

DESIGN OF A PHASE-LOCKED LOOP CIRCUIT IN GALLIUM ARSENIDE
FOR USE IN A HIGH FREQUENCY CLOCK DISTRIBUTION CHIP

by

James H. Bandy

Submitted to the

DEPARTMENT OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

in partial fulfillment of the requirements

for the degrees of

BACHELOR OF SCIENCE

and

MASTER OF ENGINEERING

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

March, 1995

© James H. Bandy, 1995. All Rights Reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly
paper and electronic copies of this thesis document in whole or in part.

Signature of Author _____
Department of Electrical Engineering and Computer Science, March 15, 1995

Certified by _____
C. G. Fonstad, Prof. of Elec. Eng.

Certified by _____
P. J. Zdebel, Company Supervisor

Accepted by _____
F.R. Morgenthaler, Chair, Department Committee on Graduate Students

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

JUL 17 1995

Earker Eng

Design of a Phase-Locked Loop Circuit in Gallium Arsenide For Use in a High Frequency Clock Distribution Chip

by

James H. Bandy

Submitted to the Department of Electrical Engineering and Computer
Science on March 15, 1995, in partial fulfillment of the requirements
for the degrees of Bachelor of Science and Master of Engineering

Abstract

This thesis discusses the design of a phase-locked loop (PLL) in GaAs for clock distribution applications. To reduce the skew between input and output clock signals from a clock distribution chip, a PLL may be used. By designing the PLL for this application in GaAs, higher frequency performance is possible than with a CMOS design. Basic PLL behavior is described, along with a detailed explanation of the motivations for designing with Source-Coupled FET Logic (SCFL) for an E/D MESFET process in GaAs. The work completed included the design of the major blocks for a PLL: a phase detector, a loop filter, and a voltage-controlled oscillator (VCO). The design considerations taken into account for each of these blocks are described in detail. The phase detector design was based on a novel implementation of a three-state phase detector. The VCO design was a ring oscillator with a frequency range of 200-500MHz. An active filter was used for the loop filter. Finally, a behavioral model was constructed to simulate the performance of the full PLL. The steps taken to construct the behavioral model, as well as the results of the simulations, are described.

Thesis Supervisor: Clifton G. Fonstad
Title: Prof. of Electrical Engineering

Company Supervisor: Peter J. Zdebel
Company: Motorola, Semiconductor Products Sector

Table of Contents

1. Introduction	7
1.1 Project Motivation	7
1.2 Scope of Thesis Work	8
1.3 Content and Organization of Thesis	10
2. MESFETs and Source-Coupled FET Logic	13
2.1 Introduction	13
2.2 Fundamentals of MESFETs	13
2.3 Source-Coupled FET Logic	19
3. The Theory of Phase-Locked Loops	23
3.1 Introduction	23
3.2 The Phase Detector	24
3.3 The Voltage Controlled Oscillator	26
3.4 Linear Performance Characteristics	27
3.5 Limits on Linear Behavior	37
3.6 Frequency Acquisition	40
3.7 Final PLL Architecture	47
4. The Phase Detector	49
4.1 Introduction	49
4.2 Design Alternatives	49
4.3 Phase Detector Circuits	58
4.4 Simulated Performance	64
5. The Voltage-Controlled Oscillator	69
5.1 Introduction	69
5.2 VCO Performance Criteria	69
5.3 Design Alternatives	71
5.4 Effects of Using Divider Circuits	72
5.5 Ring Oscillator Circuits	75
5.6 Simulated VCO Performance	79
5.7 Noise Considerations	81
6. The Loop Filter	83
6.1 Introduction	83
6.2 Charge Pumps	83
6.3 Loop Filter Design Trade-offs	86
6.4 Active Loop Filter Architecture	87

6.5	Op Amp Design	91
6.6	Op Amp Performance and Effects on PLL Performance	93
7.	Behavioral Modeling	95
7.1	Introduction	95
7.2	Motivations for Behavioral Model	95
7.3	Construction of Behavioral Models	95
7.4	PLL Performance in the Behavioral Model	100
8.	Conclusions	103
8.1	Summary of Work	103
8.2	Suggestions for Continuing Work	103
	Bibliography	105
	Acknowledgements	107
	Appendices	109
	Appendix A: Schematics for phase detector circuits	111
	Appendix B: Simulated output waveforms for phase detector	117
	Appendix C: Schematics for VCOs	123
	Appendix D: Simulated output waveforms for VCO	129
	Appendix E: Schematics for loop filter	135
	Appendix F: Behavioral model code for mixer	139

Chapter 1 - Introduction

1.1 Project Motivation

As the speed of integrated circuits has increased due to advances in process and device technologies, it has been desirable to increase the rate of communication between chips by operating at higher clock rates. In order to communicate properly, however, the clock must arrive near synchronously at all chips. This requirement has led to demands for smaller and smaller clock signal skews between when the clock signal reaches each chip on a board. In order to address this problem in workstations and personal computers, manufacturers have resorted to using clock distribution chips to synchronize the clock pulse on each board. One main function of the clock distribution chip is to buffer the main clock signal which is normally heavily loaded, causing signal integrity problems. A second function is to redistribute the clock signal through each of its outputs with minimal skew between the outputs.

Initially clock distribution chips performed these functions through sophisticated use of normal buffer circuits with careful matching of the delay paths through the circuit to minimize the output skew. Recently, however, in order to improve the precision of these chips, internal phase-locked loop (PLL) circuits have been used to further reduce the output skew. These PLL circuits provide a negative feedback loop which minimizes the phase error between the edges of the input and output clock signals. Thus, this architecture can more accurately match the phase of the input clock signal than a scheme using normal input buffers. The use of a PLL also provides the clock distribution chip with some additional functionality. Since a PLL can be used for frequency synthesis, it is possible to generate output clock signals from the clock distribution chip at multiples of the input clock frequency. This means the master clock signal distributed throughout the system can be at a lower frequency. Lowering the frequency of this signal permits a lower-cost crystal to be used, reduces the RF interference produced, and simplifies board design. [1]

Currently, clock distribution chips for the workstation and PC market are produced almost exclusively using silicon CMOS processing technologies. These chips operate at relatively low frequencies where CMOS circuitry performs well while consuming small

amounts of power. The introduction of new high speed microprocessors, such as Intel's Pentium (66MHz) and Motorola's Power PC (100MHz), stretch the maximum frequency range of CMOS circuits. Also, CMOS power consumption, which is a function of operating frequency, becomes less competitive. This has led to the investigation of alternative process technologies. Two processing technologies which are being considered because of their inherent speed / power performance ratios are silicon BiCMOS and gallium arsenide E/D MESFET. The GaAs process is particularly attractive because it is a lower complexity process which would allow the chips to be produced at a lower cost.

