INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

ProQuest Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
800-521-0600

UMI[®]

Design Considerations for High Speed Clock and Data Recovery Circuits

by

Michel Beshara

A thesis submitted to the
Faculty of Graduate Studies and Research
in partial fulfilment of the requirements
for the degree of
Master of Applied Science

Ottawa-Carleton Institute for Electrical Engineering,
Department of Electronics,
Carleton University,
Ottawa, Ontario, Canada

April 2002

© Michel Beshara 2002

**National Library
of Canada**

**Acquisitions and
Bibliographic Services**

**395 Wellington Street
Ottawa ON K1A 0N4
Canada**

**Bibliothèque nationale
du Canada**

**Acquisitions et
services bibliographiques**

**395, rue Wellington
Ottawa ON K1A 0N4
Canada**

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-72063-2

Canada

The undersigned recommend to the Faculty of Graduate
Studies and Research acceptance of the thesis

*Design Considerations for High Speed
Clock and Data Recovery Circuits*

submitted by Michel Beshara
in partial fulfilment of the requirements for
the degree of Master of Applied Science.

Thesis Supervisor
Chairman,
Department of Electronics

Ottawa-Carleton Institute for Electrical Engineering,
Department of Electronics,
Carleton University,
Ottawa, Ontario, Canada

April 2002

Abstract

Design considerations for high speed clock and data recovery circuits are presented. Linear and bang-bang clock recovery architectures are examined and it is shown that the performance of the circuits is related to the performance of the flipflops used to build them. The design of high speed, emitter-coupled logic flipflops is investigated and two flipflops operating at speeds of up to 17GHz and 44GHz in simulation are designed in an experimental SiGe process. The design of voltage-controlled ring oscillators for use in clock recovery circuits is also investigated and an oscillator is designed and fabricated in a 0.18 μ m CMOS process. The measured centre frequency of the oscillator is 1.43GHz and the phase noise is -80.2dBc/Hz at a 1MHz offset.

Acknowledgements

Thanks to John Rogers and especially to Rony Amaya for help with measurements and lab equipment. Thanks to Malcolm Stewart for help with circuit design and layout. Thanks to Ronald Fernando for his views on all things non-engineering. Thanks to Tuan Ho for the sharp insight and for pausing to answer my questions no matter how busy things got.

Thanks to Nortel Networks for research funding and for access to tools and technology. Thanks to the Natural Sciences and Engineering Research Council for financial assistance.

Thank you to my supervisor, Calvin Plett, for his patience and for always being available to answer questions, for help with circuit design, for sending me to ISSCC in San Francisco, for the many courses, for funding, and for all the thoughtful comments and advice.

Thank you to Tina for love and encouragement and help with sea moss design.

Finally, to my family, I could write a whole other thesis about all the things I should thank you for. Mercifully, I know I don't have to. Thank you most of all.

Contents

Introduction	1
1.1 Fibre-Optic Receiver	2
1.2 Clock and Data Recovery	3
1.2.1 Open-Loop Filters	3
1.2.2 Closed-Loop Synchronizers	4
1.3 Thesis Outline	5
Clock Recovery Circuits	7
2.1 Jitter and Phase Noise	8
2.2 NRZ Data	10
2.3 PLL	13
2.3.1 PLL Overview	13
2.3.2 Frequency Response	17
2.3.3 Jitter Generation	19
2.3.4 Data-Dependent Jitter	20
2.4 Discussion	21
High Speed Clock Recovery Architectures	23
3.1 Phase Detectors for Clock Recovery Circuits	23
3.2 Clock and Data Recovery Circuit with Linear-Type PLL	24
3.2.1 The Hogge Phase Detector	25
3.2.2 Jitter Transfer Function	28
3.3 Clock and Data Recovery Circuit with Bang-Bang PLL	29
3.3.1 The Alexander Phase Detector	30
3.3.2 Jitter Transfer Function	32
3.4 Discussion	33
High Speed Flipflop Design	34
4.1 ECL Flipflop	34
4.2 Delay Components	36
4.2.1 R_L Delay Component	37
4.2.2 τ_F Delay Component	38
4.2.3 C_{DE} Delay Component	38
4.2.4 r_{bi} Delay Component	39
4.2.5 Parasitic Resistance Delay Component	40
4.3 Design and Simulation	40
4.4 Layout	45
4.5 Measured Results	46
4.6 Design Guidelines	47
4.7 Discussion	48

Ring Oscillators and Phase Noise	49
5.1 Ring Oscillator Performance Criteria.....	50
5.1.1 Tuning Range.....	50
5.1.2 Spectral Purity	50
5.1.3 Supply and Substrate Noise Rejection.....	51
5.1.4 Linearity.....	51
5.2 Ring Oscillators.....	51
5.3 VCO Phase Noise.....	53
5.3.1 Sources of Noise.....	54
5.3.2 Flicker Noise.....	56
5.3.3 Thermal Noise	57
5.3.4 Small-Dimension Effects.....	58
5.3.5 Phase Noise in Ring Oscillators	59
5.4 Discussion	59
Ring Oscillator Design	60
6.1 VCO with Delay Variation by Positive Feedback	60
6.1.1 Negative Resistance.....	61
6.1.2 Delay Stage.....	62
6.2 Design and Simulation	67
6.3 Measured Results	75
6.4 Higher Power Design	78
6.5 VCO Performance	81
6.6 Design Guidelines	82
6.7 Discussion	82
Conclusion	83
7.1 Contributions	84
7.2 Future Work.....	84
ECL Basics	86
A.1 The Emitter-Coupled Differential Pair	86
A.2 ECL and CML	89
A.3 Voltage Levels	90
References	93

List of Figures

1.1	Block diagram of a fibre-optic receiver.....	3
1.2	Clock and data recovery using an open-loop filter.....	4
1.3	Clock and data recovery using a closed-loop synchronizer	5
2.1	Phase noise appears as skirts in the frequency spectrum	9
2.2	(a) NRZ data (b) RZ data.....	10
2.3	Spectrum of NRZ data.....	11
2.4	Edge detection of NRZ data	12
2.5	Two circuits that can be used for edge detection: (a) an XOR gate with a delayed input and (b) a differentiator followed by a squaring circuit	13
2.6	A basic charge pump PLL	13
2.7	Phase detector, charge pump, and low pass filter.....	14
2.8	PLL closed-loop frequency response	18
3.1	Linear PLL-based CDR incorporating a frequency detection loop.....	24
3.2	The Hogge phase detector	25
3.3	Timing diagram for the Hogge phase detector with the clock and data optimally aligned	27
3.4	Timing diagram for the Hogge phase detector with the input data ahead of the clock	28
3.5	A bang-bang PLL	30
3.6	The Alexander phase detector	31
3.7	Sampling points (A), (B), and (C) for early, on time, and late clocks	31
4.1	An ECL latch with two pairs of emitter-follower output buffers.....	35
4.2	An ECL master-slave D flipflop.....	36
4.3	ECL circuit delay components as a function of collector current	37
4.4	Latch with minimum length input transistors and single stage output buffers.....	41
4.5	Testbench arrangement for flipflop simulation	42
4.6	Input and output waveforms for the divide by 2 circuit	42
4.7	High speed latch with two stage emitter follower output buffers	44
4.8	Latch layout	46
5.1	A ring oscillator realized using 3 inverters.....	52
5.2	A differential ring oscillator realized with 4 stages.....	53
5.3	Frequency spectrum of (a) an ideal oscillator, (b) an oscillator with phase noise, and (c) the phase noise itself	55
6.1	A negative resistance circuit.....	61
6.2	Delay stage with variable negative resistance load	62
6.3	Delay stage with differential control voltage	64
6.4	Delay stage with folded PMOS control voltage transistors.....	65
6.5	Output frequency as a function of differential control voltage	71

6.6	VCO layout.....	72
6.7	Output buffer for the oscillator.....	73
6.8	Circuit to generate bias voltages.....	74
6.9	Phase noise simulation results using post-extracted view.....	75
6.10	Measured output tone for the ring oscillator	76
6.11	Measured results of the output frequency as a function of differential control voltage	77
6.12	Phase noise simulation results for the redesigned oscillator with higher current	79
6.13	Output frequency as a function differential control voltage for the higher current VCO	80
6.14	Phase noise simulation results for a differential pair ring oscillator	81
A.1	A bipolar differential pair.....	87
A.2	Plot of I_{C1} and I_{C2} with respect to V_{id}/V_T	88
A.3	A CML inverter	90
A.4	An ECL inverter	91

List of Tables

6.1	Transistor widths and lengths for delay stage devices	71
6.2	Transistor widths and lengths for the output buffer	73
6.3	Transistor widths and lengths for the bias circuit.....	74
6.4	Summary of measured results for the 4-stage ring oscillator	78
6.5	Transistor widths and lengths for the redesigned delay stage	79

List of Acronyms

CDR	Clock and data recovery
CML	Current-mode logic
CMOS	Complementary metal oxide semiconductor
ECL	Emitter-coupled logic
NRZ	Nonreturn-to-zero
PLL	Phase-locked loop
RZ	Return-to-zero
SAW	Surface-acoustic-wave
VCO	Voltage-controlled oscillator

Chapter 1

Introduction

The recent growth of the internet has been accompanied by a steady growth in demand for bandwidth. Fibre-optic networks that support data rates of 10Gb/s have become the norm for long-haul, backbone communication and 40Gb/s capable networks are set to become the standard in the near future. For shorter distances, 2.5Gb/s networks are rapidly becoming ubiquitous. While optical fibre can support data rates in the THz range, it is the electronic circuitry at the photon/electron boundary that limits the performance of this channel of communication. To obtain the best performance at the highest speeds, engineers design with relatively expensive technologies such as gallium arsenide or indium phosphide. However, 2.5Gb/s and 10Gb/s electronic components intended for high-volume, low-cost systems become economically viable and more profitable when implemented in silicon bipolar or CMOS technology. Such components are now the subject of much research and some are already commercially available.

The subject of this thesis is the design in silicon technology of clock and data recovery (CDR) circuits for fibre-optic receivers. In implementing such circuits, engineers are confronted with many decisions ranging from the choice of system architecture to the size of transistors and resistors used to optimize performance. Because good CDR cir-

cuit design requires knowledge that spans from the system level, through the component level, and down to the device level, much of the information on the subject is scattered in the literature. Research publications do not explain why and how most design decisions are made or how to make the most appropriate decisions for future designs. In this thesis, two different clock recovery architectures are explored and it is shown that the performance of both circuits is related to the performance of the flip-flops used to build them. As well, the design of voltage-controlled oscillators for use in CDR circuits is examined and the main considerations are identified. As part of this work, several circuits were designed and fabricated in two different silicon processes. In an experimental very high speed silicon germanium (SiGe) process, flipflops suitable for use in high speed CDR circuits were designed, simulated, and fabricated. In a standard $0.18\mu\text{m}$ CMOS process, a differential 2.5GHz voltage controlled oscillator (VCO) was designed, fabricated, and measured.

1.1 Fibre-Optic Receiver

A block diagram of a fibre-optic receiver appears in Figure 1.1. The receiver consists of a photodetector at the front-end, usually a p-i-n diode or an avalanche photodetector, that converts the transmitted light signal into an electrical signal. The electrical signal is amplified by a low-noise preamplifier and then by another amplifier with automatic gain control (AGC). A clock and data recovery (CDR) circuit extracts timing information from the input data and uses it to sample the data stream. Finally, a demultiplexer converts the retimed serial data into parallel data at a lower rate where it can be digitally processed by other circuitry.

FIGURE 1.1 Block diagram of a fibre-optic receiver.

1.2 Clock and Data Recovery

When information is transmitted over a fibre-optic network, there is no separate clock or reference signal that is also transmitted. Instead, a clock recovery circuit must process the input data stream and extract timing information in order to generate a clock that is used to sample the data at the centre of the data pulse.

CDR circuits used in fibre-optic receivers can be broadly grouped into two categories: open-loop filters and closed-loop synchronizers.

1.2.1 Open-Loop Filters

Clock recovery using the open-loop filter method is illustrated in Figure 1.2, where the timing clock is directly extracted from the received signal. As will be seen in Chapter 2, data is transmitted over fibre-optic networks in nonreturn-to-zero (NRZ) format which contains a null frequency component at the bit rate. However, timing information can be extracted from the data by using a nonlinear edge detecting circuit. The extracted signal is then filtered by a high Q surface-acoustic-wave (SAW) filter or a phase-locked loop (PLL) and the filtered clock is used to sample the received data stream. The filter used must be highly selective in order to minimize the phase jitter of the clock signal. A phase shift block at the output of the filter allows for manual tuning to centre the clock edge in the bit interval.

The advantage of the open-loop technique is that it is not subject to instability or non-linear problems like frequency acquisition and cycle slipping. The disadvantage is that the one-time adjustment of the phase shift block will not track phase offsets due to temperature variations and component aging. Furthermore, when a SAW filter is used, it cannot be integrated with the rest of the receiver, which results in a more expensive and less reliable system [1].

FIGURE 1.2 Clock and data recovery using an open-loop filter.

1.2.2 Closed-Loop Synchronizers

Closed-loop synchronizers are more complex than open-loop filters but they offer several advantages that make up for the added complexity. A block diagram of a closed-loop synchronizer appears in Figure 1.3. As with the open-loop method, a timing signal at the bit rate is extracted from the data by a nonlinear edge detector. A phase detector then compares the timing signal's phase with the phase of the voltage-controlled oscillator (VCO) output and generates an error signal that drives the VCO to oscillate at the bit rate. Feedback from the decision circuit is used to align the recovered clock and the data to minimize the bit error rate. In practice, the decision circuit, edge detector, and phase detector are often implemented in one block. Examples will be discussed in Chapter 3.

FIGURE 1.3 Clock and data recovery using a closed-loop synchronizer.

Closed-loop synchronizers do not require components that are difficult to integrate so they can be used in fully integrated receivers. Furthermore, they have the advantageous property of being able to continually compensate for changes in the environment and input bit rate. Their disadvantage is that they are more difficult to design than open-loop filters as designers have to contend with issues like cycle slipping and frequency acquisition.

1.3 Thesis Outline

In this chapter, two types of clock and data recovery circuits were discussed, namely open-loop filters and closed-loop synchronizers. The remainder of this thesis deals only with closed-loop synchronizers since almost all modern CDR circuits are of this type. PLLs form the basis of most closed-loop synchronizer CDR circuits and they are discussed extensively throughout the thesis.

In Chapter 2, the properties of PLLs and the design constraints imposed by their application to clock and data recovery are considered. Jitter transfer and jitter generation are examined carefully as output clock jitter is one of the most important specifications given for CDR circuits. The format and properties of transmitted data are also examined and their effect on the design of CDR circuits is discussed.

Examples of linear and bang-bang PLL-based CDR circuits are examined in Chapter 3. The advantages and disadvantages of each architecture are considered with a detailed

examination of Hogge and Alexander phase detectors. The performance of both phase detectors is shown to depend very strongly on the speed of the flipflops and logic gates used to build them.

In Chapter 4, the design of very high speed flipflops in silicon bipolar technology is investigated. Guidelines are given based on the results drawn from circuits that were designed and fabricated in a high speed SiGe process.

