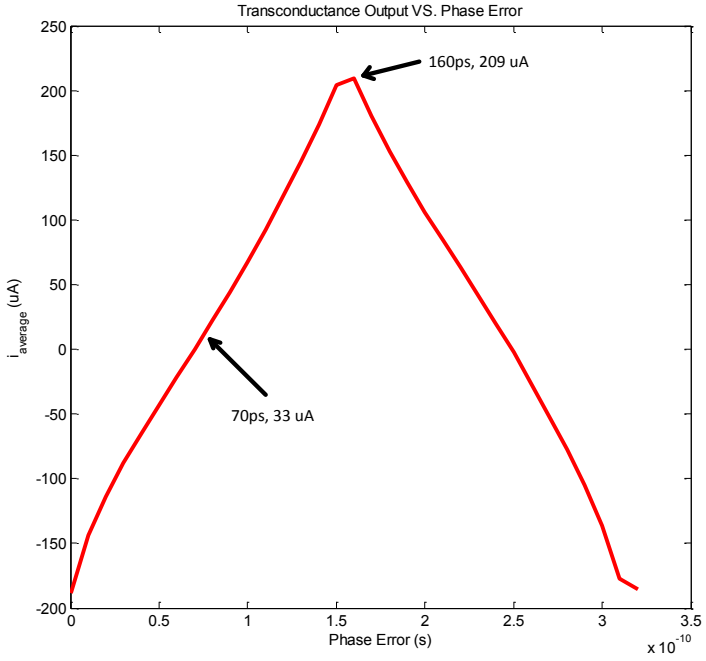

Figure 4-6: Current output of transconductor.

4.2.2 Loop Filter

The current moved in to and out of the loop filter by the transconductor makes the frequency tuning of the VCO possible. The voltage that develops on the loop filter, V_{LF}, is presented in Figure 4-7. This rippling voltage is due to the current ripple of the transconductor. It is seen that the PLL becomes phase locked in approximately 7 ns.

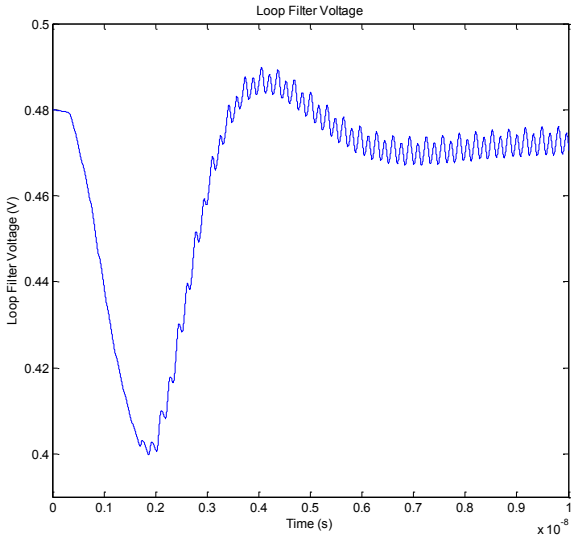

Figure 4-7: Control line of sub-VCO bank.

4.3 VCO

The oscillation frequency and phase difference of each stage plays an important role in the correct functionality of a sub-rate, multi-phase CDR. Figure 4-8 illustrates a transient analysis for eight of the sixteen output phases.

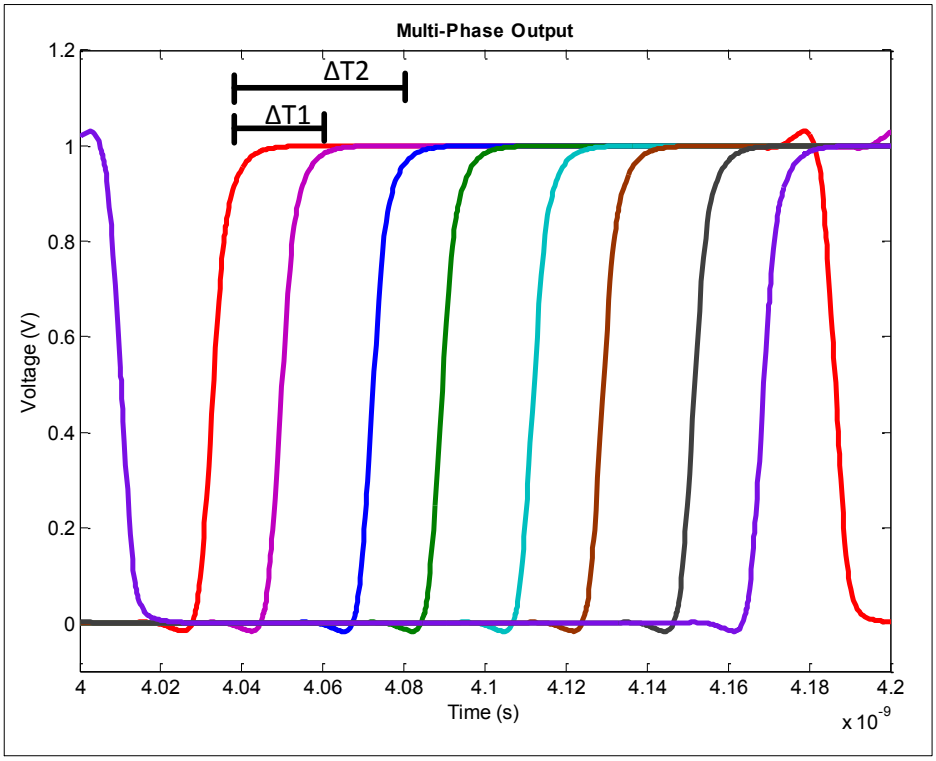

Figure 4-8 Multi-phase output at common node

In order to further analyse the multi-phase alignment,

Table 5 below was assembled. It is a measure of the difference in time that each phase is relative to phase one, relative to the ideal value theoretically expected.

Table 5: Difference between phases of oscillator relative to phase 1.

Phases	ΔT from Ideal Sampling Point (ps)	Phases	ΔT from Ideal Sampling Point (ps)
1 to 2	1.76	1 to 10	3.69
1 to 3	-0.47	1 to 11	-1.56
1 to 4	1.42	1 to 12	3.14
1 to 5	-1.6	1 to 13	-3.56
1 to 6	-0.5	1 to 14	-0.27
1 to 7	-3.1	1 to 15	-4.2
1 to 8	0.35	1 to 16	0.6
1 to 9	-1.9		

4.3.1.1 Post-Layout Results

The sub-VCO block was the only block to show a large decrease in performance. Other blocks did not show a significant degrade in performance during parasitic extraction simulations. The reason for decreased performance of the sub-VCO is due to the capacitive loading on the nodes of the ring VCO. Variable feed-forward (FF) paths were implemented in the design in order to account for extra capacitance post fabrication. This included two inverters that were able to be independently activated. During schematic simulations, a 5 fF capacitor was placed at each stage in the ring in order to simulate some wiring and node capacitances. The following frequencies were recorded using the different feed-forward (FF) activation schemes using the schematic. The term *FF0* refers to the feed-forward inverter of drive strength 0 and the *FF1* refers to the feed-

forward inverter of drive strength 1.Next, the same simulation was run with the extracted layout values, where parasitic capacitances were present. The following values were recorded. Both tables were done in an open loop simulation with a control voltage of 0.5 V.

Table 6: Feed-forward strengths and schematic vs. extracted simulated results.

	Oscillation Frequency (schematic simulation)	Oscillation Frequency (extracted simulation)
FF0 =0 ; FF1=0	1.84 GHz	995 MHz
FF0 =0 ; FF1=1	3.24 GHz	1.55 GHz
FF0 =1 ; FF1=0	2.54 GHz	1.29 GHz
FF0 =1 ; FF1=1	4.37 GHz	1.92 GHz

This large change in performance can also be seen when comparing the tuning curves. Figure 4-9 shows the simulation results using schematic simulation values along with the parasitic extracted ones. It is evident that the overall frequency and tuning dropped due to the extra wiring capacitance and other device parameters that were not taken into account during schematic simulations. In order to obtain the simulation curve, only one of the possible two feedforward inverters were activated. For the parasitic extracted case, both feedforward inverters were needed in order to improve the frequency of oscillation as it was heavily degraded by the extra loading capacitance.

In order to improve these results, schematics and the circuit layout can be improved. In the schematic, a better job needs to be done in approximating the node parasitic

capacitances. This can be done possibly by attempting layout at an early stage of the design, and finding parasitic capacitances and adding them to the schematic during simulations. In order to improve the layout, work must be done to ensure that wire lengths and loading are as identical as possible on each stage of the ring.

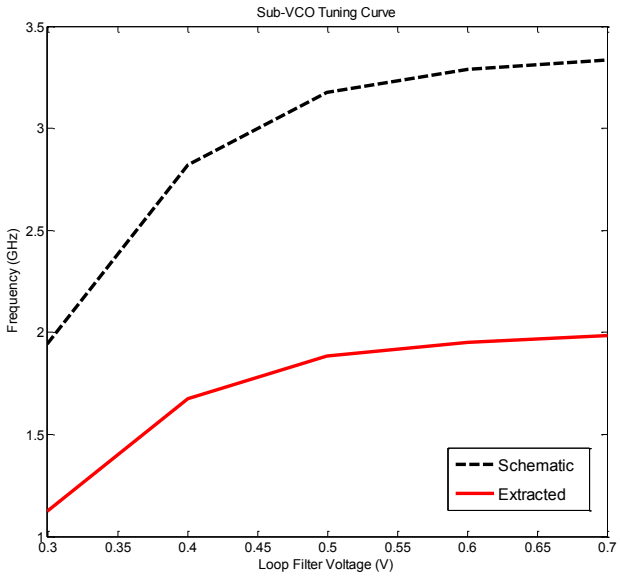

Figure 4-9: VCO tuning curves before and after parasitic extraction.