The critical element in these new clock distribution chips will be the PLL circuitry. Optimization of the logic which surrounds the PLL in these circuits will only improve performance to a certain point. The performance of the PLL will be the predominant determinant of what output skew can be achieved. The motivation behind the work done for this thesis was to explore the alternatives for designing a high performance PLL suited to this application, and to develop a design for such a PLL in GaAs.

1.2 Scope of Thesis Work

The primary focus of this thesis was the circuit design work required to build a phase-locked loop in gallium arsenide for use in a high frequency clock distribution chip. To form a basis for doing this work, solutions to this type of design in other semiconductor processing technologies were investigated and evaluated as to whether they could be realized in GaAs. This process involved both researching the relevant literature and speaking to other engineers within Motorola who were familiar with the design of PLLs. A number of the different design alternatives which were considered as part of this process, as well as the important design considerations that became apparent, are discussed within this thesis.

At the end of this research process, a set of designs were selected for the composition of a complete PLL circuit. In order to construct a PLL, three major circuits must be designed: a phase detector, a loop filter, and a voltage-controlled oscillator (VCO). Each of these circuits is described in detail later in this thesis. The arrangement of these circuits within the PLL is illustrated in Figure 1.1. The designs selected to

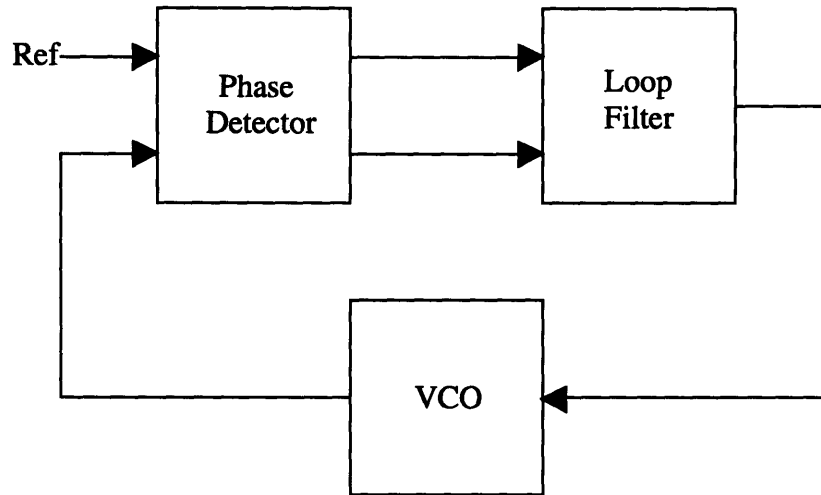


Figure 1.1 - Basic phase-locked loop design for clock synchronization

implement these circuits were then redesigned in Source-Coupled FET Logic (SCFL) in order to be implemented in an E/D MESFET GaAs process. The resulting circuits were simulated with HSPICE to verify their functionality and performance. The interactions between the phase detector circuit and the loop filter circuit and between the loop filter circuit and VCO circuit were also tested via simulations with HSPICE.

As a final step in the design process, the overall performance of the PLL was simulated through the construction of a behavioral model. The majority of the PLL circuits were modelled using standard elements available in the SABER behavioral modeling software package which was used. For the one circuit which could not be constructed out of standard elements, a new behavioral model was developed based on detailed HSPICE simulations of that circuit. The final behavioral model was used to perform extensive tests of the full PLL's performance.

While the phase detector and VCO circuits designed for this thesis were also laid out for fabrication, this work is not discussed here since it was done by other members of the design team. In addition, changes which were being made in the digital GaAs process flow used to fabricate these circuits delayed the return of fabricated circuits. This prevented the testing of fabricated circuits during the period when this thesis work was done. Some subsequent testing of the VCO circuits, however, has been done. These tests

show the performance of the actually circuits to be close to the performance obtained in HSPICE simulations. [2]

It should be noted that this work was part of the first attempt within Motorola at designing a PLL in E/D MESFET GaAs technology. It was also part of the first attempt at designing a PLL-based, monolithic clock distribution integrated circuit with a frequency range of 100MHz to 200MHz in E/D MESFET GaAs technology.

1.3 Content and Organization of Thesis

The content of this thesis has been written in the expectation of sufficiently instructing the reader such that they could continued the work begun here. Chapters 2 and 3 are intended to provide the reader with the necessary background for understanding the work discussed in the rest of the thesis. First, Chapter 2 examines the basic characteristics of MESFETs, as well as the basic elements of Source-Coupled FET Logic. This chapter also further explains the motivations behind using a GaAs E/D MESFET process and the reasons for designing in SCFL. Chapter 3 introduces the reader to the fundamentals of PLL design. It begins by describing the basic characteristics of both the PLL and its major blocks. This is followed by the derivation of the major equations for describing PLL performance.

Chapters 4, 5, and 6 are devoted to the details of the design of the phase detector, VCO, and loop filter, respectively. In each chapter, the design alternatives which were considered are presented first, along with the justification for the approach which was chosen. This comparison process also points out the important design parameters associated with each of the circuits. Next, the important details of the circuit architectures which were designed are examined. Finally, the simulated performance of each circuit, as well as any affects of this performance on the full PLL's behavior, are discussed.

Chapter 7 discusses the behavioral modelling work which was performed, and the results which were obtained. The final chapter, Chapter 8, summarizes the work which was done for this thesis, and makes suggestions for continuing the work.

References

- [1] D. Chen, "Designing On-Chip Clock Generators," *IEEE Circuits and Devices*, vol. 8, pp 32-36, July 1992.
- [2] R. Sodhi, "Development of a DC Pulse Test Circuit for an Electromigration Study," Thesis for Bachelor of Science, MIT, Cambridge, MA, 1995.