In Chapter 5, the main performance criteria affecting the design of VCOs for clock recovery circuits are discussed. A general explanation of ring oscillators is given followed by a discussion of phase noise as it affects the oscillator in a CDR circuit.

The design and implementation of a CMOS VCO for operation at 2.5GHz is the subject of Chapter 6. An explanation of the VCO is presented and details of the design are followed by simulated and measured results.

Conclusions and suggestions for future work appear in Chapter 7.

Chapter 2

Clock Recovery Circuits

Data is transmitted over fibre-optic networks in non-return-to-zero (NRZ) format without a reference clock signal. At the receiving end, a clock and data recovery (CDR) circuit must generate the clock from the input data stream and then use that clock to correctly sample the data.

In most modern receivers, a phase-locked loop (PLL) based circuit is used to perform clock and data recovery. Such circuits have the advantageous property of being able to dynamically compensate for changes in the environment and input bit rate. Furthermore, PLLs are highly amenable to monolithic implementation as they do not require components which are difficult to integrate. A fully integrated clock and data recovery circuit, which may be part of a fully integrated receiver, is less costly and more reliable than a solution comprised of discrete components. This is particularly important at high frequencies where the package for each component must be separately and carefully designed often at an excessive cost and involving a lot of work. As well, a fully integrated implementation invariably consumes less power than a discrete solution as it reduces the number of large buffers needed to drive signals off-chip.

In the first part of this chapter, jitter and phase noise will be defined and discussed in relation to CDR systems. Next, the properties of NRZ data will be reviewed and con-

sideration will be given to their effect on the design of clock recovery circuits. The remainder of the chapter is devoted to a discussion of PLLs, paying particular attention to the constraints imposed by the application to clock and data recovery.

2.1 Jitter and Phase Noise

The term jitter is used to describe the time-varying offsets in the phase of the output clock measured with respect to a specified reference signal. Ideally, the clock generated by the CDR circuit should sample the incoming data exactly in the centre of the data pulse. In reality, various noise sources including thermal noise and supply and substrate noise, as well as jitter that is already present in the incoming data cause the recovered clock to wiggle about its ideal phase. Too much wiggling degrades the receiver's bit error rate so minimizing jitter is one of the most important specifications given for CDR circuits.

Output jitter can be measured in a number of ways depending on which signal is taken as the reference. Absolute jitter is measured with respect to a periodic signal having the same average frequency as the output signal. Input-to-output jitter, also called tracking jitter, is measured, as the name implies, with respect to the input signal. Finally, cycle-to-cycle jitter is measured from one output clock period to the next output clock period. Jitter measurements are often reported as RMS or peak-to-peak values.

Jitter transfer and jitter generation are frequently given as specifications in the design of CDR circuits. Jitter transfer is defined as the ratio of output jitter to input sine wave jitter. As will be seen later, some PLLs exhibit jitter peaking and can amplify input jitter at frequencies near the loop bandwidth. Clearly this is to be avoided as it is detrimental to CDR system performance. Jitter generation refers to jitter that is contributed by the CDR circuit itself. It is measured on the recovered clock with jitter-free data applied to the input.

Jitter is measured in the time domain. Its frequency domain equivalent is called phase noise and appears in the frequency spectrum as skirts around the output signal tone. This is shown in Figure 2.1. Phase noise is expressed in terms of dBc/Hz at a specified

frequency offset from the recovered clock frequency. To measure phase noise, the total noise power in a unit bandwidth at the offset frequency is calculated and divided by the signal power at the clock frequency.

FIGURE 2.1 Phase noise appears as skirts in the frequency spectrum.

There are several sources of noise that can cause jitter in a CDR system. Thermal noise generated by the devices used in the PLL can be significant at low bias currents. Supply and substrate noise generated by on-chip sources external to the PLL, including chip output drivers and digital blocks, can cause frequency and phase shifts in the recovered clock. These shifts can endure for several cycles until the noise subsides or the PLL can correct for the error. Furthermore, supply noise can modulate the source or drain voltages of switching devices, affecting threshold voltages and diffusion capacitances. Substrate noise can couple into PLL components by modulating device threshold voltages.

Off-chip filters can be used to filter out most off-chip noise but they do little to attenuate on-chip noise. On-chip filters can be effective but they are not very good at filtering low frequency noise since capacitors usually cannot be made large enough on-chip to give a useful bandwidth. To be effective in reducing jitter in PLLs, the filter cut-off frequency must be below the loop bandwidth because noise that is outside the loop bandwidth is filtered out by the loop itself. In addition to filtering, separate supply pins should be used for the PLL whenever possible to isolate its power supply. Though this

is an effective technique, isolated supplies still experience noise from coupling to other supplies through off-chip paths and coupling to the substrate through well contacts and diffusion capacitances [3]. Ultimately, a combination of techniques is necessary to minimize the output jitter in CDR systems.

2.2 NRZ Data

Nonreturn-to-zero or NRZ refers to a format where the data signal does not return to zero between adjacent pulses representing ones. This is in contrast to the return-to-zero or RZ format where the data signal always returns to zero between adjacent pulses. As an example, consider the data sequence 1, 1, 0, 1, 0, 0, 1, 1 shown in Figure 2.2 in the NRZ and RZ formats. Each bit has a duration T_b , called the bit period and is equally likely to be zero or one. The reciprocal of the bit period, $r_b = 1/T_b$, measured in bits/s, is called the bit rate.

FIGURE 2.2 (a) NRZ data (b) RZ data

Since NRZ data contains less transitions than RZ data when representing the same bit pattern, it consumes less bandwidth and is the preferred format for transmission over fibre-optic channels. Unfortunately, NRZ data has two properties that complicate the task of clock recovery. First, the data may consist of long strings of consecutive ones or

zeros. During these periods, in the absence of data transitions, the CDR circuit must 'remember' the bit rate and continue to produce an accurate clock signal with negligible frequency drift [2]. This can be accomplished with proper design of the CDR circuit and will be considered later in this chapter.

Second, the spectrum of NRZ data does not contain a component at the bit rate or at frequencies that are integer multiples of the bit rate. This can be seen by examining the power spectral density function of NRZ data which is given by [2].

$$P(\omega) = T_b \left[\frac{\sin(\omega T_b/2)}{\omega T_b/2} \right]^2 \quad (2.1)$$

The function is plotted in Figure 2.3 and has a value of 0 for $\omega = 2m\pi/T_b$ where m is any integer.

FIGURE 2.3 Spectrum of NRZ data.

Without a spectral component at the bit rate, a clock recovery circuit may lock to spurious signals or not lock at all. Therefore, it is necessary to generate a tone at the bit rate using a nonlinear operation such as edge detection. Any linear operation on the data will always contain a null at the bit rate and is ineffective in generating the required spectral component.

As illustrated in Figure 2.4, edge detection generates a pulse whenever there is a data transition from zero to one or from one to zero. The resulting signal contains twice as many transitions as the original NRZ data and consequently a frequency component at the bit rate. It should be noted that timing information is only contained in the transitions between adjacent, non-identical bits. A constant stream of ones, for example, will contain no discernible timing information.

FIGURE 2.4 Edge detection of NRZ data.

There are a number of circuits that can be used for edge detection. These include an XOR gate with a delayed input or a differentiator followed by a squaring circuit, as shown in Figure 2.5 [2]. However, for many PLL-based CDR circuits, edge detection of the input data is performed directly by a phase detector which is a component of the PLL.

FIGURE 2.5 Two circuits that can be used for edge detection: (a) an XOR gate with a delayed input and (b) a differentiator followed by a squaring circuit.

The rest of this chapter will focus on the design of PLLs for clock and data recovery. As shown above, clock recovery from NRZ data demands that some sort of edge detection be performed on the data to allow the PLL to lock to the bit rate. Furthermore, the PLL must be able to contend with the absence of transitions in the input data and still produce a steady and accurate clock signal. With this in mind, we can proceed to a more detailed discussion of clock recovery PLLs.

2.3 PLL

2.3.1 PLL Overview

A basic PLL is shown in Figure 2.6. It is a feedback system consisting of a phase detector (PD), charge pump, low pass filter (LPF), and voltage-controlled oscillator (VCO).

FIGURE 2.6 A basic charge pump PLL.

The phase detector compares the PLL input signal with a feedback signal from the VCO and produces an UP or DOWN pulse that represents the direction of the phase error between the two signals. If the input signal edge arrives before the VCO edge, the phase detector generates an UP pulse to increase the VCO frequency and make the VCO edge arrive sooner on the next cycle. If the input signal edge arrives after the VCO edge, a DOWN pulse causes the VCO frequency to decrease delaying the arrival of the next VCO edge. Ultimately, the two edges become aligned. In some phase detector implementations, the length of the pulses varies to represent the magnitude of the phase error with larger errors generating longer pulses. The pulses control a charge pump consisting of two switched current sources to sink or source current, as shown in Figure 2.7. The charge pump drives the loop filter which integrates the current and produces the VCO control signal. The effect of the control signal is to cause the VCO output to track the input signal eventually locking in phase.

FIGURE 2.7 Phase detector, charge pump, and low pass filter.

The loop is considered locked when the phase difference between the input and output signals is constant with time, a consequence of which is that the input and output frequencies are equal. When the loop is locked, the phase detector either asserts both the UP and DOWN signals simultaneously or it asserts neither signal, depending on the implementation. Furthermore, for clock recovery circuits, phase detectors are designed

such that when there are no transitions in the input data, for example during a sequence of input ones or zeroes, neither an UP or DOWN pulse is asserted or both pulses are asserted simultaneously, as in the locked condition.

PLLs can be realized without a charge pump and they often are in discrete implementations. However, for integrated circuit PLLs, charge pumps are used widely because they offer the advantage that the phase difference between the input and output signals is theoretically zero when the loop is locked. This is not so when the charge pump is omitted. In that case, the static phase difference, Φ_e , depends on the loop gain, K , and on the frequency offset, $\Delta\omega$, between the input signal and the VCO centre frequency, so that [2]

$$\Phi_e = \theta_o - \theta_i = \frac{\Delta\omega}{K} \quad (2.2)$$

The centre frequency of some integrated oscillators can vary by a factor of two at the extremes of process and temperature [4]. Furthermore, for high frequency operation, it is usually impossible to use active filters to achieve high gain. These filters require opamps or transconductors that introduce additional parasitic poles to the system adversely affecting performance and making it very difficult or impossible to stabilize the loop. For this reason, passive filters are preferred over active filters for use in high frequency clock recovery circuits. The result is a low overall loop gain and, if the charge pump is absent, a large static phase offset when the loop is locked. A large static phase offset, particularly one that is highly sensitive to the value of the input frequency and the VCO centre frequency, is very undesirable since it increases the jitter of the output signal [5]. For integrated clock recovery circuits, charge pump PLLs invariably offer better performance than those that do not include a charge pump so the remainder of this chapter will deal only with charge pump PLLs.

It is now instructive to examine loop operation in more detail. Consider the PLL of Figure 2.6 in an unlocked state with a VCO frequency that is close to but a little higher than the input frequency. Since the PLL is not locked and the first feedback edge arrives before the first corresponding input edge, the phase detector generates a

DOWN pulse that causes the VCO frequency to be reduced. If there is no input signal transition during the cycle, neither an UP or DOWN pulse is produced and the charge on the loop filter capacitor does not change. The filter component values must be properly chosen and the circuit must be carefully laid out to minimize current leakage during long periods without transitions. Any leakage reduces the VCO control voltage causing the output frequency to drift.

After several cycles, the phase error is reduced to zero but the VCO frequency is now lower than the input frequency. The feedback edges arrive later than the corresponding input edges so the loop begins to increase the VCO frequency above the input frequency to reduce the phase error. When the phase error is again zero, the VCO frequency is now higher than the input frequency. If a capacitor alone is used as the loop filter, this oscillating behaviour can continue indefinitely. However, the addition of a resistor in series with the capacitor dampens phase and frequency overshooting eventually leading to a stable locked state. The resistor converts the charge pump current into an extra control voltage that is added to the control voltage already integrated on the loop filter capacitor. The resistance value determines the amount of damping and as the resistance increases, the oscillations decay to zero at an increasing rate. For some resistor value, the loop becomes critically damped and the oscillations disappear as the phase error approaches zero without overshooting. If the resistance is increased further, the loop becomes overdamped and the phase error approaches zero rapidly but then slows down and takes a long time to reach zero.

While many feedback systems are designed to be critically damped to achieve the fastest settling response, this is usually not the case for clock recovery PLLs. Other constraints require that the PLL be overdamped and this issue will be revisited later in the chapter.

The different types of damping behaviour can be quantified by deriving the frequency response of the PLL. However, things are complicated by the fact that the switching operation of the charge pump makes the loop a discrete-time system mandating a non-linear, discrete state-variable analysis [6]. Fortunately, a linear, continuous-time

approximation can be used if only small phase excursions are assumed under locked conditions and if the loop bandwidth is a decade or more below the input frequency. A well designed loop will track the input signal with little deviation when it is locked and clock recovery PLLs are often designed with a narrow loop bandwidth so that phase error due to external noise will be small. Therefore both assumptions are reasonable for the application considered here. It should be noted that using a continuous-time approximation essentially ignores the detailed behaviour of the loop within a single cycle and examines the average behaviour over many cycles. In some cases though, per cycle behaviour can be important even for narrow bandwidths [6].

2.3.2 Frequency Response

The charge pump PLL with an RC filter has a second-order closed-loop response. The first pole is contributed by the loop filter which integrates the charge representing the phase error while the second pole is due to the VCO which integrates the output frequency forming the output phase. The transfer function shows the relationship between the input and output phases, in contrast to familiar feedback circuits where voltages or currents are usually considered. Expressed in a form commonly used in control theory, the transfer function is given by [2] [6],

$$H(s) = \frac{\theta_o(s)}{\theta_i(s)} = \frac{2\zeta\omega_n s + \omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2} \quad (2.3)$$

where,

$$\omega_n = \sqrt{\frac{K_{VCO} \cdot d \cdot I_p}{2\pi C}} \quad (2.4)$$

$$\zeta = \frac{RC}{2} \cdot \omega_n = \frac{RC}{2} \sqrt{\frac{K_{VCO} \cdot d \cdot I_p}{2\pi C}} \quad (2.5)$$

K_{VCO} is the VCO gain in rad/sec/v, I_p is the charge pump current, ω_n is the natural frequency of the system and ζ is the damping factor. Additionally, the transition density,

denoted d , represents the probability of occurrence of a transition in the input data stream. It must be factored into the equation to avoid imposing an erroneously large error current upon the PLL. For a continuously alternating input pattern, exactly the same as the clock, d adopts its maximum value of 1 [6]. However, for a pseudorandom NRZ bitstream, the value of d is 0.5 since the bitstream has half as many transitions as the clock.

Figure 2.8 shows the closed-loop frequency response of Equation 2.3 plotted for different values of ζ and for frequencies normalized to ω_n . The plot shows that the PLL transfer function is that of a low pass filter. Input phase noise at frequencies below ω_n is unattenuated whereas input phase noise at frequencies above ω_n is filtered with a slope of -20dB/decade.

FIGURE 2.8 PLL closed-loop frequency response.