4.3.2 Interpolation

The ability to interpolate between phases presents finer granularity in phase steps. As discussed before, a secondary loop can adjust which phase of the VCO the phase detector will phase lock with. Figure 4-10 below shows the change in phase lock and the movement from phase 1 in quadrature to phase 2 and then phase 3. This can continue up to phase 16, after which, it will restart at phase 1.

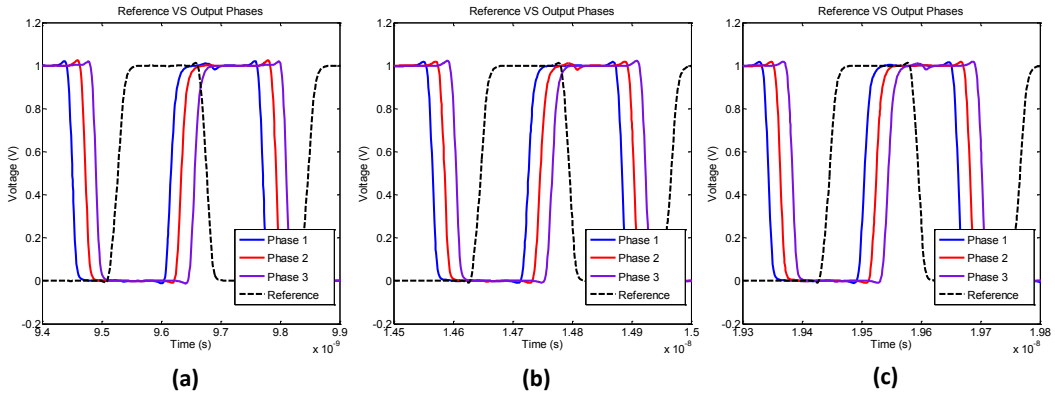

Figure 4-10: PLL phase locked to (a) phase 1 (b) phase 2 (c) phase 3.

4.4 Phase Noise and Jitter

The phase noise associated with an open loop simulation is presented in Figure 4-11. It was obtained through noise analysis methods in Cadence simulation software. The plots are of one, two and three active sub-VCOs respectively. The frequency of oscillation is 3.125 GHz, and the frequency shown on the x-axis is the offset in frequency, relative to the frequency of oscillation.

Recalling the theoretical noise reduction called for in (2-9), the following values are expected for two and three sub-VCOs respectively:

$$10 \cdot \log\left(\frac{1}{2}\right) = -3dB$$

$$10 \cdot \log\left(\frac{1}{3}\right) = -4.8dB$$

The simulations follow the predicted trend, as a phase noise improvement seen is 3 dB and 4.7 dB when activating an extra two and three active sub-VCOs, respectively.

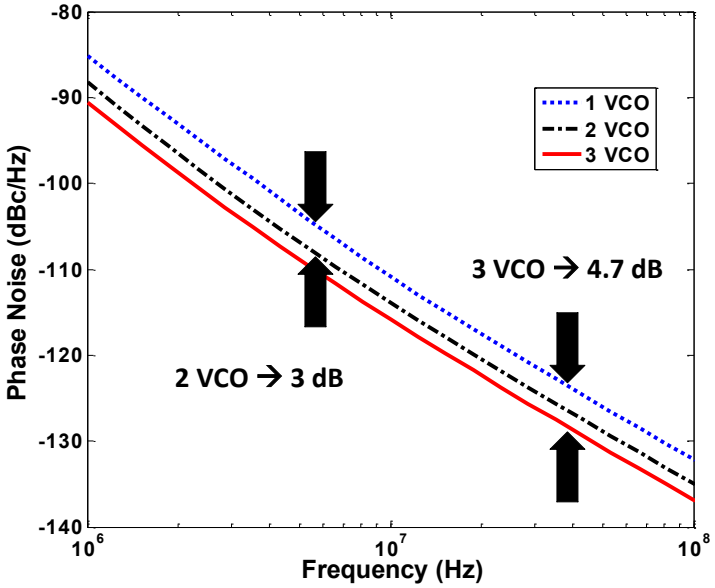

Figure 4-11: Open loop phase noise simulation.

When put into a closed loop system (such as the PLL), the behaviour seen in Figure 4-12 is observed. The open loop response was transposed on top of the closed loop phase noise plots in order to gain more insight. It is shown that the open loop and closed loop plots converge beyond the loop bandwidth, as expected. However, inside the loop bandwidth, simulation shows that there is less than the expected 4.8 dB drop in phase noise. It is believed that this is due to other noise sources being dominant in the simulation.

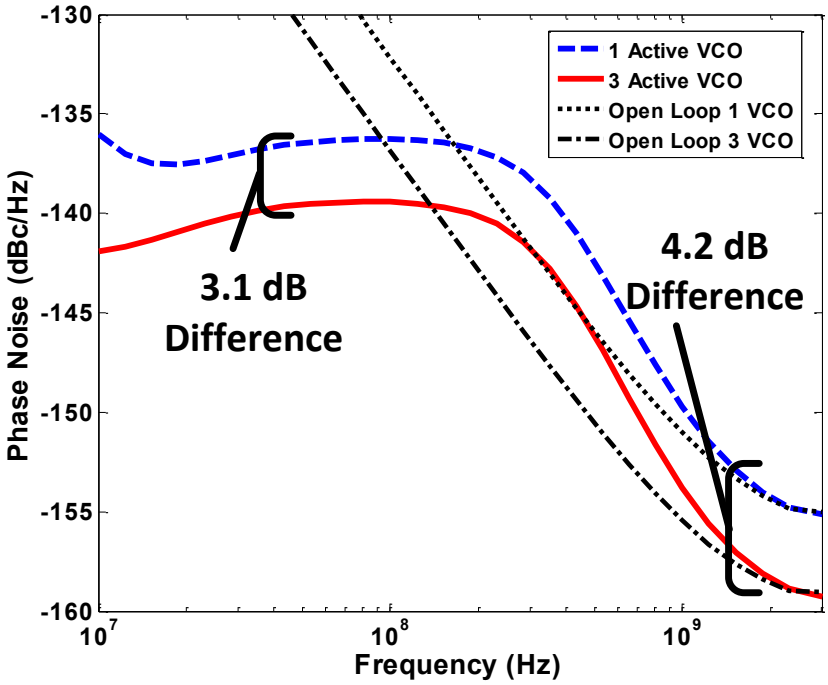

Figure 4-12: Closed loop phase noise simulation.

In order to test this hypothesis, it was necessary to see the dominant noise sources in the circuit. Figure 4-13 shows the results of this phase noise simulation.

Device	Param	Noise Contribution	% Of Total
/R3	rn	8.80262e-16	12.81
/M0	id	5.48728e-16	7.98
/I1/M3	id	3.93895e-16	5.73

Spot Noise Summary (in V^2/Hz) at 150M Hz Sorted By Noise Contributors
Total Summarized Noise = 6.87203e-15
No input referred noise available
The above noise summary info is for pnoise data

Figure 4-13: Dominant noise sources (results of PNoise simulation).

This shows that the top three dominant sources of noise are R3 (the loop filter resistor), M0 (the current starving MOSFET) and M3 (NMOS of transconductor). These results

suggest that noise sources other than the sub-VCOs contribute the most, such that the improvement was not as significant as originally expected. This situation is expected to improve when testing the fabricated chip as the VCO is expected to be the dominant source of noise.

4.5 Power Dissipation

The power that this architecture will dissipate is expected to be inversely proportional to the phase noise, as theory predicts in (2-9). Shown below is a plot of the analog supply current, functioning at 1V DC. The plot is also labelled to indicate when the new sub-VCOs are added in. This increases the current draw and thus the total power dissipation.

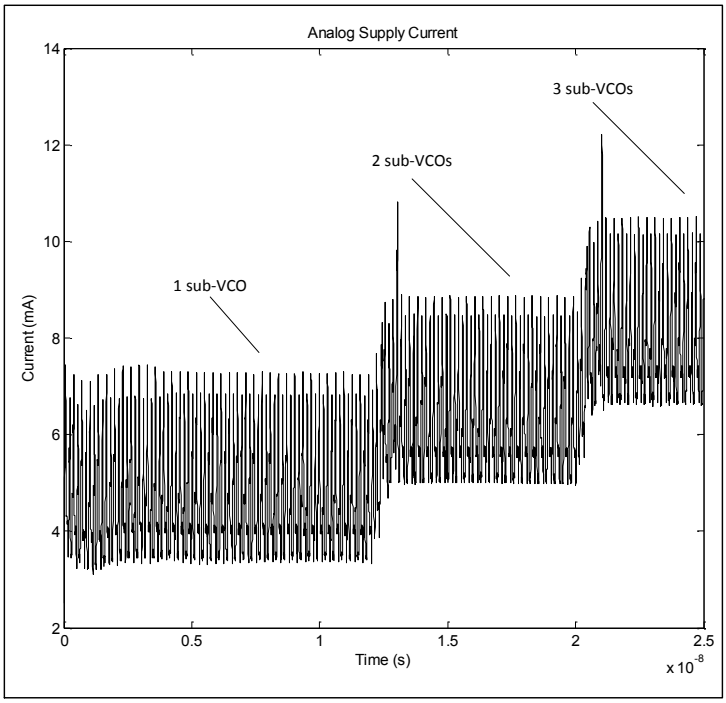

Figure 4-14: Current from analog supply.

There is a 3mA ripple on the analog supply due to the differential clock circuit. This clock circuit serves as the reference and so any noise present inside the loop bandwidth

will modulate onto the VCO control voltage. The analog supply was used so that the digital supply, which will have more switching events and feeding multiple other circuits, would not couple in the noise to the reference. In order to extract information on the actual sub-VCO component, Figure 4-15 is included. The extra activity at the transitions is due to the timing logic used for correctly enabling the sub-VCOs. This indicated that by switching from the high-performance mode (3 sub-VCOs) to the low-power mode (1 sub-VCO) a power savings of 46% can be achieved.