Chapter 2 - MESFETs and Source-Coupled FET Logic

2.1 Introduction

This section essentially looks at the environment in which the circuits for this thesis were designed. Before any attempt was made to design the circuitry for the PLL, a number of decisions had already been made. The first of these was that the design was going to be attempted in Gallium Arsenide (GaAs). Since GaAs lacks a stable native oxide this meant that MOSFETs (metal-oxide-semiconductor field effect transistors) could not be used for the design, but instead MESFETs (metal-semiconductor field effect transistors) would be used. This section will present the basic structure and behavior of these devices and will point to some of the major differences between these devices and MOSFETs. The reasons for using only enhancement and depletion devices, rather than complementary devices, will also be mentioned.

Another decision that was made before beginning the circuit design process was to design in Source-Coupled FET Logic (SCFL). This section will discuss the basic characteristics of SCFL designs, which closely resemble those of ECL designs using bipolar devices. The reasons for choosing this family of logic over Direct-Coupled FET Logic (DCFL), the other major logic style in GaAs technology, will also be discussed.

2.2 Fundamentals of MESFETs

As was mentioned above, a designer working in GaAs has no choice but to design with MESFETs instead of MOSFETs. To someone not familiar with circuit design in GaAs, the term MESFET may also be unfamiliar. In silicon, the majority of all digital circuit design is done with MOSFETs. Thus, in addition to explaining the basic behavior of MESFETs, this section will also point out the important differences between the two types of devices. This should help to make clear the trade-offs involved in choosing to design in GaAs rather than Si.

To begin, a cross-sectional view of a typical MESFET, as well as a MOSFET, are shown in Figure 2.1. Looking at this illustration, one sees that a MESFET is a three terminal device consisting of a gate contact, a source contact, and a drain contact. The source and the drain contacts are both ohmic contacts to n^+ -GaAs. The two contacts are

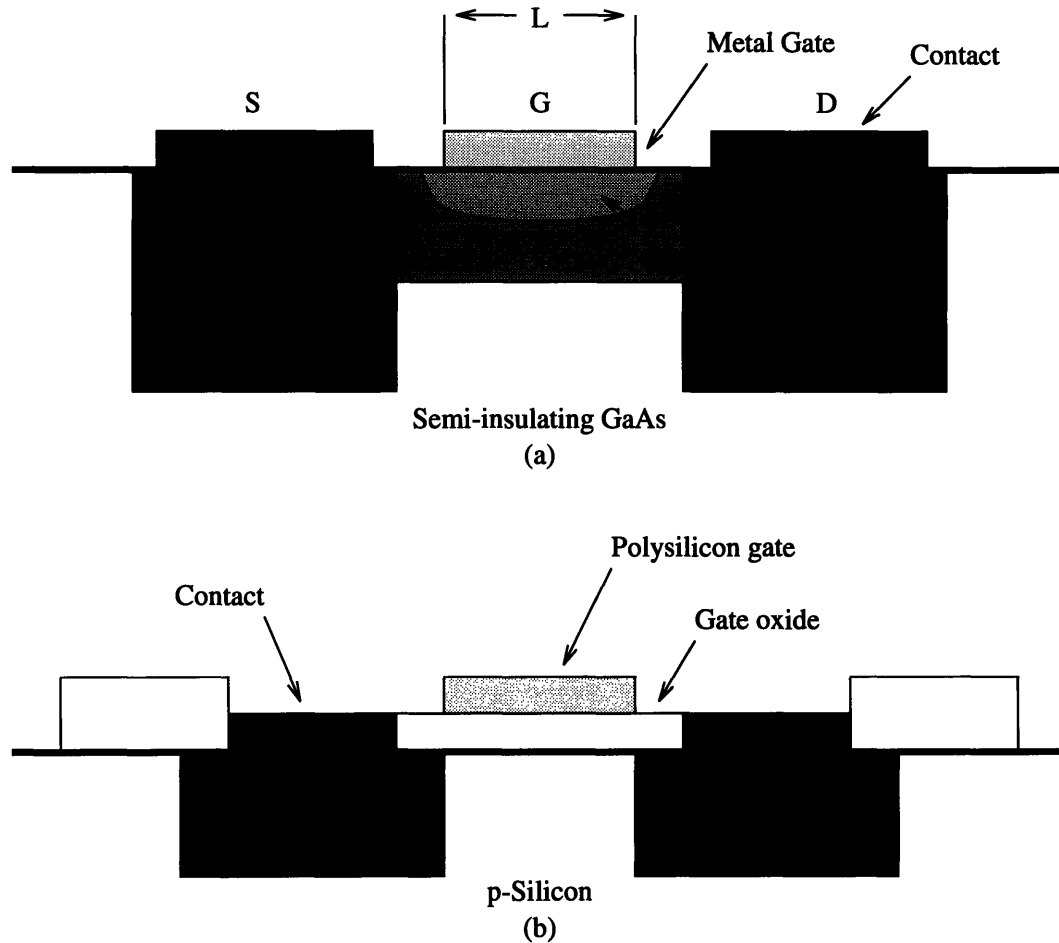


Figure 2.1 - Cross-sectional views (not to scale) of a MESFET (a) and a MOSFET (b)

physically identical because of the symmetry of the device, however, by convention they have different functions. A MESFET will always be placed in a design so that the current flows from the drain contact to the source contact. The names of the two contacts come from the electron flow within the device for this direction of current. Electrons will always be sourced at the source contact and drained away at the drain contact.

The gate contact is the control node for a MESFET. The biasing of this contact relative to the source contact will have the major effect on the magnitude of the device's drain-source current. This functionality of the three terminals is the same as that for the three terminals of a MOSFET. The critical difference between these two types of devices, however, is in how the gate contact to the channel region is formed. In a MOSFET, the gate is separated from the channel by a layer of oxide which insulates the two regions. In a MESFET, though, the gate makes direct contact with the channel region, forming a

metal-semiconductor or Schottky contact. This connection constrains the gate voltage of a MESFET because of the diode's forward conduction limit. As will be shown later, this property turns out to be an important factor in using this device in DCFL designs.

The connection of the gate to the channel also points to a second difference between the two devices. In the MESFET, the gate connects to a pre-existing channel formed by the n-GaAs region which extends between the source and drain contact regions. The width of this channel is modulated by varying the space-charge region of the Schottky diode formed by the gate contact. Decreasing the voltage of the gate will cause the space-charge region of the Schottky diode, where majority carriers are depleted, to grow wider and pinch off more of the channel. The diagram of the MOSFET, which represents an enhancement type MOSFET, however, shows that there is no pre-existing channel for this device. Instead, the gate contact, which functions like a capacitor, must be used to invert the surface region under the gate. Positive charge placed on the gate will cause electrons to be pulled into the area under the gate, and when the number of electrons is sufficient to become the dominant carrier in this region, a conducting channel will be formed. The depletion MOSFET uses this same mechanism for modulating channel width, but does include a pre-existing, doped channel.