The plot also shows phase noise amplification or ‘peaking’ at frequencies near ω_n . This can have a deleterious effect on the PLL’s output jitter. The problem is further compounded when data is transmitted through a chain of regenerative repeaters as the timing jitter of the recovered clock increases exponentially. It is important to note that there is a small amount of phase noise amplification at frequencies around ω_n for all values of ζ but the effect is most pronounced for values of ζ less than 0.7 [5]. Therefore, in order to control jitter accumulation, damping factors with a value between 4 and 6 usually are used in clock recovery PLLs [5] [7]. With a large value of ζ , the transfer function approaches that of a single pole low pass filter with a 3dB bandwidth given by [8]

$$F_{3dB} = \zeta \cdot \frac{\omega_n}{\pi} \quad (2.6)$$

2.3.3 Jitter Generation

Having examined the jitter transfer function, we now consider jitter generation within the loop itself. As was mentioned previously, charge pump PLLs have a static phase error that is theoretically zero when the loop is locked. However, non-idealities including mismatches and offset currents in the charge pump result in a non-zero static phase error and lead to additional output jitter. Shunt loading of the loop filter, represented by the resistor R_s in the following equation, causes a static phase error that is given by [6]

$$\theta_s = \frac{2\pi\Delta\omega}{K_{VCO}I_pR_s} \text{ rad.} \quad (2.7)$$

By making the shunt load impedance large, its effect on the static phase error can be minimized. Shunt loading of the filter usually comes from the input impedance of the VCO control terminal or from the current switches in the charge pump. Both these load impedances can be made extremely large, particularly if the circuit is fabricated in a CMOS process and the filter is loaded by CMOS gates.

Another contributor to static phase error is leakage current. The phase error resulting from a bias current I_b continuously injected into the filter node can be calculated as [6]

$$\theta_b = 2\pi \frac{I_b}{I_p} \text{ rad.} \quad (2.8)$$

Any mismatch between the current source and current sink in the charge pump causes the VCO control voltage to vary even if the input phase difference is zero. Thus the PLL locks with a finite static phase error in order to cancel the net charge deposited by the charge pump on the filter capacitor. Equation 2.8 indicates that the proportion of leakage current to charge pump current must be minimized to keep the static phase error close to zero.

Voltage ripple at the VCO control input is an effect that may contribute substantially to the output jitter. To suppress voltage ripple, a second capacitor can be connected from the output of the charge pump to ground, in parallel with the RC filter. As a rule of thumb, the second capacitor should be one tenth the size of the first capacitor [9]. The transfer function becomes third order with the introduction of a real pole at frequencies considerably higher than the resonant frequency of the loop. The third order system can be approximated by a second order system as the loop dynamics remain almost the same except for a slight decrease in the damping. Nevertheless, the PLL needs to be carefully simulated to ensure that the loop is still stable and that performance is not degraded by the addition of the new capacitor [9]. Finally, it should be noted that using a fully differential charge pump greatly reduces the magnitude of the voltage ripple [2].

2.3.4 Data-Dependent Jitter

As discussed earlier, the input data to a CDR circuit may consist of long strings of consecutive ones or zeros. During these periods, in the absence of data transitions, the CDR circuit must continue to produce an accurate clock signal with negligible frequency drift. Thus it is useful for the designer to be able to calculate the data-dependent jitter so that circuit parameters can be chosen to meet the necessary jitter requirements. The data-dependent jitter, ΔT_o , due to mismatches and offset currents in the charge pump is given by [8]

$$\Delta T_o = \frac{0.02 \cdot \delta_o \cdot N \cdot \pi \cdot F_{3dB}}{d_{min}} \quad (2.9)$$

with the units for the calculated jitter in picoseconds when F_{3db} is given in MHz. δ_o is the relative current offset at the charge pump output, N is the number of consecutive zeros or ones in the data, and d_{min} is the minimum expected transition density. Equation 2.9 indicates that as F_{3db} increases, the data-dependent jitter increases because the loop attenuates less of the input phase noise. The jitter also increases with the relative current offset at the charge pump output but decreases when there are more data transitions since more timing information becomes available.

The peak-to-peak jitter, ΔT_{ripple} , arising from the voltage ripple on the VCO control line can be expressed as [8]

$$\Delta T_{ripple} = \frac{0.01 \cdot K_{hop} \cdot \pi \cdot F_{3dB}}{d_{min}} \quad (2.10)$$

with the units for the calculated jitter in picoseconds when F_{3db} is given in MHz. K_{hop} is the attenuation due to higher order poles in the PLL filter and can include the shunt capacitor used to suppress ripple as discussed previously.

2.4 Discussion

In this chapter, the design of PLLs for clock recovery applications was examined. It was seen that minimizing jitter in the recovered clock is one of the main requirements for PLLs used in CDR systems.

To ensure that the static phase error in the recovered clock is zero, for optimum sampling of the input data, a charge pump PLL is preferred over a PLL that omits the charge pump. The charge pump should have a tristate output and minimal offset current so that the recovered clock frequency does not stray during long periods without an input data transition.

The jitter transfer function of the PLL showed that the loop exhibits jitter peaking which can amplify input jitter at frequencies around the loop bandwidth. Thus, a judi-

cious choice of loop parameters, which usually includes a large damping constant, is necessary to meet stringent jitter requirements.

The PLLs described by the equations in this chapter have a VCO control voltage that is proportional to the phase error. From Equation 2.10, it can be seen that the ripple on the control voltage depends on the input data transition density. For high operating frequencies, the analog control signals may not have time to settle during the bit period introducing high nonlinearity into the loop. Alternate clock recovery schemes based on sampling techniques result in uniform digital control pulses and are well suited to support high data rates. Digital control pulses are not proportional to the amount of phase error and result in so-called bang-bang operation in the locked state. With proper design, the increased jitter arising from bang-bang operation can be made smaller than the jitter introduced by data-dependent and nonlinear analog pulses at high frequencies. In the next chapter, practical implementations of linear PLL-based CDRs and bang-bang PLL-based CDRs will be examined. The main issues affecting the design of linear and bang-bang phase detectors for high speed operation will be identified and discussed.

Chapter 3

High Speed Clock Recovery Architectures

In this chapter, two specific architectures, namely linear PLL-based CDRs employing a Hogge phase detector and bang-bang PLL-based CDRs employing an Alexander phase detector will be examined. The two architectures are similar and differ mainly in the type of phase detector they employ. Linear-type PLLs use a phase detector that produces a continuous phase error measurement whereas bang-bang PLLs incorporate a phase detector that quantizes the error. This difference has a large impact on the dynamics of the loop and each implementation has certain advantages and disadvantages when compared to the other.

It will be seen that the critical component in the design of both the Hogge and Alexander phase detectors is the flipflop. In the next chapter, the design of emitter-coupled logic (ECL) flipflops in high speed SiGe bipolar technology will be investigated and some design guidelines will be given.

3.1 Phase Detectors for Clock Recovery Circuits

A key requirement of phase detectors used in clock recovery PLLs is that they correctly process input data that is missing transitions. In the absence of transitions, a conventional phase detector will interpret the VCO frequency as being higher than the

input frequency and drive the control voltage to correct the perceived error. For clock recovery applications, the phase detector must not generate an error signal when there is no transition in the input data. The Hogge and Alexander phase detectors both meet this requirement and neither circuit requires a separate edge detector at the input as they both inherently perform this function. Furthermore, for both phase detectors, the decision circuit is an integral component of the architecture and the recovered clock is automatically aligned to the centre of the bit interval.

To ensure good pull-in performance and lock acquisition, both Hogge and Alexander clock recovery PLLs require some type of frequency acquisition aid. A PLL-based clock recovery circuit is only capable of pulling-in a frequency error of the same order as the closed-loop bandwidth, which is typically a factor of 1000 less than the bit-rate. Therefore, without frequency acquisition aids, the VCO centre frequency has to be stable to within 0.1% over all processing and temperature variations, which is quite a stringent requirement [1].

3.2 Clock and Data Recovery Circuit with Linear-Type PLL

A CDR PLL circuit with a linear phase detector usually includes a frequency detector, as shown in the block diagram of Figure 3.1. The arrangement consists of two loops, one for frequency alignment and the other for phase alignment.

FIGURE 3.1 Linear PLL-based CDR incorporating a frequency detection loop.

A lock detector determines if the VCO output is close enough to the required frequency by comparing it to a reference signal. If not, the frequency loop is selected and remains enabled until the VCO is oscillating at the correct frequency. When the correct frequency is established, the phase detector is engaged to adjust the output phase.

A Hogge phase detector is suited for clock and data recovery applications and will be discussed in the following section. Examples of CDR circuits that employ a Hogge-type phase detector appear in [7] and [8].

3.2.1 The Hogge Phase Detector

The Hogge phase detector [10], illustrated in Figure 3.2, is a simple circuit composed of two flipflops, two XOR gates, and a summing block. In actual implementations, the necessary summations usually are performed by steering currents into and out of a capacitor with a differential charge pump so the summing block is not needed [7].

FIGURE 3.2 The Hogge phase detector.

Operation of the Hogge PD is as follows. When there is a transition in the input data, the D input and Q output of flipflop U1 become unequal so the output of XOR gate U3

goes high. It remains high until the next rising edge of the clock when the input data value is clocked through U1. At the same time, the output of XOR gate U4 goes high because the D and Q lines of flipflop U2 now become unequal. The output of U4 remains high until the next falling edge of the clock when the input data is clocked through U2 [7].

The output of XOR gate U3 is a positive pulse when there is an input data transition. Its width depends on the phase error between the input data and the clock. For a 50% duty cycle clock, the width is given by [1]

$$t_{U3} = T/2 + t_d + t_\phi \quad (3.1)$$

where T is the clock period, t_d is the delay through the flipflop, and t_ϕ is the timing error in the clock. The output of XOR gate U4 is also a positive pulse whose width is always

$$t_{U4} = T/2 + t_d \quad (3.2)$$

With the output of U4 acting as a reference, a phase error pulse of width t_ϕ is obtained by subtracting the output of U4 from the output of U3.

Figure 3.3 is a timing diagram for the phase detector with the input data and the recovered clock optimally aligned. As indicated on the diagram, the output of the phase detector has a zero average value.

FIGURE 3.3 Timing diagram for the Hogge phase detector with the clock and data optimally aligned.

A timing diagram with the input data ahead of the clock appears in Figure 3.4. In this case, the output of the phase detector has a positive average value indicating that the VCO output frequency should increase to allow the clock to catch up to the data. If the input data were behind the clock, the phase detector's output would have a negative average value.

There are two issues that must be considered carefully when designing the Hogge phase detector for high speed operation in a clock recovery circuit. First, the output has a strong ripple component that can introduce excess clock jitter. The outputs of the XOR gates can drive the UP and DOWN inputs of a charge pump directly allowing the summing block in the phase detector to be replaced by the charge pump steering currents into and out of a capacitor. The problem is that the XOR outputs do not overlap. Referring to Figure 3.3, it is clear that even when the clock is optimally aligned with the data, the charge pump receives an UP pulse followed by an equally wide DOWN pulse instead of no pulses at all or both pulses simultaneously. To minimize ripple, a shunt capacitor can be added to the loop filter as discussed in Chapter 2. Another solution is to delay the output of XOR U3 by $T/2$ so that the outputs of the two XOR gates

line up at the input of the charge pump. Any errors in the delay will result in a residual ripple causing only second order degradation of the clock phase jitter [1].

FIGURE 3.4 Timing diagram for the Hogge phase detector with the input data ahead of the clock.

The second issue involves the delay, t_d , through the flipflops. In the phase detector, the first flipflop is sampling the data signal with a clock that is nominally aligned to the centre of the bit interval. The second flipflop is sampling a retimed data waveform where the phase error of the sampling depends on the delay of the first flipflop. This means that if the first flipflop is sampling with an optimal clock phase, then the second flipflop is sampling with a clock timing error of t_d . In the extreme, if t_d is half the bit period, for example 50ps in a 10Gb/s system, the second flipflop will be sampling at the data transition rather than in the middle of the bit interval [1]. It is clear that the flipflops must be designed to minimize throughput delay.

3.2.2 Jitter Transfer Function

The equations presented in Chapter 2 are directly applicable to linear PLLs employing a Hogge phase detector. As discussed previously, the most important requirements for a clock recovery PLL are minimizing jitter transfer and jitter generation. Jitter is well

predicted by the equations and loop parameters can be selected to satisfy jitter specifications.

3.3 Clock and Data Recovery Circuit with Bang-Bang PLL

CDRs incorporating a bang-bang PLL use an Alexander-type phase detector [11] instead of the Hogge-type phase detector used with linear PLLs. Examples of bang-bang PLL-based CDRs appear in [12] and [13].

Practical implementations of the bang-bang PLL consist of two loops to control the VCO, as shown in Figure 3.5. A frequency-control loop provides coarse adjustment by setting the centre frequency, f_c , of the VCO. If the frequency does not match the input data rate, a long train of UP or DOWN pulses is generated by the phase detector. The pulse train is integrated with a charge pump and capacitor to generate the control voltage that is used to adjust the VCO frequency. Fine adjustment is provided by a phase-control loop which uses the UP and DOWN pulses directly to shift the output frequency of the VCO in discrete steps, Δf_{bb} , until the clock and data signals are optimally aligned.

The VCO used in the bang-bang PLL must have two control inputs. A typical control voltage input sets the centre frequency and a bang-bang input accepts the UP and DOWN pulses directly from the phase detector. The UP and DOWN pulses cause the VCO to oscillate at frequencies $f_c + \Delta f_{bb}$ and $f_c - \Delta f_{bb}$, respectively. To minimize jitter, a charge pump with tristate output can be used. In that case, the phase detector output includes an ENABLE signal that is only high when there has been a data transition. In the absence of a transition, the VCO maintains its current frequency. A more detailed explanation of the Alexander phase detector appears in the following section.

FIGURE 3.5 A bang-bang PLL.

3.3.1 The Alexander Phase Detector

The Alexander phase detector, shown in Figure 3.6, works by comparing the previous data value (A), the current data value (B), and a sample at the transition point (T). If the current (B) and previous (A) data values are the same, then there has been no transition and the ENABLE output remains low. When (A) and (B) are different and the transition sample (T) is the same as the previous data value (A), the clock is early. If (T) is the same as the current data value (B), then the clock is late. When the clock is on time, the sample (T) occurs exactly at the transition point and randomly equals (A) or (B), causing the phase detector to generate an UP or DOWN pulse. The result is that the VCO generated clock 'bangs' around the optimal sampling point, speeding up half of the time and slowing down the other half. Figure 3.7 shows the sampling points (A), (B), and (T) for early, on time, and late clocks [1]. In some cases, the values of (A) and (B) will be the same but the value of (T) will be different. This means that there has not been a transition but that somehow (T) got the wrong value. The occurrence of this situation is usually due to noise or poor signal integrity and is an indication of bit errors. A loss of signal output can be added to the phase detector to be asserted if too many bit errors occur within a certain time.

FIGURE 3.6 The Alexander phase detector.