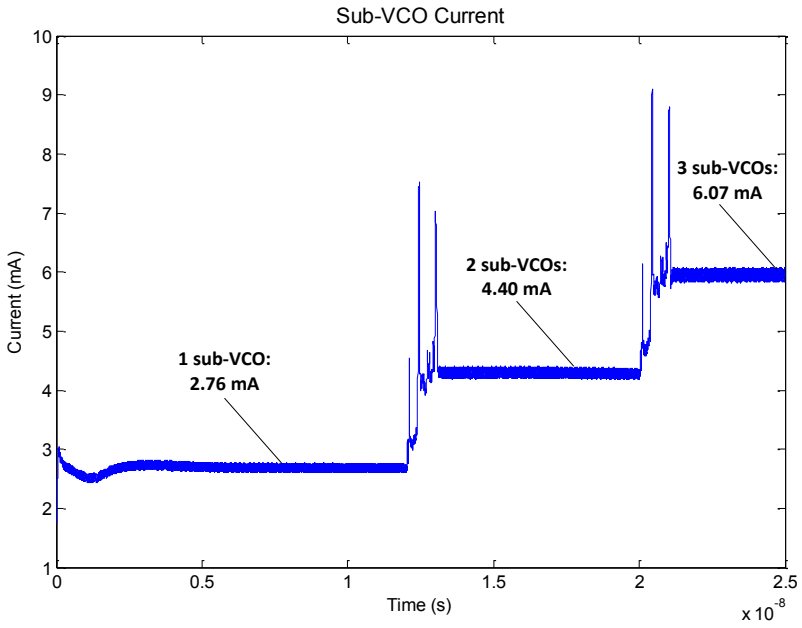

Figure 4-15: VCO bank current draw.

The PLL in Toifl et al., consumed 19.8 mW of power for a 2.5 GHz signal, using 90 nm CMOS SOI technology [10]. In the work presented in this thesis, the power of the high-performance is roughly about half of this. The phase noise simulations also appear to be about 10 dBc/Hz lower than the published results in [10], albeit from measured results

4.6 Excursion Plots

The compensation methods discussed previously will now be tested for efficiency. Figure 4-16 below is an excursion plot showing non-compensated sub-VCO activation as well as a compensated sub-VCO activation. Both simulations came from the same circuit, the only difference was that for the non-compensated case the activation sequence was not used, and all stages were started up together. Also, the variable capacitors were deactivated and variable buffers became fully activated upon activation instead of a power-proportional control scheme.

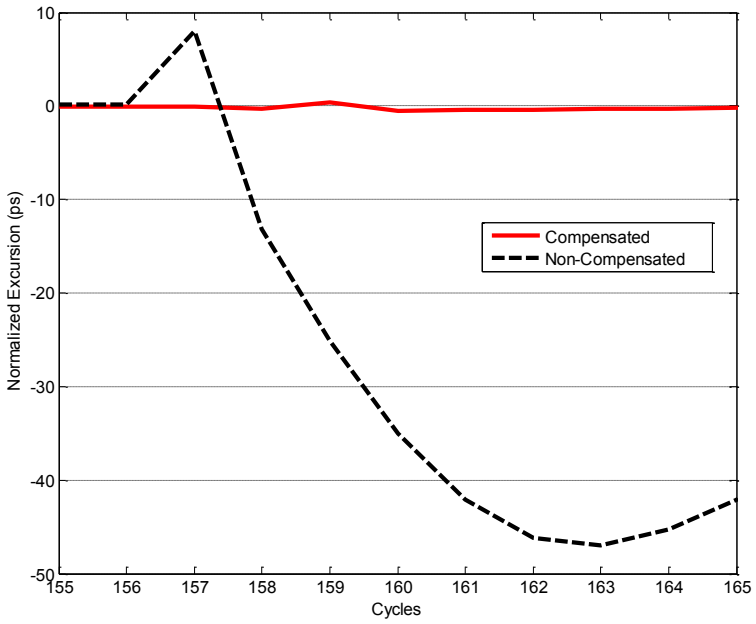

Figure 4-16: Excursion plot of 1 to 3 enabled sub-VCOs.

The results of this simulation are interesting and promising. The non-compensated case gets perturbed and undergoes a phase excursion of almost 50 ps. On the other hand, the compensated case has an excursion of only 3 ps. In other tests it was found that the actual structure of the ring helped the phase excursions. This is to say that in simulations where

only switches were used to dynamically connect sub-VCOs, the PLL would lose lock in most cases or experience very large phase excursions.

Table 7 has been assembled in order to show the phase excursions and how they are affected by the compensation methods. It is clear that the feedforward ring architecture has had a large impact on the phase excursions, but it is not clear yet as to the exact reasons for this.

Table 7- Phase excursions due to various compensation methods

Compensation Method	Excursion (Compensation Enabled)	Excursion (Compensation Disabled)
None	60 ps + (lost lock)	60 ps + (lost lock)
Variable capacitors	16 ps	58 ps
Variable capacitors + variable buffers	6.2 ps	58 ps
Variable capacitors + Variable buffers + FF Ring architecture	3 ps	47 ps

The recalibration of the system from 3 sub-VCOs to 1 sub-VCO happens in a reverse manner, with the phase buffers getting disconnected first. The resulting phase excursion, shown in Figure 4-17, is however slightly higher than the 1 to 3 sub-VCO case, at 18 ps. This is not a large problem as this reconfiguration can be done at the lower data rate and

thus the larger unit interval will allow for a larger phase excursion. This can be further improved upon by adjusting and improving upon the delay path logic, as it is the timing of the deactivation of each stage that causes these excursions.

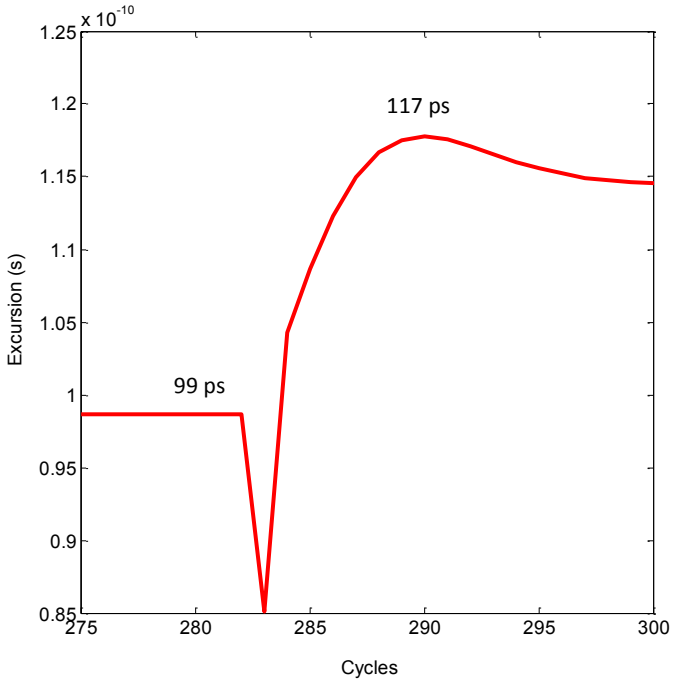

Figure 4-17: Excursion due to reconfiguration from 3 sub-VCOs to 1 sub-VCO.

4.7 Mismatch

The issue of mismatch was deemed from early on to be an issue with the proposed architecture. It was therefore tested with DC voltage supplies on each of the gates of the current starving transistors to simulate threshold differences. For one particular simulation, a DC source of +50mV was added to the gate of sub-VCO #2 and -50mV to sub-VCO #3. Figure 4-18 shows the current through each of the current starving

transistors for each of the sub-VCOs. It demonstrates that sub-VCO #1, which is always active, adjusts itself to compensate for the other rings in the closed loop PLL when they activate.

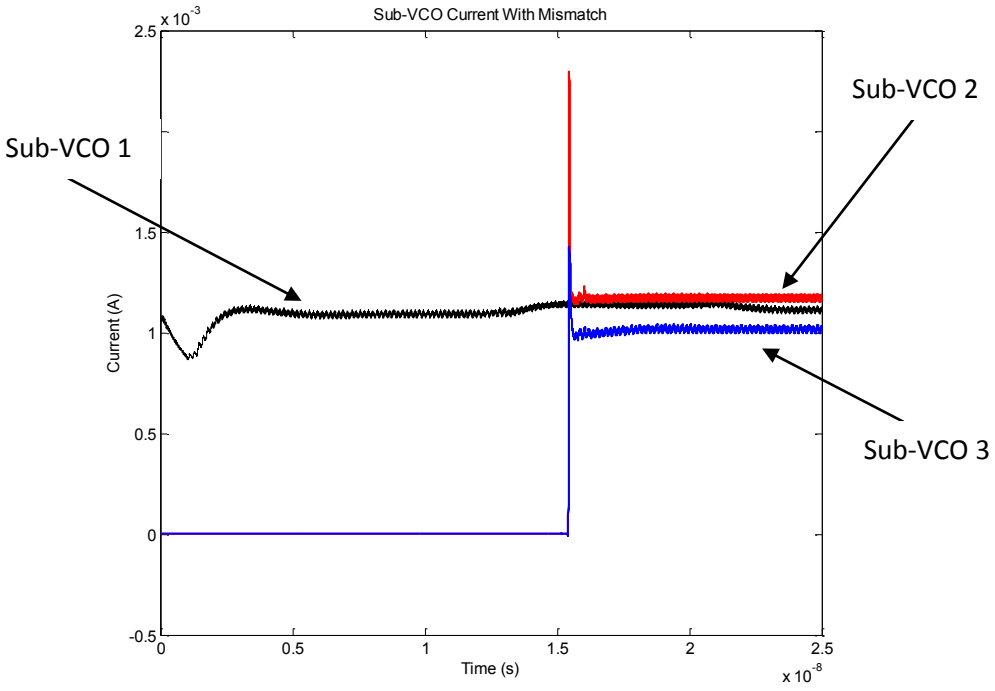

Figure 4-18: Current per sub-VCO under effects of threshold mismatch.