One last difference between the two devices in terms of construction is that a MESFET is built in a semi-insulating substrate while the n-type MOSFET is built in a p-type substrate. On the plus side for MESFETs, this leads to lower parasitic device capacitances, since MOSFETs must deal with the space charge region associated with a reverse biased pn junction. This pn junction, however, provides better isolation between devices. While the resistance of the semi-insulating substrate is quite high, typically on the order of 10^6 to $10^8 \Omega\text{cm}$, the possibility of charge flowing through the substrate and causing device interaction leads to less dense spacing rules than for Si MOSFETs.

Since choosing to work with MESFETs or MOSFETs is essentially a direct function of choosing whether to work in GaAs or Si, it is also worth noting the underlying differences in materials that one tries to take advantage of by designing in GaAs. The primary advantage of GaAs over Si is the dramatically higher mobility of electrons in GaAs for low electric fields. Translated, this means higher speeds at lower power levels. A disadvantage, however, is the slightly lower mobility of holes in GaAs. In Si, where




Figure 2.2 - Example $I_{DS} - V_{DS}$ characteristic for a MESFET ($V_{GS2} > V_{GS1}$, etc)

these two mobilities are much closer together, the advantages of CMOS circuit design techniques can be utilized with a more acceptable penalty in speed. In GaAs, the complementary approach is avoided in order to preserve the speed advantage. For this reason, designs in GaAs generally use n-type enhancement and depletion devices and no p-type devices. This is similar to Si NMOS design. By avoiding using p-type devices, a second drawback of these devices is also avoided. The low barrier voltage of most metals on p-type GaAs makes fully depleting the channel difficult. This leads to higher leakage currents with p-type GaAs MESFETs.

Having outlined the motivations for designing with GaAs MESFETs, the behavior of these devices can now be looked at in more detail. A graph illustrating the I_{ds} vs. V_{ds} characteristic for a typical MESFET is shown in Figure 2.2. I_{ds} is the current from the drain contact to the source contact, while V_{ds} is the drain to source voltage. This graph includes a number of curves for different gate-source voltages, V_{gs} . The performance of a MESFET is broken into three regions: cut-off, linear, and saturation. In the cut-off region, the V_{gs} is such that the entire channel is pinched off by the depletion region of the Schottky diode without any V_{ds} being applied to the device. The V_{gs} for which this condition occurs is generally called the threshold voltage, V_T . While some current can flow through the device under this condition, called the subthreshold current, this current is very small compared to the magnitude of the currents in the other regions of operation.

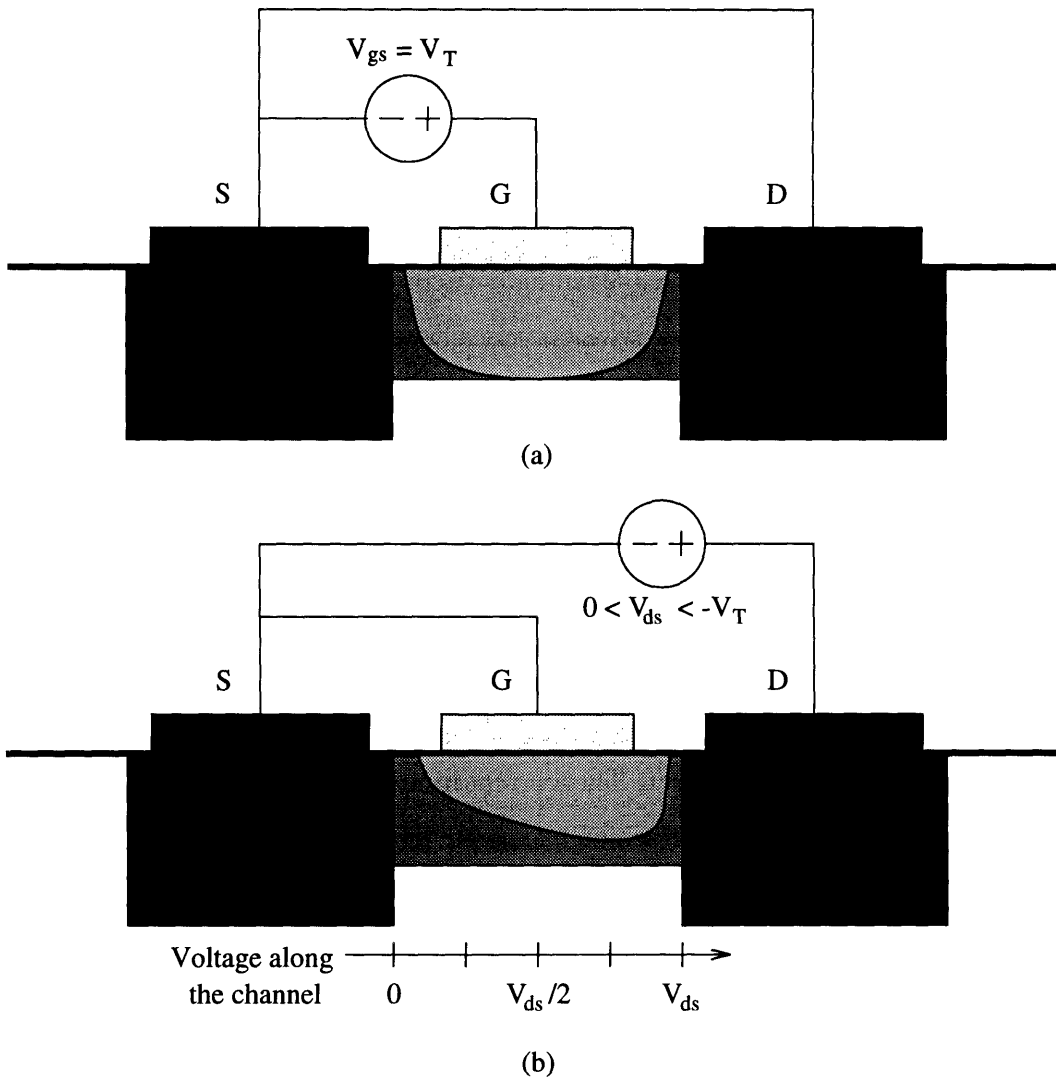


Figure 2.3 - Illustration of MESFET channel in cut-off region (a) and in linear region (b).