The main concern in designing the Alexander phase detector is to minimize the setup and hold times of the flipflops. During normal operation, one of the flipflops is sampling data at the transition point on every clock cycle. The smaller the setup and hold region, the more accurately a decision can be made about whether the clock is early or late and the lower the generated jitter. The quality of the data transition itself is also important. The AGC amplifier that precedes the CDR circuit in the receiver must produce an output with sharp transitions between bits.

FIGURE 3.7 Sampling points (A), (B), and (C) for early, on time, and late clocks.

As with the Hogge phase detector, the Alexander phase detector is arranged so that there are two flipflops sampling data in series. The second flipflop is sampling a retimed data waveform where the phase error of the sampling depends on the delay of the first flipflop. Therefore, the delay, t_d , through the flipflops must be minimized. In

the Alexander phase detector, the situation is complicated by the fact that the flipflops have unequal loads. Referring to Figure 3.6, the first flipflop that samples *Data In* and generates the current data value (B) is loaded by another flipflop and by the inputs to two logic gates. The other flipflops in the phase detector all have smaller loads. At high frequencies, small differences in loading or in the length of interconnect can significantly misalign signals. Furthermore, adding even one gate to the load of a flipflop can appreciably increase the delay through the cell. The designer must be careful to ensure that signals needing careful alignment still are aligned after the flipflops have been loaded with other cells and with interconnect lines.

3.3.2 Jitter Transfer Function

Bang-bang PLLs are generally less demanding on the technology than linear PLLs. In the bang-bang architecture, the phase detector is, in principal, the only critical component. The loop can be designed with a wide bandwidth to filter internally generated noise so it is more tolerant of VCO phase noise. Furthermore, non-idealities in the charge pump are not as important as in the linear PLL [12].

Jitter transfer in the bang-bang PLL is nonlinear and cannot be analysed using the linear PLL theory discussed in Chapter 2. Analysis of the loop will not be presented here but the system behaviour is similar to a double integration delta modulator with prediction acting in the phase domain [12]. The jitter transfer function is close to a single pole response with the bandwidth inversely proportional to the input jitter amplitude [12]. The 3dB bandwidth of the transfer function is given by

$$F_{3dB} = \frac{k_I \cdot d \cdot \Delta f_{bb}}{\Delta T_{pp}} \quad (3.3)$$

where k_I is a proportionality constant, d is the data transition density, Δf_{bb} is the bang-bang frequency step, and ΔT_{pp} is the peak to peak jitter amplitude.

The jitter generated in the bang-bang PLL is proportional to the frequency step Δf_{bb} and to the delay in the PLL loop measured as a number, m , of clock periods required to propagate a signal from the phase detector output back to its input [12] so that,

$$\Delta T_{gpp} = k_2 \cdot m \cdot \Delta f_{bb} \quad (3.4)$$

Equation 3.4 shows that a small bang-bang frequency step and a short delay around the loop will minimize jitter generation.

For the linear PLL, jitter transfer characteristics are independent of the jitter amplitude and can be predicted according to linear PLL theory. For the bang-bang PLL, the 3dB bandwidth of the jitter transfer function is inversely proportional to the jitter amplitude making jitter prediction more complicated. However, the jitter transfer function does not exhibit jitter peaking as in the linear case and the relationship in Equation 3.3 can be used to fit jitter characteristics to those given in system performance specifications.

3.4 Discussion

In this chapter, linear and bang-bang PLLs suitable for CDR applications were investigated. The main design considerations for high frequency phase detector design were identified. It was seen that the key component in the design of both the Hogge and Alexander phase detectors is the high speed flipflop with minimal throughput delay and metastability region. In the next chapter, the design of ECL flipflops in high speed SiGe bipolar technology will be investigated and some design guidelines will be given.

In the two chapters after that, the discussion will move on to the VCO and the design and implementation of a 2.5GHz VCO will be presented. Low phase noise performance of the VCO is critical to achieving low output jitter and good overall performance in the PLL.

Chapter 4

High Speed Flipflop Design

In this chapter, the design of high speed emitter-coupled logic (ECL) flipflops in a SiGe bipolar process will be considered. High speed flipflops are the key component for good performance of both the Hogge and Alexander phase detectors discussed in the previous chapter.

The emitter-coupled differential pair is the basic building block of all ECL circuits. With this circuit at the core, a number of other high speed circuits including flipflops, logic gates, multiplexers, and demultiplexers can be constructed. This chapter will focus on the design of flipflops as these form the basis of the phase detectors used in clock recovery circuits. An introduction to ECL circuit design including an explanation of voltage levels and the difference between emitter-coupled logic and current-mode logic (CML) appears in Appendix A.

4.1 ECL Flipflop

ECL flipflops are composed of two latches arranged in a master-slave configuration. An ECL latch, shown in Figure 4.1, consists of an input differential pair, Q_1 - Q_2 , and a latching pair, Q_3 - Q_4 . Another differential pair, Q_5 - Q_6 , is connected to the emitters of the input and latching pairs and serves to steer current between them based on the value

of the input clock. When $clk+$ is high, current is steered through the input pair and the latch is in transparent mode. When $clk+$ goes low and $clk-$ goes high, the input value is latched as current is steered through the latching pair. The circuit also contains two pairs of emitter follower output buffers. A CML variant would omit these output buffers.

FIGURE 4.1 An ECL latch with two pairs of emitter-follower output buffers.

To create an edge-triggered flipflop, two latches are configured as shown in Figure 4.2. The output of the first latch is connected to the input of the second latch and the $clk+$ input of each latch is connected to the $clk-$ input of the other latch. This arrangement of clock inputs ensures that while one latch is transparent the other is always locked, resulting in edge-triggered operation.

FIGURE 4.2 An ECL master-slave D flipflop.

To obtain the fastest performance from a flipflop, transistors and resistors must be correctly sized and an appropriate bias current must be chosen.

4.2 Delay Components

In order to choose the size of transistors to optimize a circuit for speed and power, it is important to understand which parts of the circuit contribute the most delay and how to reduce the delay without expending any more power than is necessary. For digital circuits, the main concern is switching delay or how quickly the circuit can toggle between a one and a zero or vice versa.

Switching delay can be expressed as a linear combination of all the time constants of a circuit with each time constant weighed by a factor that is determined by the detailed arrangement of the circuit [14][15][16]. Thus, the switching delay can be written as

$$T_{delay} = \sum_i K_i R_i C_i + \sum_j K_j \tau_j \quad (4.1)$$

The first sum is over all the resistances and capacitances of the transistors in the circuit and also includes the load resistors. The second sum includes the transistor forward transit times. K_i and K_j are the corresponding weighing factors. For the ECL flipflop, all the transistors in the circuit are biased in the forward-active region so the reverse transit times are zero and do not appear in the equation.

A sensitivity analysis can be carried out to determine the value of the weighing factors. The magnitude of R_i and C_i is varied over the required parameter range in such a way that the contribution of the $R_i C_i$ time constant to the propagation delay is isolated and then the corresponding K factor is evaluated [15].

It is not feasible for a circuit designer to conduct a delay component analysis every time an ECL gate is to be designed. Fortunately, several such analyses appear in research literature [14][15][17], the results of which will be discussed in the following section.

In [17] the authors analysed a flipflop with 2 emitter followers and expressed the propagation delay as the sum of 102 time constants! Weighing factors were extracted for 22 dominant terms using sensitivity analysis. The time constants can be grouped together and are responsible for five delay components as shown in Figure 4.3.

FIGURE 4.3 ECL circuit delay components as a function of collector current.

4.2.1 R_L Delay Component

The load resistance delay component is due to all the RC time constants that are formed from the combination of parasitic and load capacitances with the load resistors R_L . If the load capacitance is large, its associated time constant may be the dominant contrib-

utor to this delay component. In that case, minimizing the load capacitance by using appropriately sized buffers will reduce the delay.

As illustrated in Figure 4.3, the R_L delay component is inversely proportional to the bias current. This is easy to understand in light of the fact that the output voltage swing is equal to the product of the bias voltage and the load resistance (as explained in Appendix A).

$$V_{SWING} = I_{BIAS} \cdot R_{LOAD} \quad (4.2)$$

To keep V_{SWING} constant, R_L must be reduced as I_{BIAS} is increased. Since power consumption is not the dominant specification when designing fibre-optic receivers, the ECL circuitry used in receivers usually can be operated at large currents to minimize the R_L delay component [16].

4.2.2 τ_F Delay Component

The delay component due to forward transit time is one of the dominant delay components for ECL circuits. It is constant at low collector currents but increases rapidly at higher currents when base widening occurs [16].

4.2.3 C_{DE} Delay Component

The C_{DE} or emitter diffusion capacitance delay component is proportional to the collector current at low currents but also increases rapidly at high currents when base widening occurs. This can be understood by examining the effect of base widening on the emitter diffusion capacitance.

Stored charge associated with forward biasing the emitter-base diode can be written as [16].

$$Q_{DE} = \tau_F \cdot I_{EE} \quad (4.3)$$

Although C_{DE} is a nonlinear capacitance, it is approximately equal to the stored minority-carrier charge divided by the average change in input voltage when the gate changes state (which is just the logic swing of the gate) [16],

$$C_{DE} \approx \frac{2Q_{DE}}{\Delta V} = \frac{2\tau_F \cdot I_{EE}}{\Delta V} = \frac{2\tau_F}{R_L} \quad (4.4)$$

At low currents, the C_{DE} delay component is proportional to I_{EE} since base widening is negligible and τ_F is independent of current. At high currents, τ_F increases with I_{EE} and the delay component increases in proportion to the product $\tau_F I_{EE}$.

Equation 4.4 indicates that increasing R_L will decrease the C_{DE} delay component. However this would increase the R_L delay component described in Section 4.2.1. ECL circuits normally should not be biased at currents high enough to cause base widening and at the lower currents, the R_L component is considerably larger than the C_{DE} component. Therefore, increasing R_L in order to decrease C_{DE} is not beneficial as it will increase the overall delay by increasing the R_L component.

4.2.4 r_{bi} Delay Component

The r_{bi} delay component is due to the RC time constants associated with the intrinsic base resistance of the circuit transistors. At low collector current density, r_{bi} is constant and independent of current but once base widening occurs, r_{bi} decreases rapidly as current density increases [16].

The r_{bi} delay component is small in modern high speed digital circuits. Hence there is no advantage in operating an ECL circuit at high current in order to minimize r_{bi} , since the increase in other delay components will substantially increase the overall switching delay.

4.2.5 Parasitic Resistance Delay Component

The parasitic resistance delay component is generally small in modern high speed digital circuits [16]. It is due to the RC time constants associated with the extrinsic base resistance, the collector resistance, and the emitter resistance. The parasitic resistors are, to a first order, independent of the operating current so their associated delay component is independent as well.

Having established the principal delay components that affect the speed of ECL flip-flops, we can proceed to some design examples and establish some guidelines for effective high speed design.

4.3 Design and Simulation

Two flipflop test circuits were designed in a high speed SiGe process having a transition frequency, f_T , of 70GHz. By varying transistor and resistor sizes and bias currents, operation at up to 44GHz was achieved in simulation.

As discussed in the previous section, operating a transistor with current greater than the critical current density results in a large increase in the switching delay of the device. For this reason, the transistors were all operated at or below their critical current density, which in the case of the process used here was $1.5\text{mA}/\mu\text{m}^2$. Assuming a single emitter stripe on the transistors, the current density, J_c , is equal to the emitter current divided by the emitter area and can be expressed as

$$J_c = \frac{I_E}{W_E \cdot L_E} \quad (4.5)$$

where W_E is the emitter width and L_E is the emitter length. With the emitter width fixed at $0.4\mu\text{m}$ so that only the length is variable, the optimum current, I_{opt} , becomes 0.6mA per μm of length.

The first flipflop was designed with the switching and clock input transistors set to the process minimum length of $1.2\mu\text{m}$. Transistors with a length of $8.4\mu\text{m}$ were used for the single stage emitter follower output buffers. A circuit diagram of the latch used to

make the flipflop appears in Figure 4.4. All transistor and resistor sizes as well as bias currents appear on the diagram. With a bias current of 1.5mA through the switching devices and a load resistance of 150Ω, the single-ended output voltage swing was 225mV. V_{CC} and V_{EE} were set to 0V and -5V, respectively. The power consumption of the flipflop was 120mW not including the bias circuitry. It should be mentioned that by redesigning for a 3.3V supply, the power consumption could be reduced by approximately 33%.

FIGURE 4.4 Latch with minimum length input transistors and single stage output buffers.

In order to simulate the circuit, it was configured as a divide by 2 with the positive output of the flipflop connected back to its negative input and its negative output connected back to its positive input. The testbench arrangement is shown in Figure 4.5. The input buffer is identical to the emitter followers at the output of the latch and serves to level shift the input clock by one diode voltage drop. The divide by 2 flipflop is loaded by its own inputs and by the differential pair that makes up the buffer load circuit.

FIGURE 4.5 Testbench arrangement for flipflop simulation.

All simulations were conducted with a post-layout netlist that included parasitic capacitances and resistances. Transmission line effects were not considered since the circuit did not contain any long metal tracks except from the pads to the inputs or outputs. The metal lines that connected the flipflop outputs back to the inputs were short enough that transmission line effects could be safely ignored. Cadence tools were used throughout the design including the use of SpectreRF for simulation. A 200mVpp differential sine wave was fed into the clock inputs and the outputs were probed to verify that the output frequency was at exactly half the input frequency. Representative waveforms appear in Figure 4.6. The output was also examined in the frequency domain by doing a DFT on the waveform.

FIGURE 4.6 Input and output waveforms for the divide by 2 circuit.

To determine the maximum operating frequency, the input frequency was increased until the output frequency was no longer half of the input value. Beyond the maximum operating frequency, the output spectrum did not include a component at half the input frequency but instead displayed several different output tones that resulted from the mixing of signals in the flipflop.

The first flipflop functioned correctly with a maximum input frequency of 17GHz. The small internal voltage swing helped achieve high speed operation as did the use of small load resistors. However, keeping the input transistors small meant that a low bias current had to be used so that the critical current density would not be exceeded. At high speeds, the decoupling ability of the emitter follower output buffers is greatly reduced so large output buffers had to be used. The combination of low current through the switching devices and emitter followers that are seven times larger than the input pair resulted in heavy loading of the internal nodes and limited the operating frequency.

To improve performance in the next design, the bias currents were increased with a commensurate increase in the size of the switching devices. A second emitter follower was also introduced to provide better decoupling between the internal switching nodes and the external load. Using two emitter followers meant that the first one could be made small to minimize loading of the switching transistors while the second one could be made large to provide good decoupling and drive at higher frequencies [18].

For the second flipflop, the switching and clock input transistors had a length of $5\mu\text{m}$. Transistors with a length of $5\mu\text{m}$ were also used for the first stage of the emitter follower output buffers. The second stage of the buffers used transistors with a length of $18\mu\text{m}$. A circuit diagram of the latch used to make the flipflop appears in Figure 4.7. All transistor and resistor sizes as well as bias currents appear on the diagram. With a bias current of 5.3mA through the switching devices and a load resistance of 50Ω , the single-ended output voltage swing was 265mV. V_{CC} and V_{EE} were set to 0V and -5V, respectively. The power consumption of the flipflop was 240mW not including the bias circuitry, which was double the power of the first flipflop.