The results of the excursion simulations are shown in Figure 4-19. It was found that the even with mismatch, the phase excursions were kept to within 3 ps.

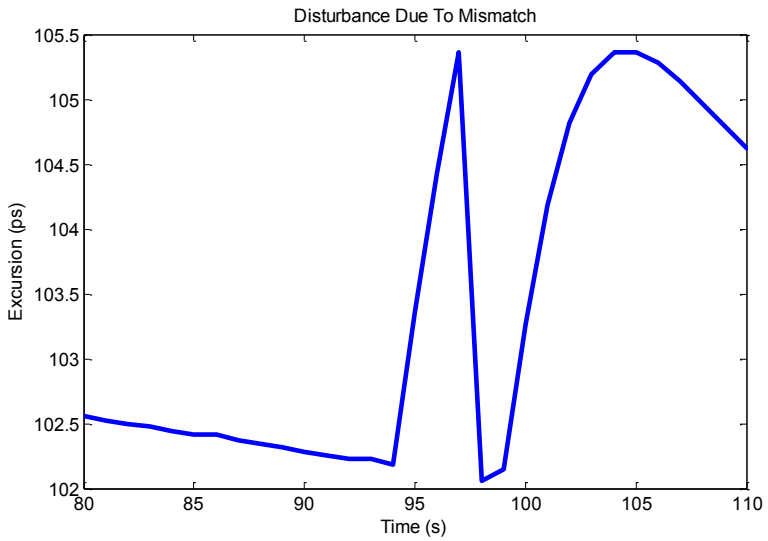

Figure 4-19: Excursion plot with effect of threshold mismatch.

Further testing was carried out to values of $\pm$ 100 mV to simulate large variations, without causing the architecture to fail. The excursion plot did however show that beyond $\pm$ 100 mV the excursions became bigger than 10 ps, which would still be acceptable. Threshold values are not expected to deviate to that extent however.

CHAPTER 5

Circuit Layout

The layout for this design was completed and submitted for fabrication in TSMC 65nm technology. Odd layers of metal were routed vertically and even layers were routed horizontally. The blocks presented in this chapter were designed and laid out by me, except for the following:

- Latch design and layout by Monir Moayedi and Michael Segev
- Barrel register design assisted by Ted Obuchowicz
- Barrel register layout by Michael Di Perna
- State register design and layout by Michael Segev
- IC power grid design and layout by Glenn Cowan and Michael Segev
- Pad ring design and layout by Glenn Cowan and Monir Moayedi

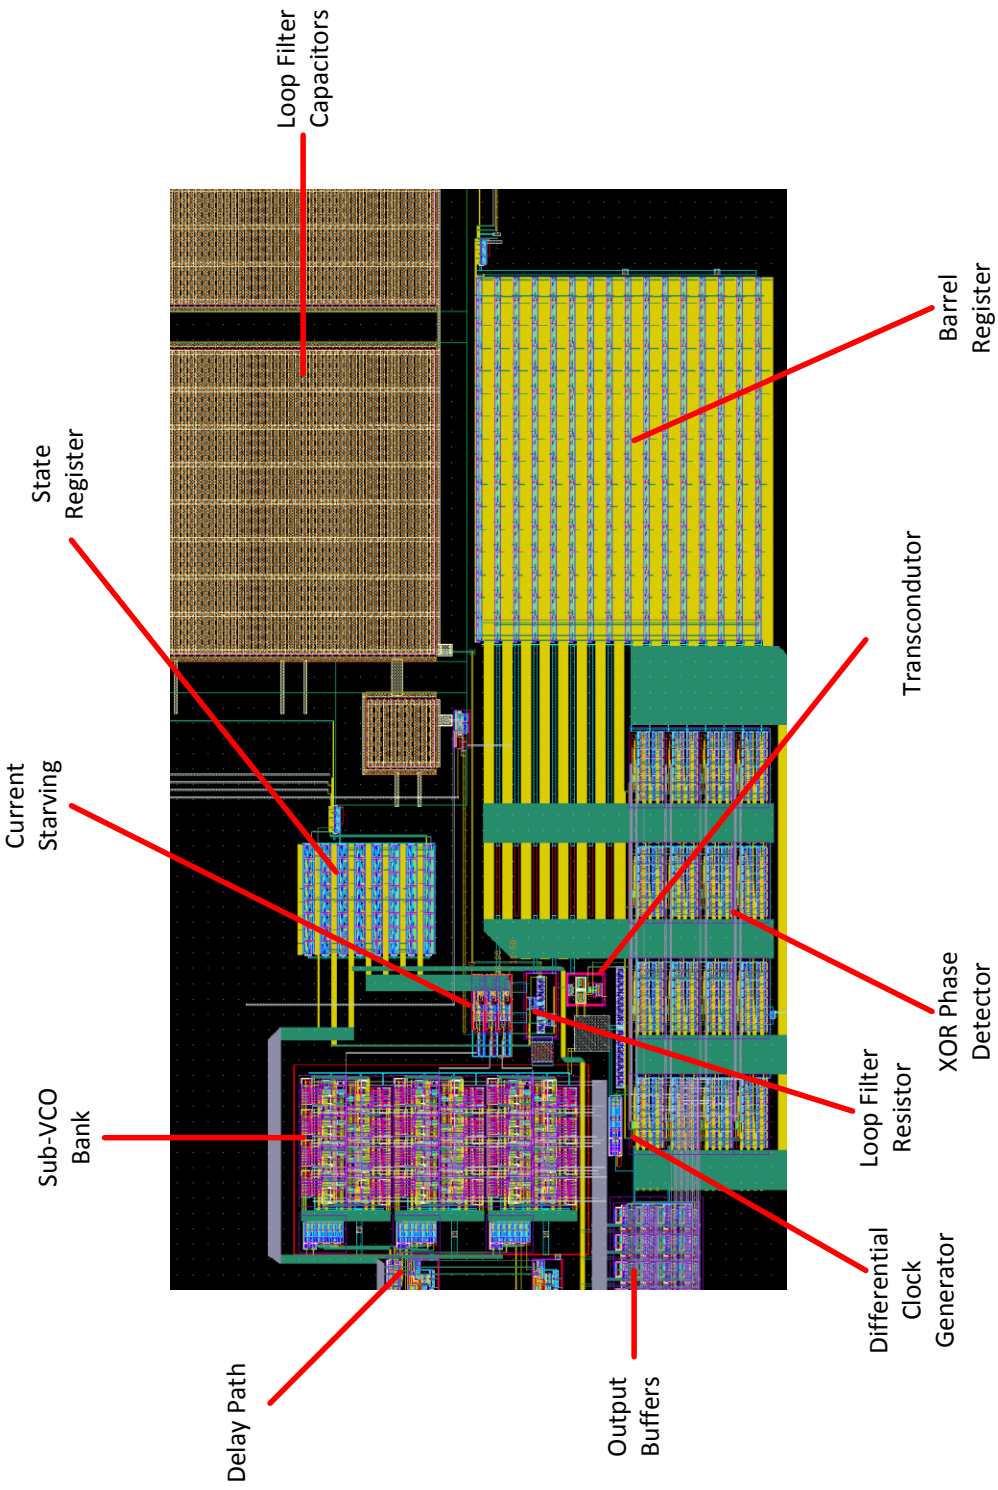

Figure 5-1: Overall layout of PLL.

5.1 Single-To-Differential Converter

The single-to-differential converter was made using CMOS logic from TSMC 65nm standard cell library. Care was taken in the layout process of this block to ensure symmetry between both differential signal paths. This was done in order to minimize duty cycle errors due to non-symmetrical loading. The layout of this block was approximately 4.30 µm by 23 µm.

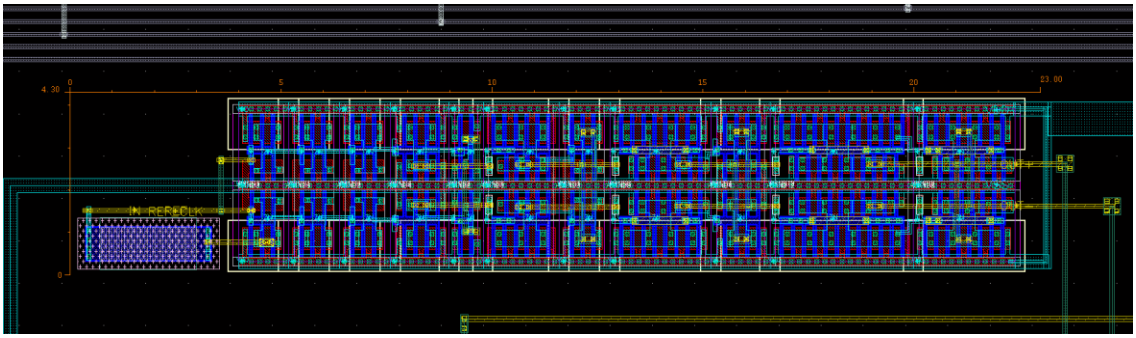

Figure 5-2: Single-to-differential clock converter.

5.2 XOR Phase Detector

The CML XOR phase detector was implemented using TSMC 65nm NMOS devices. N-wells sandwiched between p-type regions surround most of the area as a protective barrier. This was done with the hopes of ensuring a properly grounded substrate around the CML circuits, and a barrier from surrounding noise sources. The many wires are buses that connect the barrel register to the phase detector. The wire lengths are long, but it was determined that the capacitance and resistance associated with these run lengths were somewhat beneficial. As these are DC full rail signals going to gates of MOSFETs, the resistance would not affect the signal as no current would be flowing in steady state. The capacitance of the wires would also help reduce the noise associated with these nodes and thus reduce modulation into the phase detector. The phase detector has approximate dimensions of 135 µm by 46 µm.