A diagram illustrating this condition is shown in Figure 2.3 (a). The difference between a depletion device and an enhancement device when expressed in terms of V_T is that an enhancement device has a positive V_T , while a depletion device has a negative V_T .

The next region of performance is the linear region. Here V_{gs} is greater than V_T , meaning that the channel will not be pinched off when no V_{ds} is applied. If a finite V_{ds} is now applied, current will flow through the channel. This region is called linear because for small V_{ds} , the increase in current through the channel is directly proportional to the increase in V_{ds} divided by the resistance of the channel. The range of V_{ds} over which this increase in current is linear is fairly small, however, because of a second effect that starts

to limit the increase in current. The V_{ds} applied to the device must be dropped along the channel between the drain and source. This results in a decrease in the gate to channel voltage near the drain contact which causes the channel to narrow at this end. Thus the increase in current through the channel for higher V_{ds} is being countered by the narrowing of the channel. This effect is illustrated in Figure 2.3 (b) for a depletion device. If V_{ds} continues to increase, the channel will eventually become completely pinched-off at the drain end of the device. This voltage, which varies depending on V_{gs} , is called the saturation voltage, $V_{ds,sat}$. The equation for $V_{ds,sat}$ is

$$V_{ds,sat} = V_{gs} - V_T \quad (2.1)$$

For devices with short gate lengths ($L < 2\mu\text{m}$), the current saturation point is moved even lower due to another effect that limits the current flow. In GaAs, the maximum electron velocity is reached for fairly low electric fields. For short gate length devices, this field strength will be reached at a rather low V_{ds} . This V_{ds} will be below the V_{ds} required to pinch off the channel. This current limiting effect plays a significant role with the devices used in this thesis, which had a gate length of $0.7\ \mu\text{m}$.

The final region of operation is the saturation region. In this region, V_{ds} is greater than $V_{ds,dat}$, so that V_{ds} no longer significantly affects the current through the device. There is some increase in current for higher V_{ds} , which translates into a finite output conductance. This is the desirable region to operate a MESFET in because its gain is highest here. As a side note, the gain of a short gate length MESFET is roughly 3-4 times higher in this region than that of a comparable Si MOSFET. Since this is the region where the circuits in this thesis were primarily operated, a simplified equivalent circuit model for a MESFET operating in this region is shown in Figure 2.4.

This model points out the major concerns a designer must keep in mind during the design process. The equation for I_{ds} in this model is

$$I_{ds,sat} = \beta (V_{gs} - V_T)^2 \quad (2.2)$$

$$\beta = \frac{2\epsilon_s \mu_n v_{sat} W}{b (\mu_n V_{po} + 3v_{sat} L)} \quad (2.3)$$




Figure 2.4 - Simplified equivalent circuit model for a MESFET operating in saturation region.

Rather than getting into the details of the equation for β , from the designers point of view, all the parameters for β , except for the devices width, W , are fixed for a process used. While gate length is sometimes deliberately increased in order to lower currents in non-speed critical portions of a design, all high performance areas of a circuit will be designed using the minimum gate length possible with the process being used. Taking this simplified view, a few notes should be made about this current equation. First, V_{ds} is completely absent. Second, the devices current gain is a function of the square of V_{gs} . Finally, the gain of the device can be increased by increasing the width of the device. There is a trade-off here, however, in that the input node capacitance, C_{in} , increases proportionally with gate width. As a last note, the diode in this model points out the limit on the gate voltage of this device which is generally 0.6 to 0.7 V.

2.3 Source-Coupled FET Logic

This section explains the reasons for designing the circuits for this thesis in Source-Coupled FET Logic (SCFL). The basic elements of SCFL design are presented. Throughout this section, a comparison is also made to Direct-Coupled FET Logic (DCFL) design, since this is the other major style of logic design which was considered. To begin, schematics for both a basic DCFL and SCFL gate are shown in Figure 2.5.

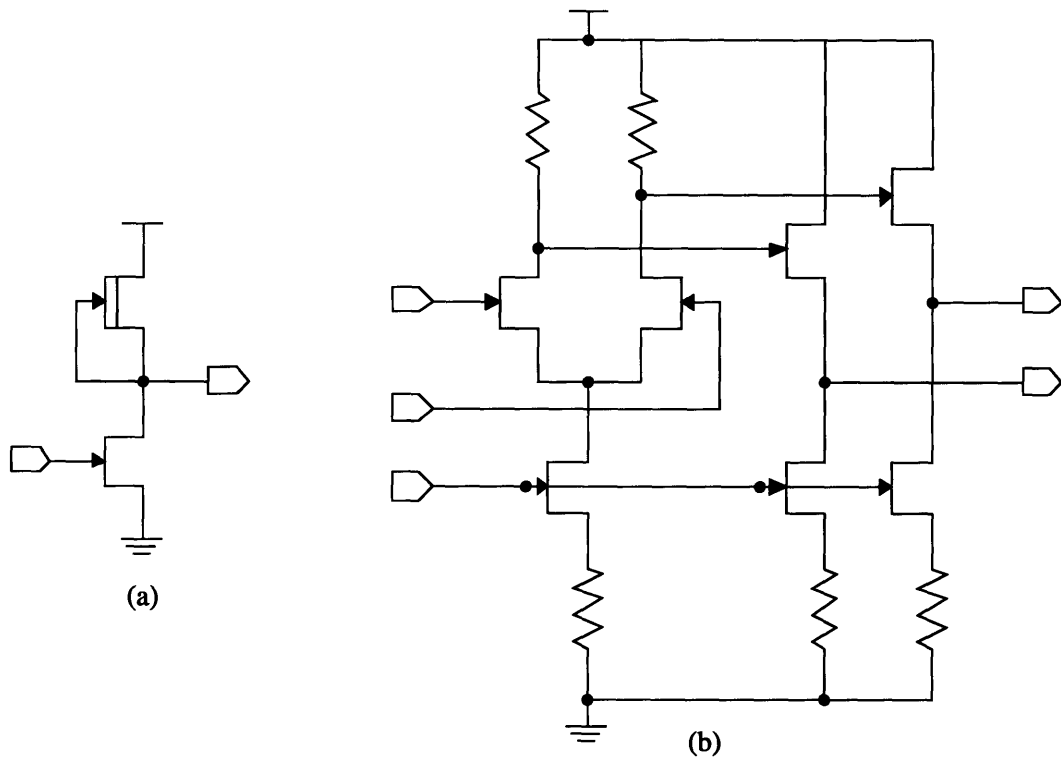


Figure 2.5 - Schematic of a DCFL gate (a) and a SCFL gate (b).