The flipflop functioned correctly with an input frequency as high as 40GHz. Furthermore, when 5nH inductors were added between the load resistors and V_{CC} , the maximum frequency increased to 44GHz. The 5nH value is reasonable for a bondwire inductance and provides additional gain which results in a higher operating frequency. However, other parasitics that would result from packaging the circuit were not considered and would likely lower the maximum operating frequency.

FIGURE 4.7 High speed latch with two stage emitter follower output buffers.

Comparing the two flipflop designs, it is clear that the addition of a second output buffer stage and an increase in the switching transistor bias current results in a significant increase in speed. While power consumption is not a dominant specification for clock recovery circuits, the possibility of overheating, and the price associated with keeping the circuitry at the correct operating temperature must be considered. For this reason, if it is sufficient to use a single output buffer and lower current, then this should be done.

As discussed in Section 4.2, when transistors are operating with a bias current less than the critical current density, the load resistance delay component is dominant and should be reduced to increase switching speed. In the second flipflop, this was achieved in two ways. First, the value of the load resistors was reduced to 50Ω from the 150Ω of the first design and the bias current was increased to achieve the desired voltage swing. Second, a two stage output buffer was used allowing the first stage to be made small so as not to excessively load the switching devices. Simulation showed that for the process used here, setting the load resistance to a value between 50Ω and 150Ω is optimum. Also, increasing the size of the first stage emitter followers beyond the size of the switching transistors did not provide any appreciable benefit in speed and slowed down the circuit if it was made too large.

4.4 Layout

A layout of the latch of Figure 4.4 appears in Figure 4.8. The dimensions are approximately $40\mu\text{m}$ by $40\mu\text{m}$. Two important considerations for an effective layout are to keep the lines in the signal path short and to maintain symmetry between differential signals. Bipolar layouts typically are less dense than CMOS layouts because of minimum spacing requirements between devices. In the layout of Figure 4.8, with minimum spacing between the transistors, lines in the signal path are kept as short as possible. The differential signal path from input to output is perfectly symmetrical except for the use of a second metal layer when it was necessary to cross over the first metal layer. Pairs of resistors that have the same value are oriented the same way in order to improve matching. Additionally, a large number of ground contacts is used to minimize resistance to ground which may result in common mode noise. Although not shown in this layout, on-chip decoupling of the supply voltage can help reduce power supply noise.

FIGURE 4.8 Latch layout.

4.5 Measured Results

The first flipflop of Figure 4.4 was fabricated in an experimental high speed SiGe process provided by Nortel Networks. The flipflop was used in the Alexander phase detector circuit discussed in the previous chapter. Unfortunately, when the circuit was probed in the lab, it did not function at all. After powering up the circuit and providing an input signal, the output was examined with a high speed spectrum analyzer but no output signal was present. When the design was checked in simulation and the layout verified again, no errors were found that could account for the circuit not working. However, while the circuit was being designed, the models used for the transistors changed several times and it is believed that the final models used were not as accurate

as expected. To the author's knowledge, not a single circuit fabricated in the same run of this process, including VCOs and amplifiers designed by other engineers, has been shown to work.

For economic reasons, the research group that originally sponsored the use of this technology for designing the flipflops switched to a different process. With access to the technology becoming very limited and the additional complication that Nortel Networks sold its fabrication facility to another company, it became impossible to continue investigating the fabricated circuits or to proceed with a redesign.

4.6 Design Guidelines

The design of high speed ECL flipflops is an iterative process but some guidelines are given here to facilitate the exercise.

1. Transistors should not be operated with current greater than the critical current density. Beyond the critical current density, speed performance decreases while power consumption continues to increase.
2. For a modern process, the optimum load resistance is somewhere between 50Ω and 150Ω .
3. After the bias current is chosen, the current switch transistors should be set to the smallest size that does not make them exceed the critical current density.
4. If speed requirements can be satisfied with a single stage emitter follower output buffer then that configuration should be used to lower power consumption. Flipflops with a single stage buffer can usually be designed for a lower power supply than circuits with two stage buffers.
5. If pushing the technology to obtain the highest speeds, two emitter followers are required.
6. Setting the first emitter follower to the same size as the current switch transistors is a good starting point as this minimizes the load on the current switch transistors. The first emitter follower is only weakly loaded by the second stage so it does not need to be very large.
7. The second emitter follower should be made three or four times bigger than the first buffer stage. This will permit the circuit to handle larger loads at high speeds and will provide a lower output impedance.

4.7 Discussion

In this chapter, the design of high speed ECL flipflops was investigated. The dominant delay components were identified and two flipflops were designed. A lower power flipflop with a single-stage emitter-follower output buffer operated at speeds of up to 17GHz in simulation. A higher power version with a two-stage output buffer operated at speeds of up to 44GHz in simulation.

In the next chapter, the discussion moves on to the VCO, an important block that greatly affects the overall performance of a PLL clock recovery circuit.

Chapter 5

Ring Oscillators and Phase Noise

A voltage-controlled oscillator (VCO) is used to generate the recovered clock in a PLL-based CDR circuit. PLL dynamics force the VCO to oscillate at the input bit rate and align its phase for optimum sampling of the input data. Since the VCO generates the recovered clock, proper design is critical for achieving low output jitter and good overall performance of the CDR system.

A type of VCO often used in clock recovery PLLs is the ring oscillator. Ring oscillators are popular because they are easy to integrate and do not require any special components like varactors or on-chip inductors. Although they have poor phase noise characteristics, their output spectrum can be cleaned up enough by including them in a PLL to make them viable for use in CDR applications. A well designed ring oscillator can generate frequencies over a wide range with good linear control over frequency.

In this chapter, the performance criteria affecting the design of VCOs used in CDR circuits will be discussed. A general explanation of ring oscillators will be given followed by a detailed discussion of phase noise in this type of VCO.

In the next chapter, the design and implementation of a CMOS VCO for 2.5GHz operation will be presented.

5.1 Ring Oscillator Performance Criteria

The VCO used in a clock recovery circuit must produce a stable, periodic signal which is used to sample the input data stream. The most important specifications affecting the design of the oscillator are tuning range, spectral purity, supply and substrate noise rejection, and linearity. As in any other design, the specifications are not independent of each other and trade-offs are required.

5.1.1 Tuning Range

The tuning range is the range between the minimum and maximum values of the VCO frequency. It must be chosen to accommodate the input bit rate of the CDR system as well as process and temperature induced variations. Furthermore, it is important that variation of the output amplitude and jitter be minimized over the tuning range [2]. For clock recovery circuits, the bit rate is fixed at a precisely controlled value so variations in the VCO centre frequency are the main consideration. For example, a VCO designed to have a free-running frequency of 2.5GHz may oscillate at 2.9GHz when fabricated. This frequency would further vary with temperature and age. If the VCO's tuning range is less than $\pm 400\text{MHz}$ around the centre frequency, it would not be suitable for use in a 2.5GHz system.

5.1.2 Spectral Purity

Phase noise at the VCO output is caused mainly by the electronic noise of the active devices in the oscillator and by supply noise. In PLLs, VCO phase noise components lying close to the loop centre frequency are suppressed by negative feedback in the loop. This occurs because the phase noise lies within the loop filter bandwidth allowing the error signal at the output of the low pass filter to correct for it. For phase noise components lying outside the loop filter bandwidth, no error signal is available to correct output phase variations. To minimize the VCO's contribution to the overall phase noise, it is desirable to maximize the loop bandwidth. However, for clock recovery circuits, a narrow loop bandwidth is preferred as discussed in Chapter 2. Therefore, the VCO should be designed to keep phase noise at a minimum.

5.1.3 Supply and Substrate Noise Rejection

The supply and substrate noise sensitivity of the VCO should be as small as possible particularly if the oscillator is integrated along with digital circuits. Clock recovery circuits built in CMOS may share the same substrate and package as a digital processor further increasing the need for good noise rejection. As a rule, having fully differential signal and control paths greatly helps in achieving high rejection of supply and substrate noise [2].

5.1.4 Linearity

A linear frequency control characteristic is generally desired in applications like FM demodulation where the VCO periodically needs to adjust to frequency steps at the PLL input. However, for timing recovery circuits, nonlinearity of several tens of percent can be tolerated [2] since a timing recovery PLL is tracking the average frequency of the input signal and the frequency deviations of the output are small when the loop is locked. Nevertheless, VCO nonlinearity must be considered carefully because variation of K_{VCO} over the tuning range can degrade loop stability and lead to loop dynamics that vary over the operating frequency.

5.2 Ring Oscillators

There are several types of VCOs including *LC* tank circuits, relaxation oscillators, and ring oscillators. The remainder of this chapter will deal only with ring oscillators as this is the type of oscillator implemented in this work. More information about the other types of oscillators is available in several textbooks [4] [19] and in research publications.

A ring oscillator is realized by placing an odd number of inverting amplifiers or delay elements in a loop so that the output of the last element is connected to the input of the first element, as illustrated in Figure 5.1.

FIGURE 5.1 A ring oscillator realized using 3 inverters.

The odd number of inversions around the loop provides 180° of phase shift. If the system exhibits an additional frequency dependent phase shift of 180° , or 60° per inverter for the 3 inverter example, while maintaining a gain of 1 or more around the loop, then oscillation can occur. For example, referring to Figure 5.1, if the output of the last inverter changes to a 0, the change will propagate through all three inverters in a time $T/2$, at which point the last inverter's output will change to a 1. After another time of $T/2$, the last inverter's output will change back to a 0. If the gain around the loop is less than 1, this oscillation will die out. However, with a gain of 1 or more, the circuit will continue to oscillate. For a ring oscillator consisting of n inverters each with a delay T_D , we can write,

$$\frac{T}{2} = nT_D \quad (5.1)$$

Then,

$$f_{osc} = \frac{1}{T} = \frac{1}{2nT_D} \quad (5.2)$$

Equation 5.2 shows that the frequency of oscillation can be controlled by varying the effective number of inverters in the ring or by varying the delay through each inverter. If the latter method is employed, the delay through each inverter can be controlled by varying either the capacitance or the resistance seen at the output node of each stage. In general, resistive tuning provides a larger and relatively more uniform frequency variation than capacitive tuning and it lends itself more easily to differential control [2].

Since ring oscillators require an odd number of inversions in a loop, they can also be realized with an even number of stages if differential inverters or delay elements are used and one of the stages is connected so that it does not invert. This way the total number of inversions around the loop is still odd as illustrated in Figure 5.2.

FIGURE 5.2 A differential ring oscillator realized with 4 stages.

In addition to improved supply and substrate noise rejection, differential ring oscillators realized with an even number of stages can provide outputs from the middle inverters that have a quadrature phase relationship with the outputs from the last inverters [1]. This can be very useful in many communication applications and some clock recovery circuits.

5.3 VCO Phase Noise

Accurate prediction of ring oscillator phase noise relies heavily on computer simulation. Nevertheless, a good understanding of the nature of phase noise is necessary for effective circuit design.

For the purposes of this discussion, an oscillator can be thought of as a system that converts voltage and current noise into phase noise. Therefore, it is necessary to identify the dominant sources of voltage and current noise in the circuit and then to understand how this noise is transformed into phase noise. The effects of some types of noise can be attenuated by proper design of the VCO and by proper design of the PLL if the VCO is part of such a circuit.

Traditional theories of oscillator phase noise have assumed that the process by which voltage and current noise are transformed into phase noise is linear and time-invariant

[20]. More thorough analysis [21] reveals that the process can in fact be considered linear but that it is time-varying within each period of oscillation. Linearity is a reasonable assumption as long as the noise superimposed on the main signal remains much smaller than the signal, which is true of any properly designed oscillator. To demonstrate that the process is time-varying, consider a sinusoidal oscillator that is disturbed by an impulse of current. If the impulse is applied exactly at a voltage maximum, the amplitude of the signal increases but the timing of the zero crossing remains unchanged. The response to the impulse superposes exactly with the oscillations. An impulse injected at some other time affects both the amplitude and the timing of the zero crossings. The greatest effect occurs if the impulse is injected exactly during a zero crossing. A mechanical analogue is that of a swinging pendulum. If the pendulum is given a small push at the top of its arc of vibration, at the point when it stops and is changing direction, its phase will remain undisturbed. However, if it is disturbed at any other time during its period, the phase will be altered. The disturbance has the greatest effect on the phase if it occurs as the pendulum is crossing its equilibrium point.

For an electrical oscillator, the disturbances that can affect the phase come from a variety of sources. These will be identified in the following section.

5.3.1 Sources of Noise

To identify the sources of phase noise in an oscillator, it is instructive to examine the phase noise spectrum. Consider first the frequency spectrum of an ideal sinusoidal oscillator as shown in Figure 5.3(a). Without phase noise, all the power of the oscillator is focused at the output frequency, f_c . Non-sinusoidal oscillators have additional energy present at harmonics of the centre frequency but for simplicity only sinusoidal oscillators will be considered here. Figure 5.3(b) shows the frequency spectrum in the presence of phase noise which causes some of the power to spread to neighbouring frequencies.

Phase noise is usually specified in dBc/Hz measured at a given frequency offset. The dBc unit indicates that the measurement is reported in dBs relative to the centre frequency. The phase noise is obtained by taking the ratio of power in a 1 Hz bandwidth at

the offset to the total power of the oscillator. Referring to Figure 5.3(b), this is the ratio of the area of the rectangle with a 1Hz bandwidth at offset f_m , to the total area under the power spectrum curve [22]. In most publications, phase noise measurements are reported at offsets of 100kHz or 1MHz.

FIGURE 5.3 Frequency spectrum of (a) an ideal oscillator, (b) an oscillator with phase noise, and (c) the phase noise itself.

The spectrum of the phase noise itself can also be shown, as in Figure 5.3(c). For an oscillator signal given by

$$X(t) = A \cos(2\pi f_o t + \theta(t)) \quad (5.3)$$

the spectrum in Figure 5.3(b) is that of the signal $X(t)$, whereas the spectrum in Figure 5.3(c) is that of the phase angle term, $\theta(t)$. The phase noise spectrum is shown on a log-log scale and exhibits three distinct regions where the spectral density of the phase fluctuations falls as $1/f^3$, $1/f^2$, and $1/f^0$ [22]. Each of these regions represents a different noise process involved in generating the phase fluctuations. The $1/f^3$ region is caused by the flicker noise in the devices that make up the VCO. Device thermal noise, on the other hand, is responsible for the $1/f^2$ region. At some large frequency offset, the spectrum flattens out to a minimum noise floor. This is due to the noise associated with active elements, such as buffers, present in the circuit or it may reflect the limitations on the instruments used to measure the spectrum [21]

5.3.2 Flicker Noise

Flicker noise is found in all active devices and is a significant noise source in MOS transistors. It is thought to occur when the carriers that constitute DC current are captured and released by traps located in the semiconductor material. This results in random fluctuations in the number of carriers in the channel of a MOS device. The characteristic times involved in the capture and release process cover a wide range but when large numbers of such events are considered, it can be shown that a power spectral density nearly proportional to the inverse of frequency results [23]. For this reason, flicker noise is also called ' $1/f$ ' noise.