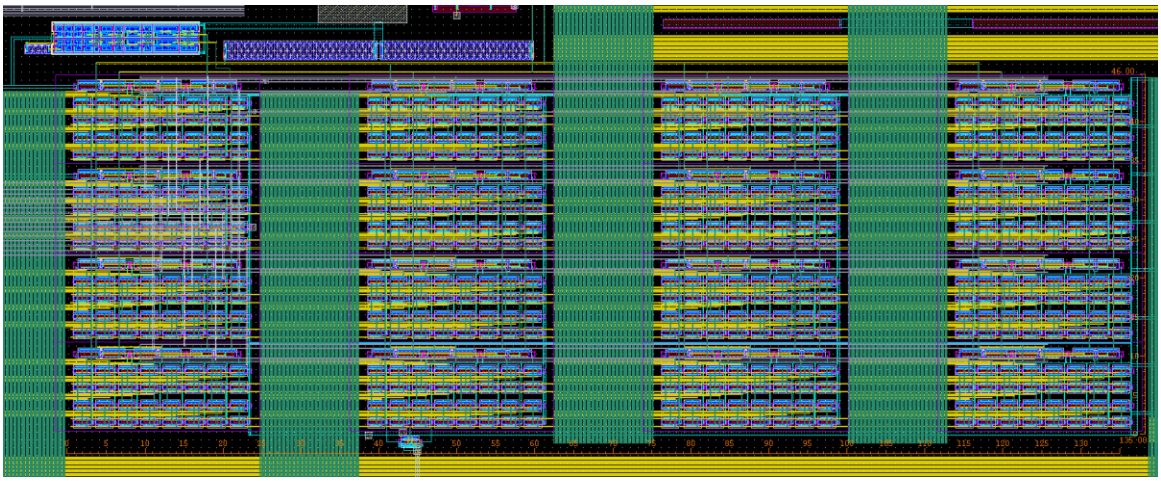

Figure 5-3: Layout of CML XOR phase detector.

5.3 Transconductor

The voltage-to-current converter was designed to be properly matched in terms of sourcing and sinking current. A common-centroid layout was used to accomplish this. In the NMOS common-centroid layout, a source becomes shared between two neighbouring transistors. The multiple fingers are spread out in order for the transistors to be subjected more equally to fluctuations due to mismatch and temperature variation [12].

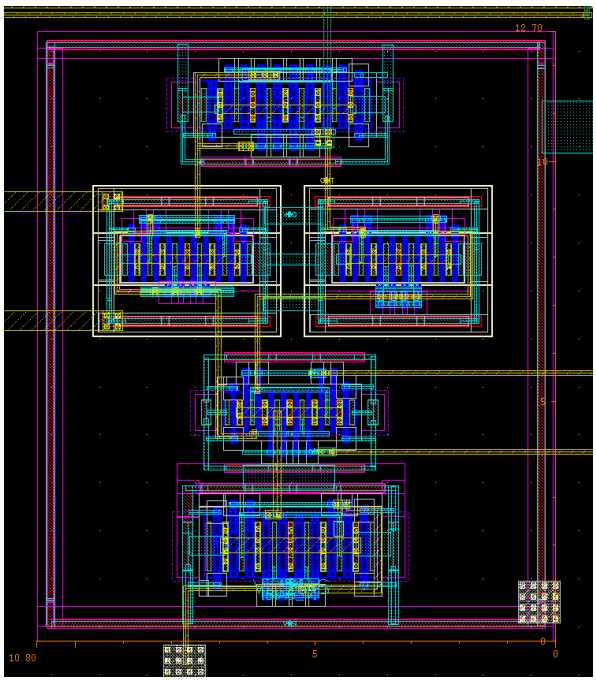

Figure 5-4: Layout of transconductor.

5.4 Loop Filter

As is normally the case, trade-offs are often present in circuit design. Analog PLLs do not suffer from quantization errors on the control line and are simpler to design than their digital counterparts. However, the downfall of the analog PLL is the size of the passive components. This becomes evident in the layout as the loop filter capacitors are about the size as all the other circuits combined in the primary loop. The area for the 40 pF (2 x 20pF) capacitor is 200 µm by 100 µm. The 1 pF capacitor is 25 µm by 25 µm. The metal wires protruding off the sides are for connections to the power grid of the top level chip.

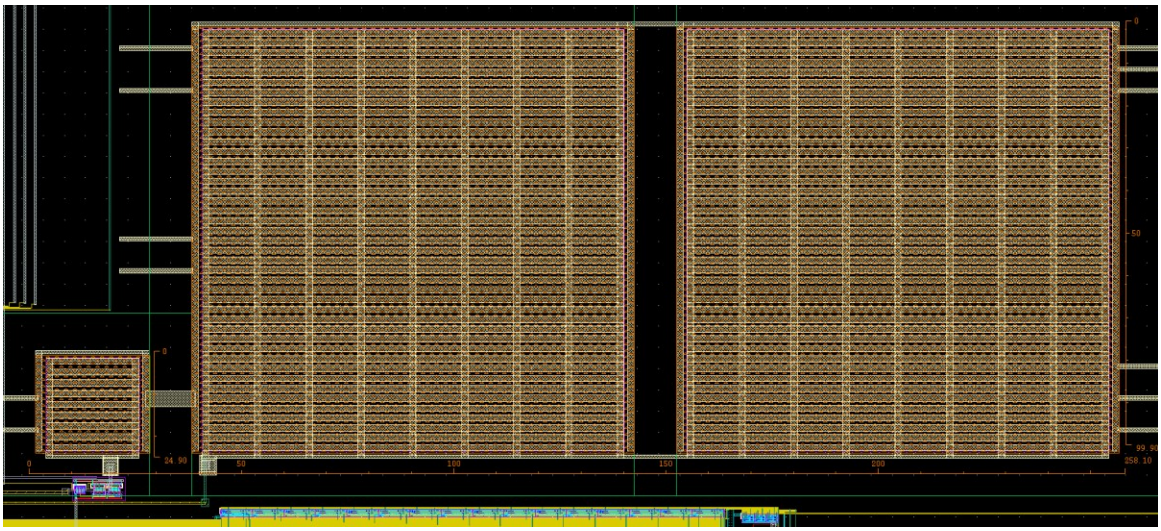

Figure 5-5: Layout of loop filter capacitors.

The variable resistance of the loop filter is shown in Figure 5-6below. The NMOS transistors on either side of the series resistors allow for a portion of the resistance to be shorted. This gives the chip tester the ability to change the loop filter's zero. The NMOS switch adds 60 Ω to the signal path, which was taken into account when designing the resistor values. The block measures 20µm by 10µm.

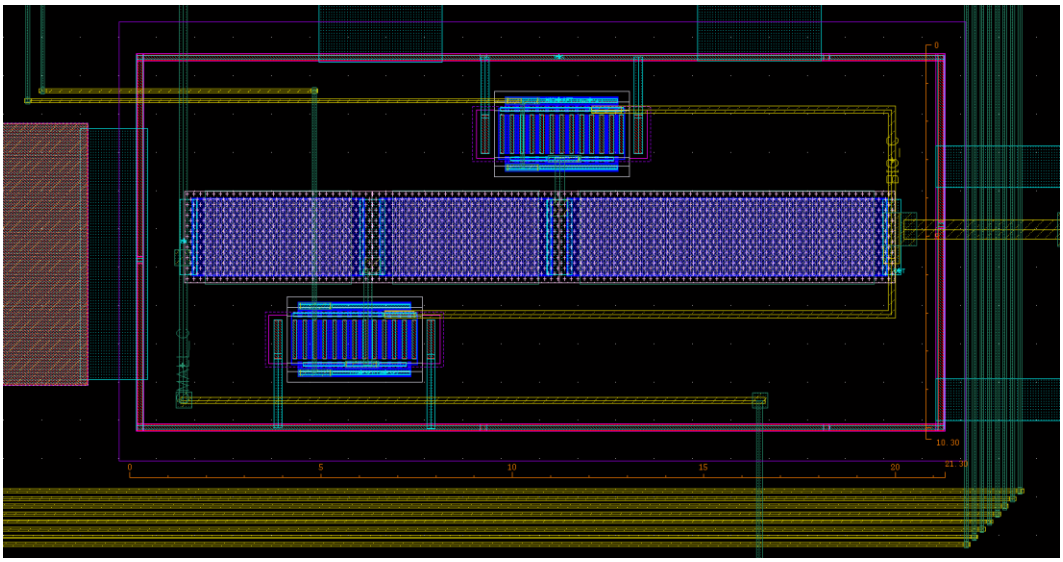

Figure 5-6: Layout of loop filter variable resistor.

5.5 Delay Line Logic

The delay line logic was implemented using standard cell inverters. The capacitors, used to increase the per stage delay in the chain, were implemented using gate capacitance of NMOS transistors. These transistors are the blue rectangles seen below the inverter chain in the bottom left corner of Figure 5-7.