On a physical level, there are a number of differences between these two gates. Most obviously, the DCFL gate is quite significantly smaller. This is clearly good for VLSI, where one wants to pack the gates as close together as possible. It also means that yields for a particular circuit should be relatively better because of the smaller die area that will be required. On a more subtle level, the smaller number of devices required for the DCFL gate means that it can operate within narrower rail-to-rail voltages. This in turn can lead to designs that require less power than a comparable SCFL design.

Another physical difference is that a standard SCFL gate uses resistors, while DCFL gates do not. In terms of process complexity, not having resistors is an advantage for DCFL design. On the other hand, while it might be possible to remove the resistors from a SCFL design by using active loads and uncompensated current sources, there are advantages associated with including them. The first is that they allow a designer more precise control over the logic swing of the gate. In a DCFL gate, the output low voltage depends on the current within the gate and the drain-source resistance of the enhancement device, while the output high voltage is clamped by the Schottky diode behavior of the

subsequent gates. In both cases, the swing is controlled by inherent parameters of the devices being used. Another important benefit of the resistors is that they provide compensation for both process and temperature variations. The output swing in an SCFL gate is controlled by a resistor ratio. Consider a case where resistor values in an SCFL design are lower than normal because of either temperature or process variation. The lower resistor values in the switching portion of the gate would mean a smaller output swing if the current through the gate remained fixed. The smaller resistor values in the current source portion of the gate, however, will produce a compensating increase in the current. Finally, the control a designer has over the output swing in this type of SCFL design has another advantage in that it allows a designer to design for smaller voltage swings. This in turn will allow for higher speeds.

On an architectural level, there are several more differences. First, the input and output for a SCFL gate is differential, while these are single-ended for a DCFL gate. This property minimizes the common-mode noise in a SCFL gate. This is a tremendous advantage since most noise that appears is common-mode. The architecture of a SCFL gate also makes it somewhat less sensitive to variations in device thresholds. Assuming that most variations in the threshold are seen across a wafer, but not locally, the switching portion of the SCFL gate should not be affected by the threshold voltage variation because it depends on a relative comparison of $V_{gs} - V_T$ in determining the current through the two sides of the current switch. The gate's overall current should not be affected since they are set primarily by the resistors. Another advantage of the SCFL gate is that the source followers give it superior fan out capability over that of a DCFL gate. Finally, one architectural similarity between SCFL and DCFL is that they are both constant current logic families, as opposed to CMOS design. This is an advantage over CMOS because it leads to reduced switching noise. At the same time, this means that when operating at the same voltage, neither SCFL or DCFL has a definite power advantage, but instead this depends on the details of the circuits. As mentioned above, though, an all DCFL circuit can be operated at lower voltages.

The eventual decision to design in SCFL reflected both circuit design considerations and concerns over the process which was going to be used. The better noise performance of SCFL circuits because of their differential inputs and outputs was

important because of the requirement for very low noise-induced jitter within the PLL. Power consumption was not as much an issue since whatever circuit was designed was still expected to be operated at 3V. Finally, the compensation properties of the SCFL gate were considered valuable since the process being worked with was still very new and was expected to have significant variations in device characteristics.

Chapter 3 - The Theory of Phase-Locked Loops

3.1 Introduction

This section looks at the basic theory surrounding the design and functioning of phase-locked loops. The object of this section is to provide the reader with a sufficient background to understand some of the higher level considerations that went into the design of the circuits for this thesis. For those readers who are interested in pursuing this topic further and designing their own PLLs, some suggestions for further reading can be made. Wolaver's book provides an excellent look at PLL design for a practicing engineer, focusing on the knowledge required to design a standard PLL. [1] Gardner's book is one of the classic texts on PLLs, providing more of the theory behind PLLs. [2]

The basic goal of a phase-locked loop is to produce an output signal whose frequency and phase are matched to the desired frequency component of its input signal. In its simplest form, the functionality of a PLL can be achieved by placing a phase detector together with a voltage controlled oscillator in a negative feedback loop. In all practical designs, however, a low-pass filter circuit is added into the forward path of the feedback loop because of the dramatic performance improvements it produces. This filter is called the loop filter because of its position within the feedback loop. Figure 3.1 shows this basic architecture.

Looking at this figure, one sees that a PLL is essentially a negative feedback system. The error signal which this system tries to minimize is the phase error between its input and output signals. Assume that the system starts with the phase and frequency of the two signals matched. Now suppose that the frequency of the input signal increases slightly. This will cause the phase of the PLL's output to start to fall behind that of the




Figure 3.1 - Basic phase-locked loop architecture

input signal. The increase in phase error, however, will cause the signal out of the phase detector to increase. Neglecting momentarily the effect of the loop filter, this higher phase detector output voltage will cause the VCO output frequency to increase. This will push the phase of the output back towards that of the input, so that the phase error is minimized again. At the same time, the frequency of the output signal will be tuned to the new frequency of the input signal. In general, one should be aware of the tight interconnection between phase and frequency, with frequency being the derivative of phase. It is this relationship that allows a change in frequency to be used to correct a phase error. Now consider the loop filters affect on this process. At low frequencies the loop filter might have a high gain which lowers the error signal the system will tolerate, while at high frequencies it may attenuate the system's response to changes in the input signal. This second effect may be desirable for eliminating the effect of higher frequency noise in the input signal. This example, however, only looks at a PLL's behavior for one region of operation.

The rest of this chapter is broken into a number of sections that address in a more detailed manner the behavior of a PLL within each of its regions of operation. One section will look at the in-lock behavior of a PLL. The term "in-lock" is used to indicate that the frequency of the PLL's oscillator has been tuned close enough to the frequency of the input that the PLL's behavior is linear. The performance of a PLL will be looked at under both steady-state and AC conditions in this section. The following section will look at the limits on the linear behavior of a PLL. This section describes the factors that determine a PLL's linear range of behavior and looks at what conditions might push a PLL out of its linear range. The next section of this chapter will deal with the frequency acquisition process, where the VCO's output frequency is tuned to within the PLL's linear range. A final section describes briefly the final architecture that was used for the PLL. First, however, the general output characteristics of the phase detector and the VCO are described and equations for their behavior are derived.