The power spectral density of an equivalent flicker noise voltage source appearing in series with the gate of a MOS transistor is given by [23]

$$S_{vf}(f) = \frac{K(V_{GS})}{C'_{ox}} \frac{1}{WL} \frac{1}{f} \quad (5.4)$$

where $K(V_{GS})$ is a bias-dependent quantity that increases approximately linearly with $|V_{GS} - V_T|$ [23], C'_{ox} is the oxide capacitance per unit area, and W and L are the width and the length of the device. To determine the noise voltage's effect on the drain current, the voltage equation can be multiplied by the transconductance, g_m , to yield an equivalent drain noise current with power spectral density

$$S_{if}(f) = g_m^2 S_{vf}(f) \quad (5.5)$$

Equation 5.4 shows that flicker noise increases when the WL product is small. In a differential ring oscillator, W and L can usually be increased in the current source transistors without an appreciable effect on the frequency of oscillation. However, the input differential transistors must be kept small for high speed operation since the input of each stage serves as the load for the previous stage.

5.3.3 Thermal Noise

Thermal noise, also called white noise because of its flat power spectral density up to extremely high frequencies, results from the random thermal motion of carriers in the channel of a MOS device. The power spectral density of an equivalent thermal noise voltage source appearing in series with the gate of a MOS transistor in the saturation region is given by [23]

$$S_{vw} = \frac{\frac{8}{3}kT}{\mu C_{ox} \frac{W}{L} (V_{GS} - V_T)} \quad (5.6)$$

where k is Boltzmann's constant, T is the absolute temperature, μ is the mobility, and C_{ox} is the oxide capacitance. Equation 5.6 shows that thermal noise in a MOS transistor is inversely proportional to the gate-source bias voltage above threshold and to the width over length ratio of the device. Not surprisingly, the thermal noise induced cycle-to-cycle jitter of a CMOS ring oscillator is also inversely proportional to the gate-source bias voltage above threshold. For a ring oscillator composed of NMOS differen-

tial pairs with triode region PMOS loads, the thermal noise induced jitter is approximately given by [22]

$$\overline{\Delta t_{vco}}^2 \cong \frac{kT}{I_{BIAS}} \cdot \frac{a_v \xi^2}{2(V_{GS} - V_T)} \cdot T_o \quad (5.7)$$

where ξ is the noise contribution factor, a proportionality constant that is related to bias conditions and increases with shrinking channel length. The equation applies to any length oscillator since the effect of the number of stages is accounted for by T_o .

5.3.4 Small-Dimension Effects

Short-channel devices can be significantly noisier than long-channel devices. High electric fields present near the drain of short-channel transistors produce ‘hot carriers’ for sufficiently large V_{DS} values. This effect can be modelled by an equivalent ‘carrier temperature’ that is higher than the lattice temperature and causes an increase in thermal noise. The noise current power spectral density can be as much as 2 to 5 times higher in short channel devices than it is in long-channel devices [23].

Furthermore, hot carriers are responsible for the drain-to-substrate current I_{DB} , which exhibits noise of its own. I_{DB} produces a voltage drop across the substrate resistance and its noise modulates the effective substrate potential. Fluctuations in the substrate voltage are coupled to the drain current through the substrate transconductance g_{mb} , contributing to an increase in the observed drain current noise [23].

There is little a circuit designer can do to mitigate the increasing noise contribution of shrinking devices. Thus it is becoming more difficult to design low jitter, high frequency oscillators as CMOS transistors are becoming smaller. Equation 5.7 indicates that it is preferable to keep $V_{GS} - V_T$ large to minimize thermal noise induced jitter. However, if $V_{GS} - V_T$ is increased, V_{DS} must be increased as well to keep devices in saturation. The increase in short-channel noise that accompanies an increase in V_{DS} may outweigh the benefits accrued by increasing $V_{GS} - V_T$.

5.3.5 Phase Noise in Ring Oscillators

Having identified the dominant sources of device noise, we now discuss the effects of this noise on the output phase of ring oscillators.

Referring to Figure 5.3(c), the $1/f^3$ region is caused by flicker noise upconverted to the region near the centre frequency. The $1/f^3$ phase noise corner can be reduced by proper design of the VCO and is not the same as the $1/f$ device noise corner. For example, symmetry of the rising and falling edges of the oscillator output signal leads to minimum upconversion of $1/f$ noise thus reducing the $1/f^3$ phase noise corner [24]. Using linear loads such as resistors also reduces flicker noise upconversion as well as substrate and supply coupling [24].

The $1/f^2$ region of Figure 5.3(c) is caused by device thermal noise. In general, increasing the number of devices in the oscillator increases the thermal noise and degrades the phase noise performance. Furthermore, if symmetry criteria are met, minimizing the upconversion of $1/f$ noise, then the jitter and phase noise of the oscillator are dominated by white noise [24].

With regards to supply and substrate noise, differential oscillators show better rejection than single-ended configurations as well as lower noise injection into other circuits on the same chip [2].

5.4 Discussion

In this chapter, the performance criteria affecting the design of VCOs used in CDR circuits was discussed. A general explanation of ring oscillators was given followed by a discussion of phase noise in VCOs. Sources of device noise were identified and their effect on oscillator phase noise was explained.

In the next chapter, the design and implementation of a differential CMOS VCO for operation at 2.5GHz will be presented. Details of the design are followed by simulated and measured results.

Chapter 6

Ring Oscillator Design

The design and implementation of a VCO for operation at 2.5GHz will be presented in this chapter. The VCO was fabricated in TSMC's 0.18 μ m single poly, six metal, 1.8V CMOS process and is compatible with the style of current-steering logic frequently used to build high speed clock recovery circuits. The chapter begins with an explanation of the VCO design and then presents simulated and measured results.

6.1 VCO with Delay Variation by Positive Feedback

The VCO is a 4-stage ring oscillator with differential outputs and differential control lines and is based on a design suggested by Razavi in [2]. It uses resistive tuning to control the delay through each stage and consequently the frequency of oscillation. The differential output nodes of every stage are connected to a pair of resistors in parallel with a circuit that acts as a variable negative resistance. The negative resistance is used to vary the overall resistance seen at the output nodes allowing control of the delay through each stage.

6.1.1 Negative Resistance

To see how a negative resistance can be implemented, consider the circuit of Figure 6.1, a pair of cross-coupled transistors sharing the same source but with the gate of each transistor connected to the drain of the other.

FIGURE 6.1 A negative resistance circuit.

With this arrangement, V_{GS1} , the gate to source voltage of transistor M_1 , is the same as V_{DS2} , the drain to source voltage of transistor M_2 , and V_{GS2} is the same as V_{DS1} . When V_{GS1} is increased by a small amount Δv_{gs} and, to maintain differential operation, V_{GS2} is decreased by the same amount, then the current in the left branch increases by an amount Δi_d while the current in the right branch decreases by the same amount. To a first order approximation, Δi_d is given by

$$\Delta i_d = g_m \cdot \Delta v_{gs} \quad (6.1)$$

Therefore, an increase in V_{GS1} which is the same as V_{DS2} , corresponds to a decrease in current through transistor M_2 and we can write,

$$\begin{aligned} -\Delta i_d &= g_m \cdot \Delta v_{ds2} \\ \Delta v_{ds2} &= -\frac{1}{g_m} \cdot \Delta i_d \end{aligned} \quad (6.2)$$

Equation 6.2 shows that transistor M_2 acts as a negative resistance with a value of $-1/g_m$. The differential circuit consisting of both transistors acts as a negative resistance with a value of $-2/g_m$.

6.1.2 Delay Stage

A variable delay stage for the ring oscillator can be formed by connecting the negative resistance circuit of Figure 6.1 in parallel with the load resistors, R_1 and R_2 , of a differential inverter as shown in Figure 6.2.

FIGURE 6.2 Delay stage with variable negative resistance load.

With R_1 and R_2 set to the same value and the cross-coupled transistors M_3 and M_4 set to the same size, the equivalent load resistance of the delay stage becomes

$$R_1 \parallel -1/g_{m3} = R_1 / (1 - g_{m3}R_1) \quad (6.3)$$

As the bias current, I_{B2} , increases so does the transconductance, g_{m3} , of each of the cross-coupled transistors. The small-signal differential resistance becomes less negative and the equivalent output resistance of the delay stage increases, lowering the fre-

quency of oscillation. As long as the absolute value of the negative resistance remains larger than the resistance R_I so that

$$1/g_{m3} > R_I \quad (6.4)$$

the product $g_{m3}R_I$ in Equation 6.3 will be less than 1 and the equivalent load resistance will be positive. When the above condition is met, the minimum effective resistance is R_I , occurring when g_{m3} is 0 or equivalently when there is no current flowing through the cross-coupled pair. The maximum value of the effective resistance depends on g_{m3} but can approach that of an open circuit if the negative resistance is set to the same value as the real resistance so that they cancel each other. In that case, the maximum effective value is a function of the amount of mismatch between the positive and negative resistances and of any parasitic resistances.

A problem with the circuit of Figure 6.2 is that the output voltage swing is not constant across the tuning range and depends on the value of the bias current I_{B2} . As I_{B2} varies, so do the currents steered by M_3 and M_4 to R_I and R_2 . To fix the problem, a differential pair can be used to steer current between M_1 - M_2 and M_3 - M_4 , as shown in Figure 6.3.

With this arrangement, the total current flowing through R_I and R_2 is fixed and equal to I_B . The control voltage does not set the value of the bias current as it did in the circuit of Figure 6.2 but rather determines what proportion of the total current flows through the input transistors and what proportion flows through the negative resistance. In addition to ensuring that the output voltage is constant across the tuning range, the new circuit benefits from improved noise immunity because it uses differential control lines. Furthermore, using only current-steering circuitry keeps the current drawn from the power supply constant and eliminates the voltage drops that are usually associated with a lot of switching activity.

FIGURE 6.3 Delay stage with differential control voltage.

To arrive at the circuit implemented in this work, the delay stage requires one more modification so that the control path uses a folded PMOS arrangement. This results in less stacked transistors and allows the circuit to operate with a lower supply voltage and a larger output swing. The modified delay stage appears in Figure 6.4.

To understand the operation of the delay stage, we first establish that M_5 and M_6 are constant current sinks. When the differential control voltage is zero so that M_7 and M_8 are equally sharing the current from the PMOS source M_9 , then the current through M_5 consists of the current drawn from the input pair plus half the current from M_9 . Similarly, the current through M_6 consists of the current drawn from the cross-coupled pair plus the other half of the M_9 current. Referring to Figure 6.4 and summing the currents at nodes A and B, we can write,

$$I_{M_5} = I_1 + I_{M_9} \quad (6.5)$$

and

$$I_{M_6} = I_2 + I_{M_8} \quad (6.6)$$

where I_1 and I_2 are the currents through the input differential pair and the cross-coupled pair respectively.

FIGURE 6.4 Delay stage with folded PMOS control voltage transistors.

The PMOS differential pair serves to steer the current through the input pair or through the cross-coupled transistors. If M_7 is fully on and M_8 is fully off, then all the current sourced by M_9 flows through M_7 and into M_5 . The current drawn through the input differential pair is lowered by an amount equal to the current provided by M_9 so that the current through M_5 remains constant. Since M_8 is off, the current drawn through the cross-coupled pair is increased by an amount equal to the current provided by M_9 in order to keep a constant current through M_6 . When the control voltage is changed so that M_7 is fully off and M_8 is on then there is an increase in current through the input pair and a decrease through the cross-coupled pair. The sum of currents flowing

through the load resistors R_1 and R_2 is constant and equal to the sum of the currents I_1 and I_2 .

The size of transistors M_7 and M_8 directly affects the sensitivity of the VCO. If a large differential control voltage is required to steer all the current from the source M_9 then the sensitivity will be lower than if all the current is steered with only a small differential voltage. If the steered current is small and the transistors have a long channel, then their transconductance will be minimized and a large differential voltage will be required for complete switching. Making the transistors longer also increases their input capacitance which may be favourable if the VCO is used in a PLL. In that case, the control voltage to the VCO comes from a charge pump and a low pass filter so a large input capacitance to the control transistors adds to the filter capacitance making it easier to achieve the desired value on-chip.

To ensure that the circuit always oscillates regardless of the control voltage, it is necessary to make the current from the PMOS source M_9 less than the constant current into the NMOS sink M_5 by an amount that will always provide enough gain to the differential pair. The maximum current that can be steered from one branch to the other is equal to the current provided by M_9 . If that amount of current is steered away from the differential pair, there still needs to be enough current left over so that the differential pair has enough gain to oscillate. Since the voltage gain of a differential pair is approximately equal to the product of the transconductance of the input transistors with the load impedance [9], we can write

$$a_v = g_{m1}(R_1 \parallel -1/g_{m3}) = g_{m1}(R_1/(1 - g_{m3}R_1)) \quad (6.7)$$

Recalling that the transconductance can be expressed as [9]

$$g_m = \sqrt{2\mu_n C_{ox} \frac{W}{L} I_D} \quad (6.8)$$

where μ_n is the electron mobility, C_{ox} is the oxide capacitance, W and L are the width and length of the transistor, and I_D is the drain current, we can approximate the minimum gain when all the current from M_9 is steered away from the input pair as

$$a_{v, min} = \sqrt{\mu_n C_{ox} \frac{W_1}{L_1} (I_{M_5} - I_{M_6})} \cdot \frac{R_1}{1 - \sqrt{\mu_n C_{ox} \frac{W_3}{L_3} (I_{M_6} + I_{M_9})} \cdot R_1} \quad (6.9)$$

The equation does not take into account load capacitance and parasitics that will further lower the gain at the operating frequency. As well, the square-law model for transistors is not very accurate for short-channel devices like the ones used in this design. Nevertheless, the equation served as a starting point in choosing relative sizes and currents before using a circuit simulator to set the final component values.

Having established the architecture and design considerations that affect the delay stage, we may proceed with the design of the ring oscillator and the choice of component sizes.

6.2 Design and Simulation

As discussed in Chapter 5, the frequency of oscillation for a ring oscillator is given by the reciprocal of two times the number of stages multiplied by the delay through each stage. Since the output of each stage is loaded by the following input, the delay is approximately equal to the RC time constant formed by each stage's load resistance and the next stage's input capacitance so that

$$f_{osc} = \frac{1}{2nT_D} \cong \frac{1}{2nR_L C_{IN}} = \frac{1}{2nR_1 C_{IN}/(1 - g_{m3}R_1)} \quad (6.10)$$

For the oscillator implemented in this work, resistors were used directly to implement R_1 and R_2 rather than PMOS transistors biased in the triode region. For high frequency operation, resistors are advantageous because they contribute less parasitic capacitance than PMOS transistors. This is particularly important here because the delay stage already has a relatively high input capacitance because of the addition of the cross-cou-

pled pair. Furthermore, since any nonlinearities in the load translate into additional phase noise [21], PMOS loads must be carefully biased to ensure that they remain deep in the triode region during operation. This can be achieved using a replica biasing circuit as discussed in [22] but would further complicate the design which is not desirable at high frequencies.