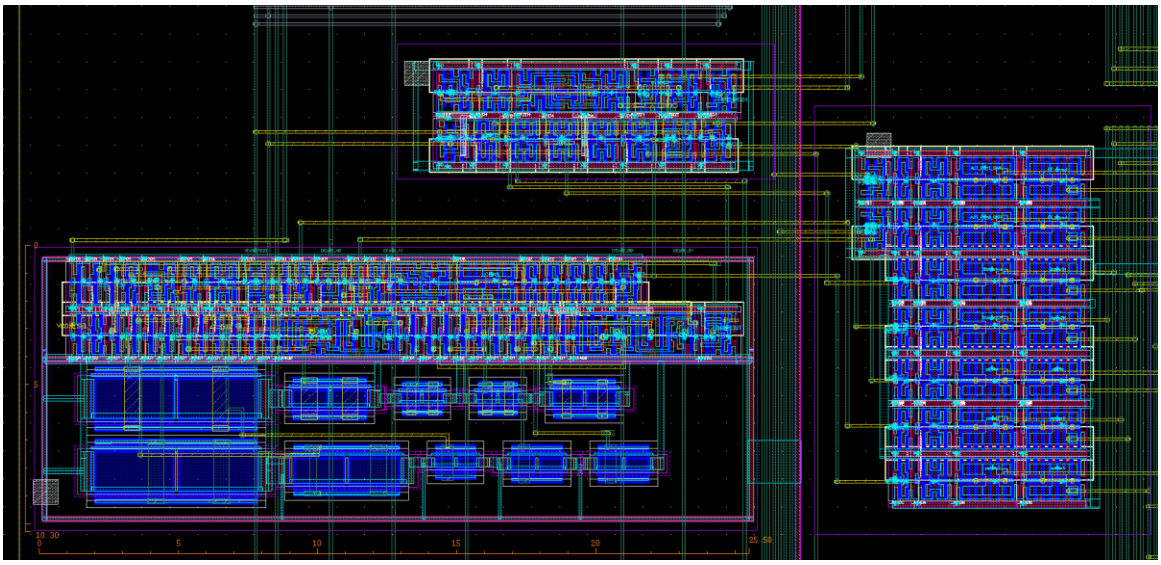

Figure 5-7: Layout of sub-VCO activation logic.

5.6 VCO

The layout of one sub-VCO is shown below in Figure 5-8. The ring was laid out with the goal of having equal loading from wires on every node. In order to do this, the top row contains four of the eight stages, which then are mirrored below for the full eight stages. This was done to minimize the length of wire in between stages. Each sub-VCO is approximately 30 µm x 40 µm. The signal path of the sub-VCO is illustrated in Figure 5-9.

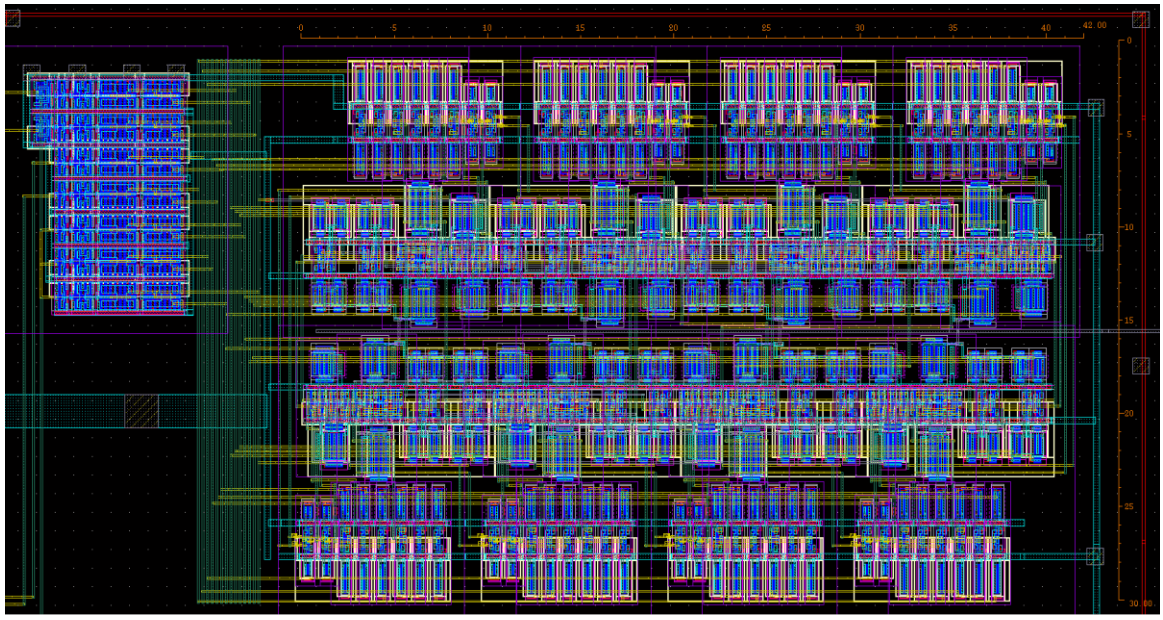

Figure 5-8: Layout of a sub-VCO.

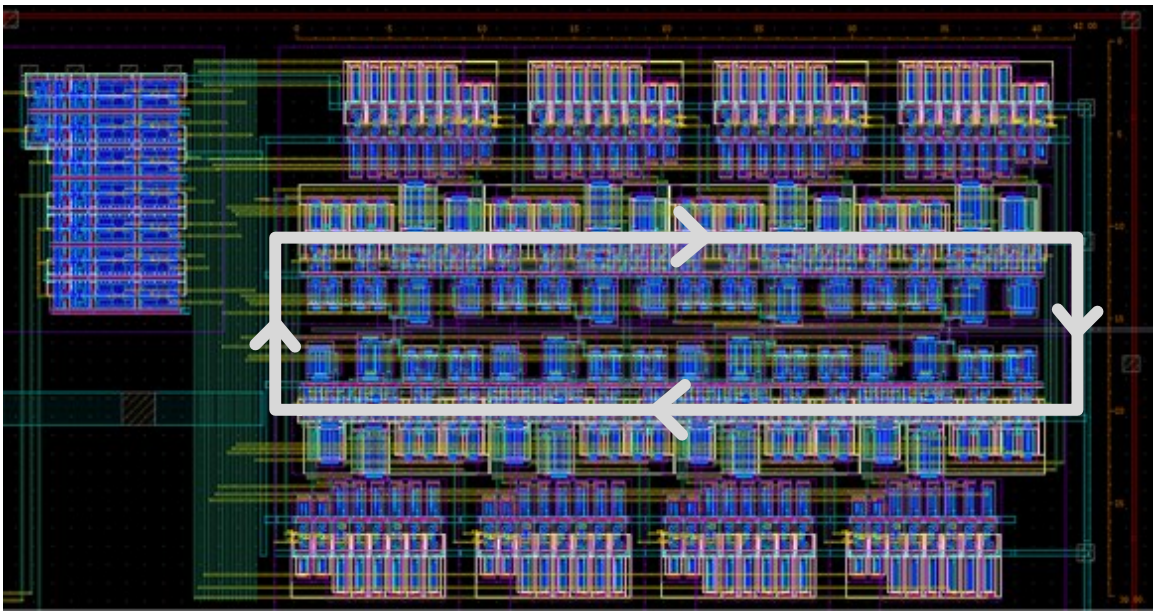

Figure 5-9: Signal path for sub-VCO.

The VCO bank, containing three sub-VCOs, is shown in Figure 5-10. Each sub-VCO used identical layouts.

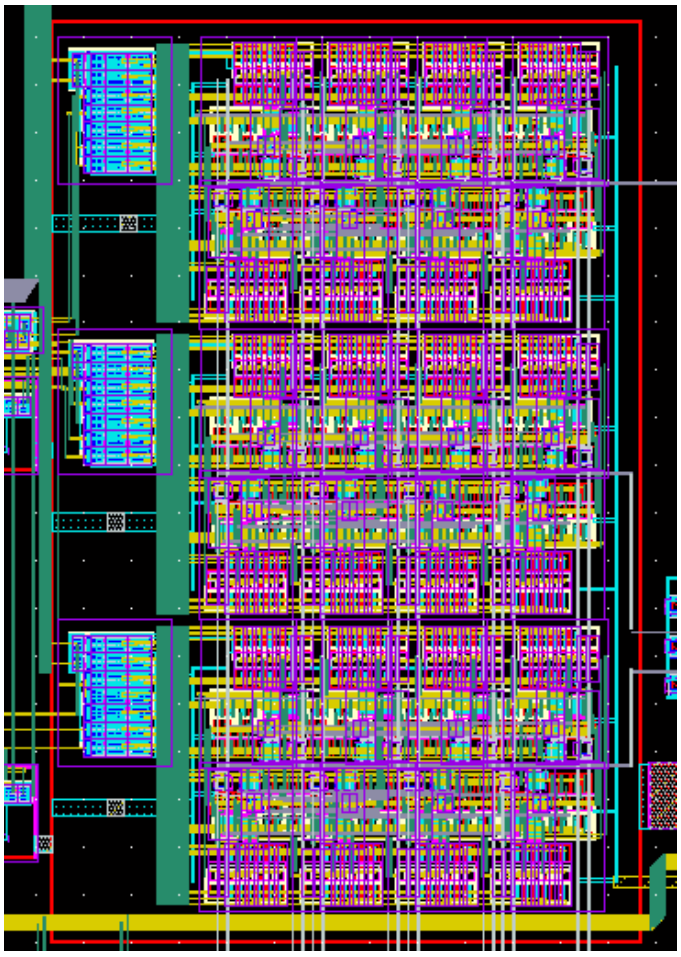

Figure 5-10: Layout of VCO bank.

Over the top of the VCO bank, metal 5 wires ran vertically. These wires allowed the connection of each of the sub-VCO phase nodes together, creating the 16 common phase nodes. The metal 5 layer was also desired as it was higher up and away from the grounded substrate and metal 1 wires, most of which presented AC grounds. This reduced the capacitance associated with these wires. Metal 4 below and metal 6 and 7 above the VCO bank were also kept clear, in order to reduce capacitance further.

5.7 Current Starving and Compensation

The layout for the current starving transistors is shown below in Figure 5-11. The large transistors on the left are the current starving ones, and the smaller, more numerous transistors on the right are for compensation schemes. This block controls the current for all three sub-VCOs. The total size of this block is 13 µm x 27 µm.