3.2 The Phase Detector

A phase detector's basic function is to compare the phases of two input signals and to produce an output voltage that is proportional to the difference in phase between the




Figure 3.2 - Example phase detector output characteristic

two signals. Assuming from Figure 3.1 that θ_i and θ_o represent the phases of the phase detector's two input signals and that v_d is the phase detector output voltage, the equation for this behavior is

$$v_d = K_d (\theta_i - \theta_o) \quad (3.1)$$

In this equation, K_d is the phase detector gain and has units of V/rad. The difference in phase between the two inputs is referred to as the phase error, θ_e .

$$\theta_e = \theta_i - \theta_o \quad (3.2)$$

This simplifies Equation (3.1) to

$$v_d = K_d \theta_e \quad (3.3)$$

Figure 3.2 shows a potential output characteristic for a phase detector that illustrates this equation.

This particular output characteristic is called a triangular phase detector characteristic because of its triangle wave shape. The first thing to notice about this phase detector characteristic is that it is periodic. This characteristic repeats with a period of 2π . While it is possible to extend a phase detector's period beyond 2π by adding memory to the phase detector, all phase detectors do in fact have some period to their output

characteristic because of the natural 2π period of phase itself. Even more limiting in this case, however, is the fact that this phase detector has the same output voltage for multiple phase errors within this 2π period. For example, the signal for a phase error of $3\pi/4$ will look the same to the rest of the PLL as that for a phase error of $\pi/4$. While not all phase detectors exhibit this second property, when present this property together with the period of the phase detector determine the range of phase error for which the phase detector's behavior is linear. To address this limitation, a linear range is defined for each type of phase detector. In this range, the output of the phase detector corresponds to only one possible phase error. The linear range for this particular phase detector would be $-\pi/2 < \theta_e < \pi/2$. When the phase error within the PLL exceeds this linear range, the PLL's behavior becomes non-linear, making it difficult to predict. These instances may cause a previously in-lock PLL to lose lock on the input signal.

A second interesting property of this output characteristic is that there is an output voltage from the phase detector when the phase error is zero. This voltage is called the phase detector offset voltage, V_{do} . This offset voltage has important consequences on both the steady-state and AC performance of the PLL. Finally, the phase detector gain, K_d , is illustrated graphically in this output characteristic. The phase detector gain is equivalent to the slope of the output voltage within the detector's linear range.

3.3 The Voltage Controlled Oscillator

The basic behavior for a voltage-controlled oscillator is to output either a sinusoidal or square wave signal at a frequency which is proportional to the voltage at its input. Figure 3.3 illustrates a potential output characteristic for a VCO. Notice that, in order to produce the desired output frequency, the PLL must generate a non-zero voltage at the input of the VCO, and that this voltage varies depending on the frequency of the input to the PLL. This voltage is called the VCO offset voltage, V_{co} . Using this offset voltage, it is possible to derive a linear equation for the deviation in VCO frequency, $\Delta\omega$, away from the input frequency.

$$\Delta\omega = K_o (V_c - V_{co}) \quad (3.4)$$

$$\Delta\omega = \omega_i - \omega_o \quad (3.5)$$




Figure 3.3 - Example VCO output characteristic

Within this equation, K_o is a new variable called the VCO gain. K_o is measured in units of rad/s/V. Using this linear model for the VCO together with the linear model for the phase detector, it is now possible to proceed with the analysis of PLL performance.

3.4 Linear Performance Characteristics

In this section, the key measures of a PLL's linear behavior will be defined and examined. A PLL's behavior is considered to be linear when all of its components are operating within their linear ranges. Since the loop filter's behavior is always linear, this means that the phase error in the system must be within the phase detector's linear range, and the frequency of the input signal must be within the VCO's frequency range. A more detailed look at the limits on linear behavior within a PLL will be taken in the next section. Within this section, two important parameters of PLL behavior are discussed, the static phase error and the frequency response. Also, the loop filter and its effect on loop performance will be discussed in more detail. In particular, it will be shown how the loop filter is used to resolve a compromise between the steady-state behavior and the frequency response of a PLL.

To begin, suppose that a PLL has been constructed like the one in Figure 3.1, where the phase detector and VCO have the output characteristics shown in Figures 3.2 and 3.3, respectively. For the time being, assume that the loop filter has a DC gain, $F(0)$, of one. This PLL will be used to illustrate how a static phase error can appear within a PLL.

Suppose that the example PLL has locked to a steady-state input signal, ω_i , at 10 MHz. Looking at the VCO characteristic, one sees that the voltage at its input, V_c , must be 3 V. This voltage must be supplied to the VCO by the phase detector. Looking at the phase detector's output characteristic, however, it becomes apparent that there must be a phase error between its inputs in order to generate an output voltage of 3 V. This phase error is called the static phase error. Thus, within this PLL, some fixed phase error, which depends on the frequency of the input signal, must be present even when the PLL has fully locked to the input signal.

In order to develop a quantitative method for determining static phase error, a DC linear model for the loop is presented in Figure 3.4. This linear model was taken from Wolaver's book. [3] Within this model, the variable $\Delta\omega$ represents the deviation of the output frequency from the input frequency. The equation for $\Delta\omega$ in this system is

$$\Delta\omega = \theta_e K_d F(0) K_o + V_{do} F(0) K_o - V_{co} K_o \quad (3.6)$$

For the system to be in lock, $\Delta\omega$ must equal zero. This, however, requires the presence of a static phase error of a sufficient magnitude to negate the voltages contributed by V_{do} and V_{co} . Equation 3.6 can be manipulated to produce an equation for this phase error.

$$\theta_e = \frac{-V_{do}}{K_d} + \frac{V_{co}}{F(0) K_d} \quad (3.7)$$

For the example PLL, $V_{do} = 3.5$ V, $V_{co} = 3.0$ V, $F(0) = 1$, and $K_d = 3/\pi$ rad/V applies. Thus, there must be a static phase error of $-\pi/6$ rads.