The input capacitance of each stage consists mainly of the gate-to-source capacitances of the input differential pair and of the cross-coupled pair plus any capacitance due to interconnect and wiring. The gate-to-source capacitance, C_{gs} , for a MOS transistor is approximately $(2/3)WLC_{ox}$ when the device is in the saturation region of operation. The gate-to-drain capacitance, C_{gd} , is not considered since it is small relative to C_{gs} . Gate-to-drain capacitance becomes important when a transistor is being used in a circuit with high voltage gain. In that case, the Miller effect causes the capacitance to be multiplied by the gain. However, for a ring oscillator the voltage gain of each stage is one, or a little greater than one in practice, so the effect of the gate-to-drain capacitance remains negligible. What is important to note for the delay cell used in this work is that the input capacitance is nearly double that of a simple differential pair delay stage. With the input capacitance as described above, a reasonable approximation for the frequency of oscillation is

$$f_{osc} \cong \frac{1}{\frac{4}{3}n \frac{R_1 C_{ox}(W_1 L_1 + W_3 L_3)}{(1 - g_{m3} R_1)}} = \frac{3(1 - g_{m3} R_1)}{4n R_1 C_{ox}(W_1 L_1 + W_3 L_3)} \quad (6.11)$$

To achieve a high operating frequency, it is necessary to keep the widths and lengths of the transistors small to minimize capacitance. Furthermore, simulation results indicated that for the circuit to oscillate at 2.5GHz, it would require a small number of stages. Since it is very difficult to obtain the gain and phase shift necessary for oscillation with a 2-stage design and 3-stage design does not provide quadrature outputs, an oscillator comprised of 4 stages was chosen. The even number of stages allowed for quadrature outputs which is useful in some clock recovery architectures. Although 0.18 μ m technology is relatively fast and capable of operating in the gigahertz range,

the delay stage used presents enough input capacitance to the previous stage to preclude the use of a large number of stages for high frequency operation.

The output voltage swing of each stage is set by the product of the R_I resistance value and the sum of the currents through the input pair and the cross-coupled pair. This sum is equal to the total current through the NMOS current sinks minus the current provided by the PMOS source. Thus, the voltage swing, V_{sw} , can be expressed as

$$V_{sw} = (I_{M5} + I_{M6} - I_{M9})R_I \quad (6.12)$$

A lower voltage swing results in faster operation but a higher voltage swing gives better noise margin. An important consideration in choosing the swing is keeping the input transistors in saturation. Level shifting buffers are not used between stages so the output of one stage directly drives the following stage. If the output voltage swing is too large it will drive the transistor receiving the high input voltage out of saturation. This can be seen by examining the voltages required to keep the device in saturation. The requirement is

$$V_{DS} \geq V_{GS} - V_T \quad (6.13)$$

In terms of voltage swing, for a transistor receiving the high input voltage, this can be rewritten as

$$(V_{DD} - V_{sw}) - V_S \geq V_{DD} - V_S - V_T \quad (6.14)$$

which simplifies to

$$V_{sw} \leq V_T \quad (6.15)$$

This is not a stringent rule but it should not be disregarded completely. If the voltage swing is too large, the input devices change regions of operation during switching which may result in an unusually shaped output waveform.

With the considerations outlined above, a voltage swing of 0.8V was chosen. The input devices have a threshold voltage of approximately 0.7V, higher than the nominal value of about 0.4V because of the body effect, so the choice of a 800mV swing was not

much greater than the threshold value. The currents were set to $200\mu\text{A}$ for I_{M5} , $200\mu\text{A}$ for I_{M6} and $50\mu\text{A}$ for I_{M9} . Using Equation 6.12, R_1 and R_2 were each set to $2.3\text{ K}\Omega$ in order to obtain a voltage swing of 800mV .

The final sizes of the transistors were determined through simulation. As explained earlier, the differential PMOS transistors were made longer than minimum length to minimize the transconductance and accommodate a large control voltage.

Sizing the input pair and the cross-coupled pair involved a trade-off between the range of output frequencies and the value of the centre frequency. Increasing the size of the cross-coupled transistors increases their transconductance and consequently the range for the load resistance as shown in Equation 6.3. Increasing the range of the variable load resistance means an increase in the output voltage range. The input pair was kept larger than the cross-coupled pair to ensure that the delay stage could be switched despite the latching action of the cross-coupled transistors. Thus any increase in the size of the cross-coupled pair was accompanied by a corresponding increase in the size of the input pair and a decrease in the output frequency because of the larger input capacitance.

To choose an appropriate size for the NMOS pairs, simulation was performed with Cadence's SpectreRF simulator using the slow/slow and fast/fast NMOS and PMOS models. With the control voltage set to 1V , the output frequency was at the low-end of the range which needed to be less than 2.5GHz . The differential control voltage was then set to -1V and the output frequency was checked to see if it was greater than 2.5GHz . If the frequency was too high the transistors were made wider and if it was too low they were made smaller. It was necessary that 2.5GHz fall within the tuning range regardless of the transistor models being used. Figure 6.5 is a graph of the output frequency as a function of the control voltage for the final circuit including output buffers and bias circuitry. The results were obtained from post-extracted simulation. The sensitivity, K_{VCO} , is equal to the slope of the graph and is approximately 320MHz/V in the case of typical model transistors. Simulation also confirmed that the circuit still oscillated when the maximum current was steered away from the input transistors.

FIGURE 6.5 Output frequency as a function of differential control voltage.

The current sink transistors M_5 and M_6 and the current source M_9 were sized to provide the required current with V_{B1} set to 700mV and V_{B2} set to 1.1V. The final device sizes for the transistors in the delay stage appear in Table 6.1.

Device	Width (μm)	Length (μm)
M1	0.8	0.18
M2	0.8	0.18
M3	0.6	0.18
M4	0.6	0.18
M5	8.0	0.36
M6	8.0	0.36
M7	5.0	1.0
M8	5.0	1.0
M9	9.0	0.36

TABLE 6.1. Transistor widths and lengths for delay stage devices.

The VCO layout consisting of four identical differential delay cells appears in Figure 6.6. The dimensions of the layout are approximately $60\mu\text{m}$ by $44\mu\text{m}$.

FIGURE 6.6 VCO layout.

In addition to the ring oscillator, a source follower output buffer and an open drain output buffer were designed and connected to the output of the oscillator to permit measurements to be taken using 50Ω probes. The source follower was connected to the output of the ring oscillator to isolate it from the larger open drain buffer and the pads. Although a source follower does not provide any gain, it does have a lower input capacitance than a differential pair amplifier which, if it were connected directly to the oscillator, would severely impact the output frequency. The buffer circuitry is shown in Figure 6.7 and the transistor sizes are shown in Table 6.2. The simulated current through each source follower was approximately $350\mu\text{A}$ and the current through the bias transistor M_{16} of the open drain buffer was approximately 2.6mA when a 50Ω load was connected. The $200\mu\text{m}$ transistors, M_{14} - M_{16} were layed out as 20 parallel devices each with a width of $10\mu\text{m}$.

FIGURE 6.7 Output buffer for the oscillator.

As well, a bias circuit was designed to provide bias voltages for the constant current sources in the delay cells. The bias circuit appears in Figure 6.8 and the transistor sizes are shown in Table 6.3. The output voltages are 700mV for V_{B1} and 1.1V for V_{B2} .

Device	Width (μm)	Length (μm)
M10	8.0	0.18
M11	8.0	0.18
M12	20.0	0.36
M13	20.0	0.36
M14	200.0	0.18
M15	200.0	0.18
M16	200.0	0.36

TABLE 6.2. Transistor widths and lengths for the output buffer.

FIGURE 6.8 Circuit to generate bias voltages.

Device	Width (μm)	Length (μm)
M17	32.0	0.36
M18	1.2	0.36
M19	4.5	0.36

TABLE 6.3. Transistor widths and lengths for the bias circuit.

Phase noise simulation was performed using SpectreRF and the results appear in Figure 6.9. The phase noise of the oscillator was very high even by the poor standards of ring oscillators. Simulation was also performed for oscillators with a higher number of stages operating at lower frequencies to determine if there was an improvement in the phase noise. The 4-stage oscillator had a centre frequency of approximately 2.4GHz and phase noise of -79dBc/Hz at a 1MHz offset. Typically, published phase

noise values for CMOS ring oscillators operating in the low gigahertz range are about -95dBc/Hz to -105dBc/Hz at a 1MHz offset [24].

A 12-stage oscillator with a centre frequency of 800MHz had a phase noise value of -89dBc/Hz at a 1MHz offset. A 20-stage oscillator had a centre frequency of 480MHz and phase noise of -94dBc/Hz at a 1MHz offset. For comparison, published phase noise values for oscillators in the 500MHz range are typically less than -105dBc/Hz at a 1MHz offset [24].

FIGURE 6.9 Phase noise simulation results using post-extracted view.

6.3 Measured Results

The fabricated 4-stage oscillator was probed under a microscope and the output was measured using a Hewlett Packard HP8593E Spectrum Analyser. With a 1.8V supply, the measured centre frequency was 1.43GHz, about 1GHz less than the simulated value. The output tone as it appeared on the spectrum analyser is shown in Figure 6.10.

A phase noise measurement was obtained by recording the power at the centre frequency and comparing it to the output power at a 1MHz offset. The measured phase noise was -80.2dBc/Hz which was very close to the simulated value of -79dBc/Hz. The measured supply current was 3.83mA, which was about 1.1mA less than the expected value. When the power supply was lowered by 10%, the centre frequency decreased to 1.41GHz. When the supply was raised by 10%, the centre frequency increased to 1.46GHz. Thus, the power supply sensitivity was calculated to be 139MHz/V.

FIGURE 6.10 Measured output tone for the ring oscillator.

The VCO sensitivity was measured by varying the differential control voltage and recording the output frequency. A graph of the output frequency as a function of the differential control voltage appears in Figure 6.11. The sensitivity of the fabricated circuit was 140MHz/V, less than half of the 320MHz/V value obtained in simulation. A summary of the measured results appears in Table 6.4.

FIGURE 6.11 Measured results of the output frequency as a function of differential control voltage.

Some differences between simulated and measured results are normal when circuits are fabricated but the discrepancy here was larger than expected. However, it is not uncommon for the centre frequency of some integrated oscillators to vary by a factor of two after fabrication [4]. Very often, oscillators are fabricated two or three times before the correct centre frequency is achieved. Because the centre frequency is closely tied to the analog characteristics of the transistors, any imperfections in the modelling of the device input capacitances can have a drastic effect. In the case of on-chip resistors, models are notoriously bad and the resistances can vary by as much as 30% from the simulated values when the devices are fabricated. The low supply current value and lower centre frequency both indicate that the fabricated resistors were in fact larger than the simulated value. The bias circuit output depends on the value of the resistor in the circuit with the current decreasing as the resistance increases. Also, larger load resistors in the delay stages increase the delay and lower the centre frequency. Resimulating the circuit with the resistor values increased by 30% and using slow/slow models for the transistors, a centre frequency of 1.7GHz was obtained. This value was closer to

the measured results than the original simulation that had assumed the resistor values would be exact.

	f_c	Sensitivity	Supply Current	Phase Noise
simulated	2.4GHz	320MHz/V	4.9mA	-79dBc/Hz
measured	1.43GHz	140MHz/V	3.83mA	-80.2dBc/Hz

TABLE 6.4. Summary of measured results for the 4-stage ring oscillator.

6.4 Higher Power Design

In order to investigate whether the phase noise performance could be improved by increasing the current through the delay stages, a higher current version of the oscillator was designed. Simulation with SpectreRF indicated that the dominant source of noise in the oscillator delay stages was thermal noise from the load resistors. Each resistor generated nearly 6% of the total noise in the first oscillator. Each transistor in the cross-coupled pair generated about 2.7% of the total noise. The current sink transistors in each stage were responsible for 2.8% of the total noise.

To decrease the size of the load resistors, the bias currents I_{M5} and I_{M6} were each increased to 540 μ A. I_{M9} was also increased to 80 μ A. Using Equation 6.12, R_1 and R_2 were each set to 800 Ω in order to maintain the same voltage swing of 800mV. The transistors were sized in the same manner as described previously. The sizes appear in Table 6.5.

Phase noise simulation of the higher current oscillator showed an improvement over the lower power design. As shown in Figure 6.12, the phase noise value measured at a 1MHz offset was -89dBc/Hz. Further simulation with higher currents showed only negligible improvement in the phase noise performance.

Device	Width (μm)	Length (μm)
M1	2.5	0.18
M2	2.5	0.18
M3	2.0	0.18
M4	2.0	0.18
M5	22	0.36
M6	22	0.36
M7	12.0	1.0
M8	12.0	1.0
M9	40.0	0.36

TABLE 6.5. Transistor widths and lengths for the redesigned delay stage.

FIGURE 6.12 Phase noise simulation results for the redesigned oscillator with higher current.

Simulation results of the higher current oscillator to determine the sensitivity and tuning range appear in Figure 6.13. The sensitivity was about 600MHz/V and the oscillator exhibited an increased range compared to the lower power design. This was expected as the increased current and larger size of the cross-coupled transistors meant a larger transconductance and consequently a larger tuning range.

FIGURE 6.13 Output frequency as a function differential control voltage for the higher current VCO.

As a baseline reference, a simple differential pair ring oscillator was designed and simulated to determine the phase noise. With a bias current of 1mA and a load resistance of 800Ω , the oscillator had a voltage swing of 800mV like the previous delay stages. The input transistors had a width of $4.5\mu\text{m}$ and a length of $0.18\mu\text{m}$. The oscillator required 12 stages in order to oscillate at 2.7GHz. As indicated in Figure 6.14, the oscillator exhibited phase noise of -107dBc/Hz at a 1MHz offset. Additional circuitry to provide frequency control and maintain a constant output voltage swing over the tuning range would necessarily result in a degradation of the phase noise performance.

FIGURE 6.14 Phase noise simulation results for a differential pair ring oscillator.

6.5 VCO Performance

Simulation with SpectreRF showed that the load resistors are the dominant source of noise in the delay stages. However, simulation of the redesigned oscillator indicated that beyond a certain limit, increasing the current in order to decrease the load resistance had only a negligible effect. In addition to noise from the resistors, each delay stage has three current sources whose noise performance directly affects the output frequency. Referring to Figure 6.4, the flicker noise of current sink M_5 is upconverted to phase noise around the output frequency by the input differential pair. The cross-coupled pair upconverts the noise from transistor M_6 , and both pairs upconvert the noise from M_9 . Any noise on M_9 causes the current steering mechanism to be noisy. Most published differential CMOS ring oscillators use only a single tail-current sink. Thus, from the outset, the tail-current noise in those circuits is three times less than in the case of the oscillator designed in this work.

6.6 Design Guidelines

Based on the oscillators designed in this chapter, the following guidelines are given here to facilitate the exercise of choosing a ring oscillator delay stage suitable for CDR applications.

1. If load resistors are used in the delay stage, their contribution to the overall noise should be carefully examined. Appropriate sizing of the resistors can result in a significant improvement of noise performance.
2. Delay stage architectures that use multiple current sinks and sources should be avoided. Architectures with a single tail-current transistor will generally have lower phase noise.
3. A large number of transistors means more thermal noise and consequently more phase noise. Thus a simple delay stage is preferred over a complex one. Any added complexity should be carefully examined to see if the proposed benefits outweigh the degradation in phase noise.
4. For high frequency operation, it is best to avoid delay stages with more than one differential pair at the input because of the increase in input capacitance.