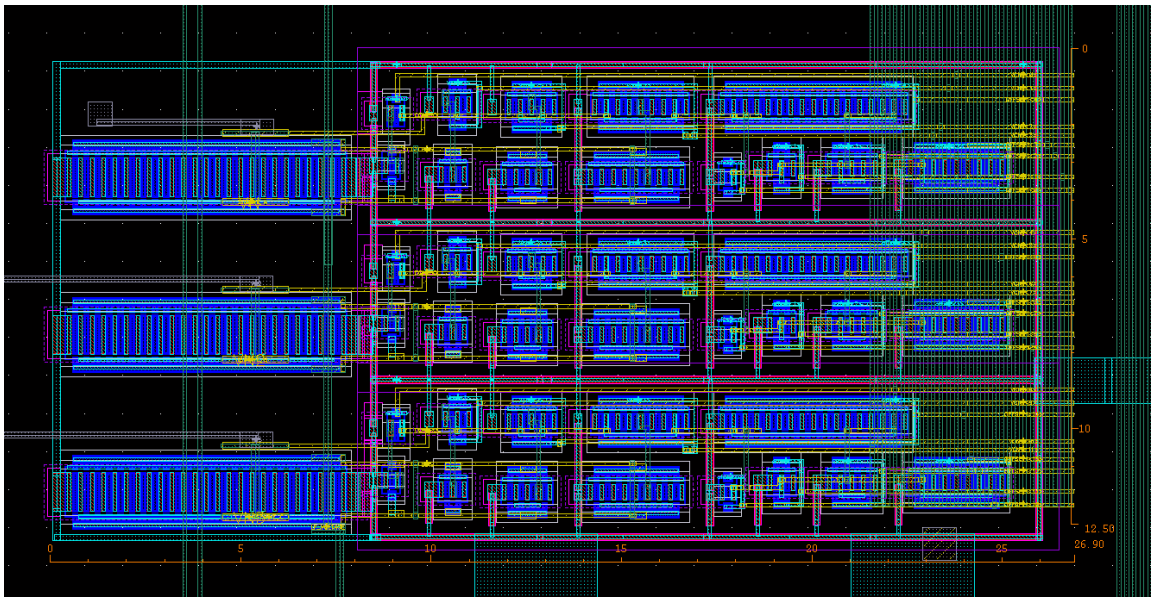

Figure 5-11: Layout of current starving and compensation transistors.

5.8 Output Buffers

The layout of the output buffers are presented in Figure 5-12below. Care was taken to ensure equal routing of metal on the output of each buffer to the phase detectors and output MUX. This was done in an attempt to make all phases as equal as possible. The total size of the layout is 28 µm x 38 µm.

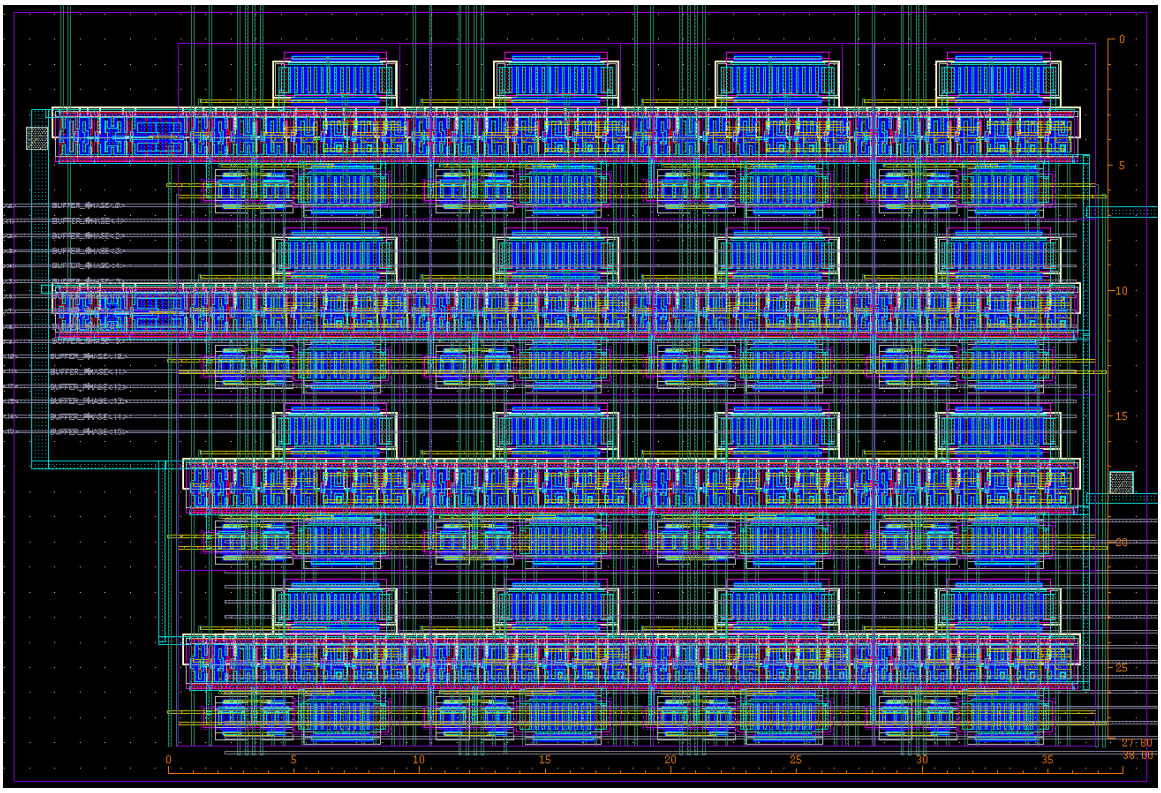

Figure 5-12: Layout of output buffers.

5.9 Barrel Register

The barrel register contains 256 D-flip flops and 256 MUXs. This connects buses that feed the phase detector to enable interpolation. This block is run off of a separate digital supply. The total size of this block is 95 µm by 123 µm.

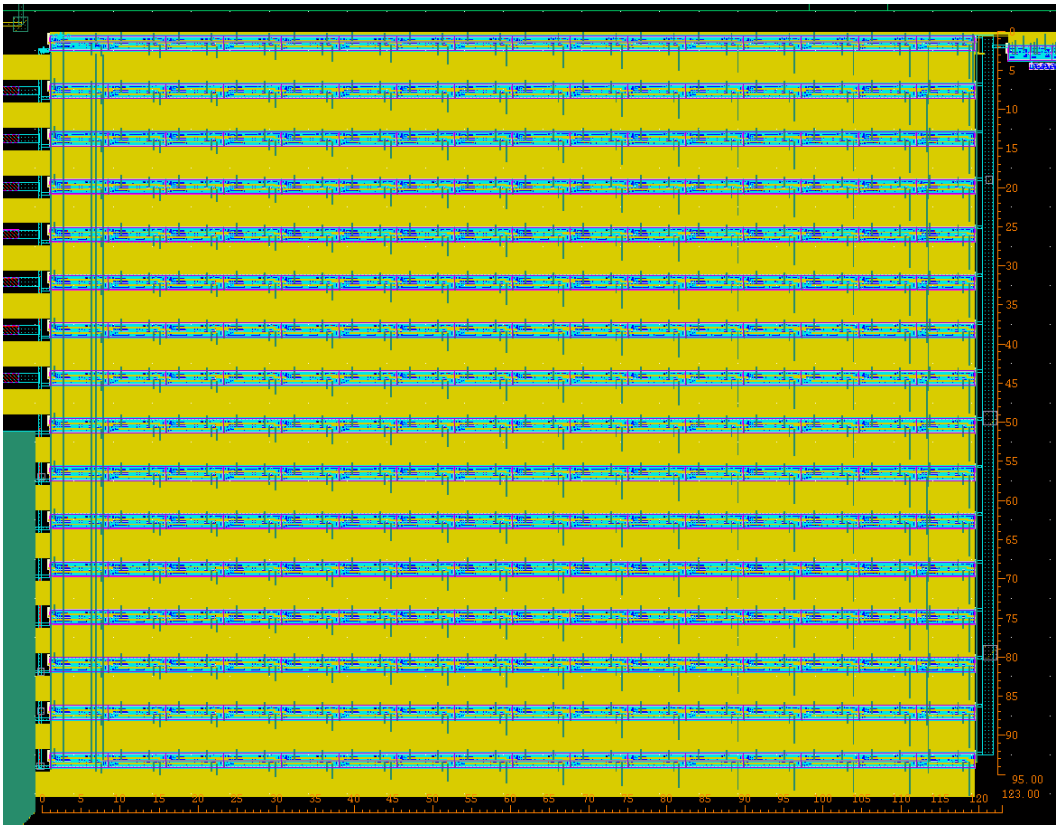

Figure 5-13: Layout of barrel register.

5.10 State Register

As there are a small number of available pins to connect the chip to the outside world, the serial register is responsible for serially taking in data and storing 64 control bits. The serial register layout is shown in Figure 5-14. This circuit contains 64 flip-flops that accept the serial data in. The flip-flop outputs then get latched into 64 registers for 64 state bits of information. The latched information is then sent to the various places on the chip that use these state bits. The size of this block is 46 µm x 37 µm.

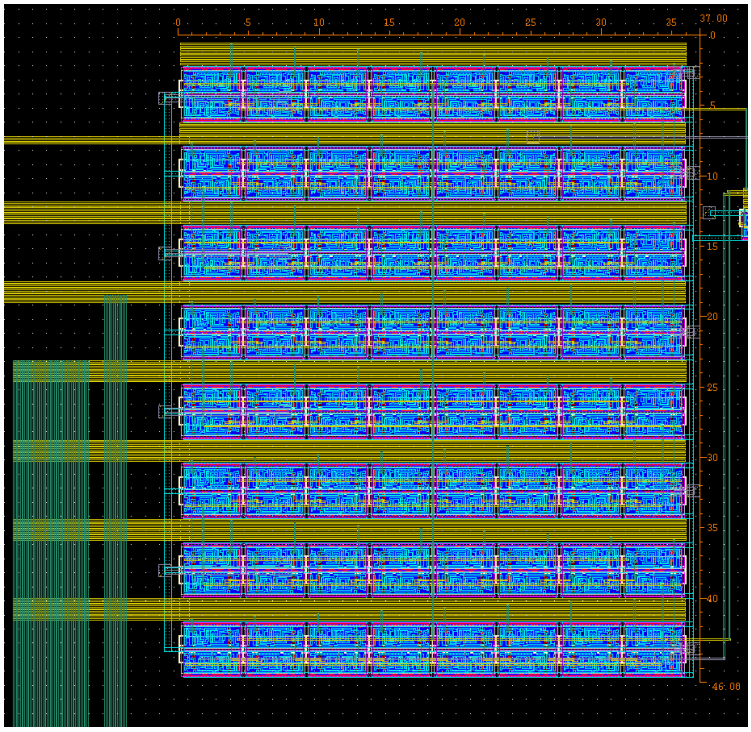

Figure 5-14: Layout of state register.

5.11 Test Chip Layout

The images in this section show the layout of the receiver test chip submitted for fabrication. These figures include a power grid to supply power to all circuits, and pads to connect up the I/O to the test board or probes. The entire chip measures 1mm by 1.5 mm.

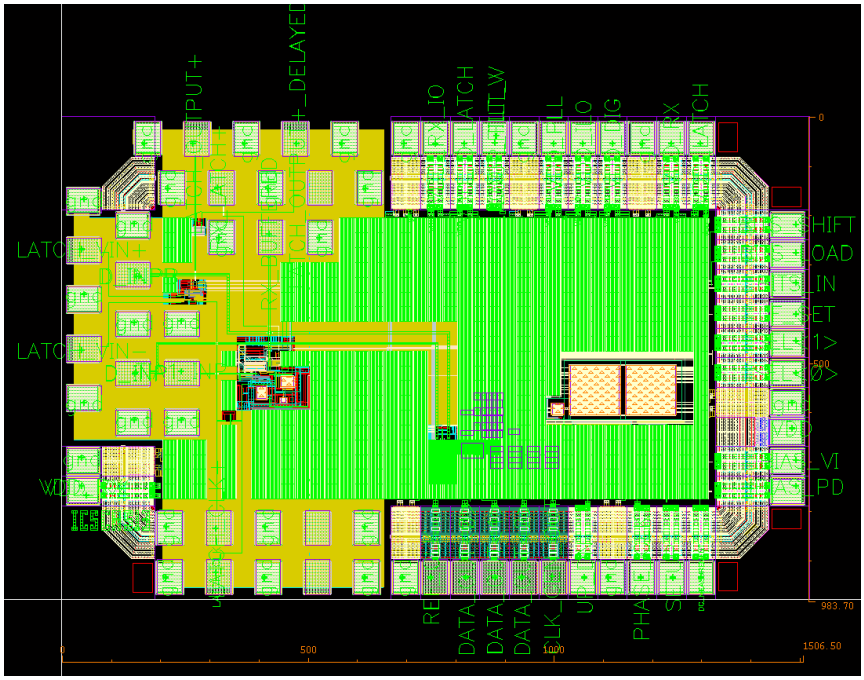

Figure 5-15: Layout of variable rate receiver IC.

Figure 5-16 is a close up view of the bottom-right corner of the test chip, which contains the PLL and clock generation circuitry. The I/O pads can also be clearly seen in this image on the bottom and right sides.

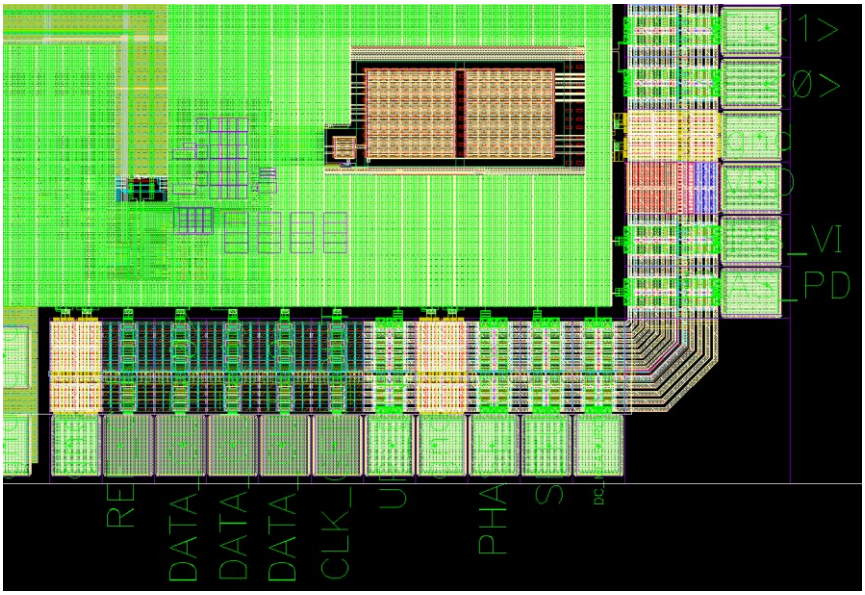

Figure 5-16: Layout of PLL with power grid.

CHAPTER 6

Conclusion

This thesis introduced a method for dynamically trading off power and noise. A new architecture was implemented that allowed sub-VCOs to be brought up to speed and connected together. This allowed for the jitter on each of the common nodes to be reduced. By introducing a timing scheme that allowed for the newly activated sub-VCOs to be brought up to speed before connection, it eliminated the problem of arbitrary phases that could cause large phase excursions. Compensation schemes were also designed in order to take care of the phase excursions that occurred upon activation of new sub-VCOs. These compensation schemes included a variable capacitance at the common nodes, as well as variable drive strength buffers. Together, the activation procedure and compensation methods, allowed us to achieve very small phase excursions during dynamic connections that were well within the jitter tolerance of a model CDR. The layout of the system in TSMC's 65nm technology was also shown. Upon receiving the fabricated chips, they will be tested and compared against theoretical predictions presented in this thesis.

Future work for this architecture may include implementing wider control of the power and noise tuning ability of the system. Increasing the number of sub-VCOs or designing binary weighted sub-VCOs would allow the range of tuning possibilities to be improved. As this work was for a proof of concept, future work may look at achieving a larger range of data rates, and incorporating the PLL into a complete receiver system.

This would include connection to latches and a TIA, along with a secondary loop with adaptive schemes that would control the implemented XOR phase detector.

References

[1] T. Toifl, et al., “A 72mW 0.03mm^{2}inductorless 40 Gb/s CDR in 65nmSOI CMOS,” *IEEE ISSCC*, pp.226 - 227, Feb. 2007.

[2] L. Rodoni, et al., “A 5.75 to 44 Gb/s quarter rate CDR with data rate selection in 90nm bulk CMOS,” *IEEE J. Solid-State Circuits*, vol. 44,no.7, pp. 1927 - 1941, Jul. 2009.

[3] M. Mansurl, et al., "A scalable 0.128-to-1 Tb/s 0.8-to-2.6 pJ/b 64-lane parallel I/O in 32 nm CMOS," *IEEE ISSCC*, pp. 402 - 404, Feb. 2013.

[4] A. Hajimiri, S. Limotyrakis and T. H. Lee, “Jitter and phase noise in ring oscillators,” *IEEE J. Solid-State Circuits*, vol. 34, no.6, pp. 790 804, Jun. 1999.

[5] M. Behbahani and G. E. R. Cowan, “Phase-Noise tuneable ring voltage-controlled oscillator in 90 nm CMOS,” *IEEE MWSCAS 2013*, pp.1031-1034, Aug. 2013.

[6] S. Kobayashi and M. Hashimoto, "A multibitrate burst-mode CDR circuit with bit-rate discrimination function from 52 to 1244 Mb/s," *IEEE Photonics Technol. Lett.*, vol. 13, no. 11, Nov. 2001.

[7] B. Razavi, *Design of Integrated Circuits for Optical Communications*, 2nd ed. New Jersey: Wiley, 2012.

[8] J. P. Hein and J. W. Scott, “z-Domain model for discrete-time PLL’s”, *IEEE Trans. Circuits Syst.*, vol. 35, no. 11, pp. 1393 1400, Nov. 1988.

[9] A. S. Sedra and K. C. Smith, *Microelectronic Circuits*, 5th ed. New York: Oxford, 2004.

[10] T. Toifl, et al., "A 0.94-ps-RMS-jitter 0.016mm^2 2.5GHz multiphase generator PLL with 360^o digitally programmable phase shift for10 Gb/s serial links," *IEEE J. Soild-State Circuits*, vol. 40, no.12, pp.2700 2712, Dec. 2005.

[11] G. Zhou, "Design of voltage-controlled oscillators based on static-CMOS inverters in 90nm process," M.S. thesis, ECE, Concordia, Montreal, QC, 2011.

[12] T. C. Carusone, D. A. Johns and K. W. Martin, *Analog Integrated Circuit Design*, 2nd ed. Wiley, 2012.

[13] C. Williams, G. E. R. Cowan and O. Liboiron-Ladouceur, "Power and noise configurable phase-locked loop using multi-oscillator feedback alignment", *IEEE MWSCAS 2013*, pp. 1023-1026, Aug. 2013.

[14] R. E. Best, *Phase-Locked Loops: Design, Simulation and Applications*, 6thed. New York: McGraw, 2007.

[15] R. C. Dorf, *Modern Control Systems*, 12thed. New Jersey: Pearson, 2011.

[16] N. Roberts. (2003, Jul. 14) *Digital noise and jitter – a primer for digital designers* [Online]. Available: http://www.eetimes.com/document.asp?doc_id=1277196