6.7 Discussion

A differential VCO using variable negative resistance for frequency control was designed and fabricated. The measured centre frequency was 1.43GHz. The VCO showed very poor phase noise performance of -80.2dBc/Hz measured at a 1MHz offset. The major contributors of noise in the circuit were identified and a redesigned VCO with higher current showed an improvement in the phase noise achieving a simulated value of -89dBc/Hz at a 1MHz offset. However, the phase noise performance of the redesigned oscillator was still not as good as that of a published circuit having similar specifications. The 4-stage ring oscillator appearing in [24] used a simpler delay stage architecture with a single tail-current transistor and achieved a phase noise of -95.2dBc/Hz at a 1MHz offset. The centre frequency was 2.81GHz and it used approximately the same net supply current as the redesigned oscillator from this work.

Because of its poor phase noise performance at the required frequencies, the VCO architecture considered in this chapter is not the best choice for use in high frequency clock recovery circuits. Alternative architectures with similar supply current and output frequency appearing in publications show better phase noise performance.

Chapter 7

Conclusion

In this thesis, design considerations for clock recovery PLLs were identified and linear and bang-bang PLLs were discussed and compared. Specifically, Hogge and Alexander phase detectors were examined and the main design issues were identified. The performance of both phase detectors was shown to be directly related to the performance of the flipflops used to build them. The design of high speed ECL flipflops was considered and two different flipflops were designed in an experimental SiGe process. The first, lower power flipflop was simulated at speeds of up to 17GHz and the second, higher power flipflop at speeds of up to 44GHz.

The design of VCOs for CDR applications was also examined and a differential VCO using variable negative resistance for frequency control was designed and fabricated in a 0.18 μ m, CMOS process. The VCO was shown to be inappropriate for high frequency clock recovery circuits because of its poor phase noise performance at the required frequencies. The dominant sources of noise in the VCO were identified and a second, higher current oscillator was designed. The second circuit showed an improvement in phase noise but still did not achieve good performance at high frequencies when compared to published circuits with similar supply current and output frequency.

7.1 Contributions

The subject of this thesis was the design in silicon technology of clock and data recovery circuits for fibre-optic receivers. The main contributions were:

1. A comprehensive overview of clock and data recovery circuit design spanning from the system level to the device level.
2. An examination of the main design issues concerning the Hogge and Alexander phase detectors.
3. An explanation, including guidelines, of how to design high speed ECL flipflops.
4. An examination of the main design issues concerning ring oscillator VCOs for use in CDR circuits, including guidelines for identifying a delay stage architecture with good phase noise performance.
5. The detailed analysis and design of a ring oscillator VCO suggested in the literature and an evaluation of its suitability for use as a component in CDR circuits.

7.2 Future Work

There is much future work that can be done in the area of CDR circuit design. Using 0.13 μm CMOS to build CDR circuits at 20GHz or 40GHz may be possible but will be challenging because of increasing device noise. Using SiGe to build circuits operating at speeds beyond 40GHz is also an area that can be explored.

As CMOS and bipolar technologies continue to shrink, designs at higher speeds become possible. However, with a decrease in size comes an increase in noise. Analog performance becomes more difficult to predict and, as was the case in this work, oscillators can show a large discrepancy between the simulated value of the centre frequency and the final result after fabrication. Thus it becomes beneficial to make circuits more 'digital'. In the case of ring oscillators, it may be better to build the ring with a large number of simple stages rather than a small number of complex stages. To centre the oscillation frequency around the desired value, some of the stages could be swapped in or out. At high frequencies, the swapping of an entire stage may be too coarse an adjustment so instead resistors or capacitors at the output of each stage could be enabled or disabled. In addition to more predictable operation, the design cycle with such an arrangement would be reduced because less layout would be necessary. With short-channel devices, the difference between schematic and post-extracted simulation

becomes more pronounced as the effect of interconnect loading increases. Multiple layouts are sometimes required before satisfactory post-extracted simulation results are achieved. This can be avoided with a design that employs swapping of delay elements. After the basic layout is completed, the required number of delay elements can be easily connected based on simulation results.

Appendix A

ECL Basics

In this appendix, a basic introduction to emitter-coupled logic (ECL) circuit design is presented. ECL is one of the fastest logic families available and it is used extensively in the design of high speed clock recovery circuits.

All ECL circuits, including flipflops and logic gates, are based on the emitter-coupled differential pair. This review will therefore begin with the differential pair.

A.1 The Emitter-Coupled Differential Pair

An emitter-coupled differential pair is shown in Figure A.1. Depending on the value of the input voltage, V_{id} , all of the bias current, I_{EE} , can be steered through either transistor Q_1 or Q_2 .

To analyse the circuit's large signal behaviour, we begin by writing an equation for the sum of voltages around the loop formed by V_{id} and by V_{be1} and V_{be2} , the base-emitter voltages of transistors Q_1 and Q_2 ,

$$V_{id} = V_{be1} - V_{be2} \quad (\text{A.1})$$

FIGURE A.1 A bipolar differential pair.

Recalling the exponential relationship for a bipolar transistor,

$$I_C = I_S \cdot e^{(V_{be}/V_T)} \quad (\text{A.2})$$

and taking the ratio of I_{C1} / I_{C2} with the assumption that I_{S1} and I_{S2} are equal, we obtain.

$$\frac{I_{C1}}{I_{C2}} = \frac{I_{S1} \cdot e^{(V_{be1}/V_T)}}{I_{S2} \cdot e^{(V_{be2}/V_T)}} = e^{(V_{id}/V_T)} \quad (\text{A.3})$$

In addition, we can also write.

$$\alpha I_{EE} = I_{C1} + I_{C2} \quad (\text{A.4})$$

where α is the common-base current gain.

Combining Equation A.3 and Equation A.4, the following equations for I_{C1} and I_{C2} are obtained.

$$I_{C1} = \frac{\alpha I_{EE}}{1 + e^{(-V_{id}/V_T)}} \quad (\text{A.5})$$

$$I_{C2} = \frac{\alpha I_{EE}}{1 + e^{(V_{id}/V_T)}} \quad (\text{A.6})$$

A plot of I_{C1} and I_{C2} with respect to V_{id}/V_T appears in Figure A.2. It shows that when the differential input voltage is 0, the bias current is evenly split between both branches of the circuit. However, when V_{id} is equal to about $4V_T$, approximately 100mV, all the current is steered through one transistor or the other. It is this property of being able to switch the current almost entirely from one branch to the other with such a small differential voltage that makes ECL circuits so fast.

FIGURE A.2 Plot of I_{C1} and I_{C2} with respect to V_{id}/V_T .

Examination of the differential pair indicates that the current drawn from the power supply is constant and equal to I_{EE} . Therefore, the differential pair consumes a constant amount of static power equal to $I_{EE}(V_{CC} - V_{EE})$. Because of the power consumption and the associated difficulties and cost of excess heat removal, ECL cannot be used for the design of very complex VLSI circuits. However, ECL is well suited for the design of high speed building blocks in fibre-optic receivers. These blocks have a relatively low gate count compared to complex VLSI circuits and power consumption, while

important, is not a predominant design specification as it is in the case of portable devices.

Figure A.2 indicates that a differential voltage of 100mV is enough to switch the current almost completely from one branch of the differential pair to the other. To margin for temperature and process variations, modern ECL circuits generally are operated with voltage swings between 300mV to 500mV. The low logic swing reduces current spikes caused by charging and discharging load capacitances. Further, the fact that the current drawn from the supply is constant and independent of the switching operation reduces switching noise on the supply lines. As a result, ECL circuits can operate reliably even with small noise margins.

Unlike conventional CMOS logic, ECL never actually switches off but rather steers current from one branch of a differential pair to the other. For this reason, ECL is also called current-steering logic.

A.2 ECL and CML

It is useful to distinguish between current-mode logic (CML) and emitter-coupled logic (ECL). The terms are sometimes used interchangeably but traditionally ECL has been used to denote circuits that include emitter-followers to act as output buffers whereas CML omits these buffers and allows the switching transistors to directly drive the following stage. The distinction between the two styles is important because the addition of output buffers affects the output voltage level and, in both styles of logic, correct biasing of every node is critical not only to high speed operation but to the circuit working at all.

As an example consider Figure A.3, a CML inverter. This is a CML circuit because it is composed only of switching transistors. An ECL inverter, shown in Figure A.4, is similar except for the addition of one stage of emitter-followers at the output. ECL circuits are also often designed with two emitter-followers at the output.

By omitting power hungry output buffers, CML circuits can be used to implement complicated gates realizing greater functionality for a given power dissipation compared to ECL circuits. However, the emitter-follower output structure of ECL gates makes it possible to drive large capacitive loads with a small delay. It also supports a large fanout which would otherwise be impossible due to the finite input impedance of the ECL gate.

FIGURE A.3 A CML inverter.

CML gates therefore should be used for on-chip logic functions when the fanout is small and the interconnect length is short. Since CML gates will be driving only smaller on-chip loads, they can be biased at lower current levels without appreciably degrading speed and they are useful when power dissipation and cell area are important issues. However, for very high speed operation, ECL wins the race hands down.

A.3 Voltage Levels

Returning to Figure A.3, assume that the output voltage swing is 400mV and that V_{CC} is 0V and V_{EE} is -5V. CML and ECL circuits are often biased with V_{CC} at 0V and V_{EE}

connected to a negative supply because the output is referenced to V_{CC} and it is easier to provide a clean ground than a clean positive supply.

Under these biasing conditions, 0V corresponds to logic 1 and -400mV corresponds to logic 0. If 0V (logic 1) is applied to V_{IN+} and -400mV (logic 0) is applied to V_{IN-} , all the bias current flows through the left branch pulling V_{OUT+} to -400mV and V_{OUT-} to 0V, effectively performing the inverse function.

FIGURE A.4 An ECL inverter.

The output voltage swing is determined by the product of the bias current with the load resistance so that,

$$V_{SWING} = I_{BIAS} \cdot R_{LOAD} \quad (A.7)$$

For example, to get a voltage swing of 400mV with a 1mA bias current requires a 400Ω load resistance.

Now consider the ECL inverter of Figure A.4 under the same biasing conditions and assuming a base-emitter voltage drop, V_{be} , of 0.8V. The voltages at the collectors of the switching transistors are still 0V and -400mV just like the CML inverter. However, the

addition of the output buffers causes another V_{be} or 0.8V drop so that the output voltages are now -0.8V for logic 0 and -1.2V for logic 1.

Voltage levels must be carefully considered when choosing bias currents and resistor values to insure that none of the transistors in an ECL circuit enter the saturation region. In normal operation, the devices should be in the forward active region so that switching occurs as quickly as possible.

References

- [1] A. Buchwald and K. Martin, *Integrated Fiber-Optic Receivers*. Massachusetts: Kluwer Academic Publishers, 1995.
- [2] B. Razavi, "Design of monolithic phase-locked loops and clock recovery circuits - a tutorial," in *Monolithic Phase-Locked Loops and Clock Recovery Circuits*, B. Razavi, Ed. New York: IEEE Press, 1996.
- [3] J. G. Maneatis, "Design of high-speed CMOS PLLs and DLLs," in *Design of High-Performance Microprocessor Circuits*, A. Chandrakasan, W. J. Bowhill, and F. Fox, Ed. New York: IEEE Press, 2001.
- [4] B. Razavi, *Design of Analog CMOS Integrated Circuits*. New York: McGraw Hill, 2001.
- [5] J. R. Long, "High frequency integrated circuit design in BiCMOS for monolithic timing recovery," Master's thesis, Carleton University, Ottawa, 1992.
- [6] F. M. Gardner, "Charge-pump phase-lock loops," *IEEE Trans. Comm.*, vol. 28, pp. 1849-1858, Nov. 1980.
- [7] T. H. Lee and J. F. Bulzacchelli, "A 155-MHz clock recovery delay- and phase-locked loop," *IEEE J. Solid-State Circuits*, vol. 27, pp. 1736-1746, Dec. 1992.

- [8] Y. M. Greshishchev and P. Schvan, "SiGe clock and data recovery IC with linear-type PLL for 10-Gb/s SONET application," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1353-1359, Sep. 2000.
- [9] D. A. Johns and K. Martin, *Analog Integrated Circuit Design*. New York: John Wiley & Sons Inc., 1997.
- [10] C. R. Hogge, "A self-correcting clock recovery circuit," *J. Lightwave Technol.*, vol. LT-3, pp. 1312-1314, Dec. 1985.
- [11] J. D. H. Alexander, "Clock recovery from random binary signals," *Electron. Lett.*, vol. 11, pp. 541-542, Oct. 1975.
- [12] Y. M. Greshishchev et al., "A fully integrated SiGe receiver IC for 10-Gb/s data rate," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1949-1956, Dec. 2000.
- [13] M. Meghelli et al., "SiGe BiCMOS 3.3-V clock and data recovery circuits for 10-Gb/s serial transmission systems," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1992-1995, Dec. 2000.
- [14] D. D. Tang and P. M. Solomon, "Bipolar Transistor Design for Optimized Power-Delay Logic Circuits," *IEEE J. Solid-State Circuits*, vol. 14, pp. 679-684, Aug. 1979.
- [15] E.-F. Chor, A. Brunnschweiler, and P. Ashburn, "A Propagation-Delay Expression and its Application to the Optimization of Polysilicon Emitter ECL Processes," *IEEE J. Solid-State Circuits*, vol. 23, pp. 251-259, Feb. 1988.
- [16] Y. Taur and T. H. Ning, *Fundamentals of Modern VLSI Devices*. New York: Cambridge University Press, 1998.
- [17] K. Ishii, H. Ichino, and C. Yamaguchi, "Maximum Operating Frequency in Si Bipolar Master-Slave Toggle Flip-Flop Circuit," *IEEE J. Solid-State Circuits*, vol. 29, pp. 754-760, July 1994.
- [18] H.-M. Rein, "Design Considerations for Very-High-Speed Si-Bipolar IC's Operating up to 50 Gb/s," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1076-1090, Aug. 1996.
- [19] T. H. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*. New York: Cambridge University Press, 1998.

- [20] B. Razavi, "A Study of Phase Noise in CMOS Oscillators," *IEEE J. Solid-State Circuits*, vol. 31, pp. 331-343, Mar. 1996.
- [21] T. H. Lee and A. Hajimiri, "Oscillator Phase Noise: A Tutorial," *IEEE J. Solid-State Circuits*, vol. 35, pp. 326-336, Mar. 2000.
- [22] T. C. Weigandt, *Low-Phase-Noise, Low-Timing-Jitter Design Techniques for Delay Cell Based VCOs and Frequency Synthesizers*. Ph.D. Thesis, University of California, Berkeley, 1998.
- [23] Y. Tsividis, *Operation and Modeling of The MOS Transistor*. New York: McGraw Hill, 1999.
- [24] A. Hajimiri et al., "Jitter and Phase Noise in Ring Oscillators," *IEEE J. Solid-State Circuits*, vol. 34, pp. 790-804, June 1999.