The static phase error found with a PLL is important because of its limiting effect on the linear range of the PLL. If $F(0)$ were very small or V_{co} was very large, the static

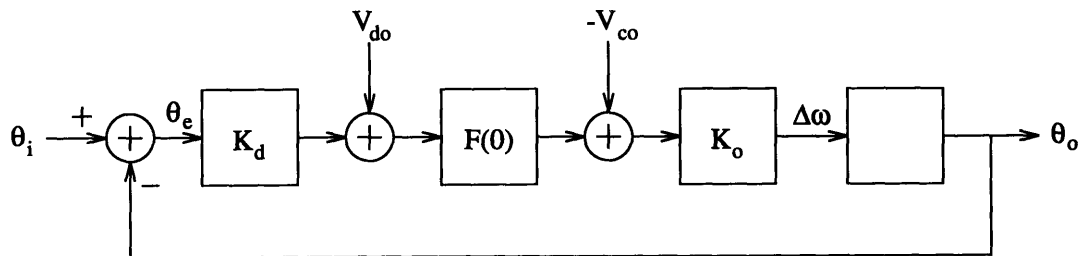


Figure 3.4 - dc linear model of PLL

phase error required could be outside the linear range of the phase detector. In such cases, the PLL would not be able to lock to the input signal. The static phase error also has an important effect on the dynamic performance of the PLL. When the frequency of the input signal changes, a dynamic phase error will generally appear in the PLL which, depending on the direction of the frequency shift, may be additive to the static phase error. The combination of these phase errors may exceed the phase detector's linear range and thus possibly push the PLL out of lock. By eliminating any static phase error, a PLL has a better chance of staying in lock, where its performance is more predictable.

As Equation (3.7) suggests, one way to get rid of the VCO's contribution to the static phase error would be to make the DC loop filter gain very large. This is, in fact, what almost all PLLs do by including a loop filter that contains an integrator. The integrator gives the loop filter a DC gain of essentially infinity, and thus completely eliminates the effect of V_{co} . This makes sense physically since the integrator in the loop filter may be charged to whatever V_{co} is required. The contribution of V_{do} is handled by making changes in the details of the phase detector circuitry that either reduce V_{do} to zero or hide it from the rest of the PLL. Examples of how this can be done may be found in Wolaver's book. [4]

As was alluded to earlier, eliminating static phase error within the PLL is not the loop filter's only purpose. The loop filter is also the one element of a PLL a designer can modify in order to control the frequency response or, alternately, the bandwidth of a PLL. Before looking at how the loop filter is used in this regard, however, the frequency response of the example PLL, before the loop filter is added, should be examined.

The terms frequency response and bandwidth are used in regard to a PLL to reflect its ability to track changes in its input signal. A PLL with a low bandwidth will have problems tracking an input signal whose phase or frequency are changing rapidly. While this may be problematic in some situations, it may be desirable in others where the rapidly changing component of the PLL's input is due to noise. Another way of looking at the bandwidth of a PLL is in terms of how much of a reaction does an error in phase produce in the PLL. If a PLL has a phase detector with a very high gain, a small phase error will produce a large signal out of the phase detector. If the PLL also has a high VCO gain, the signal out of the phase detector will produce a large step in the output frequency of the

PLL. A PLL with these characteristics will be able to respond quickly to a change in the input signal. On the other hand, this PLL will also respond dramatically to any noise on the input signal. The bandwidth a designer aims for depends on what application the PLL is going to be used in.

In order to look at a PLL's frequency response quantitatively, the system function for the PLL must be derived. The first step in this process is to find the Laplace transforms of the equations for the phase detector and the VCO. For the phase detector, Equation (3.3) becomes

$$v_d = K_d \theta_e \quad \Rightarrow \quad V_d(s) = K_d \theta_e(s) \quad (3.8)$$

For the VCO, Equation (3.4) becomes

$$\Delta\omega = K_o (V_c - V_{co}) \quad \Rightarrow \quad \omega(s) = K_o V_c(s) \quad (3.9)$$

Within the feedback loop, however, the value of interest that is fed back from the VCO is its phase, not its frequency. Phase, however, is simply the integral of frequency.

$$\theta_o = \int \omega_o dt \quad (3.10)$$

In the frequency domain this is equivalent to a division by s , so that Equation (3.9) becomes

$$\theta_o(s) = \frac{\omega(s)}{s} = \frac{K_o V_c(s)}{s} \quad (3.11)$$

Using these new equations, an AC model of the PLL's behavior can be derived. This AC model is illustrated in Figure 3.5. Notice, that since this is an AC model, the DC offsets of both the phase detector and the VCO may be neglected. Also, while this model




Figure 3.5 - ac linear model for PLL

includes the loop filter gain, K_h , it is temporarily assumed that its gain is unity across all frequencies.

The system function for a PLL may now be derived from this AC model. Looking at the AC model, one sees that a PLL is essentially a negative feedback system with unity feedback. The forward gain of this system is

$$G(s) = \frac{K_d K_o}{s} \quad (3.12)$$

Now, from feedback theory, the equation for the system function of a PLL is

$$H(s) = \frac{\theta_o(s)}{\theta_i(s)} = \frac{G(s)}{1 + G(s)} \quad (3.13)$$

Those familiar with feedback theory should be careful to note that $H(s)$ is used to represent the whole system function here, rather than the feedback gain which, in this case, is always unity. Now, substituting $G(s)$ in Equation (3.13) gives the full system function of the PLL without the loop filter.

$$H(s) = \frac{K_d K_o}{s + K_d K_o} \quad (3.14)$$

At this point, the exact way one defines the bandwidth of a PLL is somewhat arbitrary. One common measure of a system's bandwidth, however, is its -3db frequency. At low frequencies, the system function's gain is determined by the real components of the equation. For high frequencies, however, the imaginary components of the equation will dominate, and the system function's gain falls off rapidly with increasing frequency. The -3db frequency marks the transition between these two regions. The -3db frequency of this system is

$$\omega_{-3db} = K_d K_o \quad (3.15)$$

corresponding to the pole of the system function. The Bode plot of this PLL's system function is shown in Figure 3.6 to further illustrate the significance of the -3db frequency.

The problem with this PLL, where the loop filter has been left out, is that the designer has very little control over its bandwidth. The gain of the phase detector and the




Figure 3.6 - Bode plot for PLL without loop filter

VCO are largely fixed by the type of circuit chosen and the process technology used. By adding a loop filter to the PLL, the designer introduces an element into the system whose gain at high frequencies may be set arbitrarily, in order to control the PLL's bandwidth.

Assuming that an active loop filter design is used, the loop filter actually serves two purposes within the PLL. For the low frequency portion of the signal from the phase detector, the loop filter acts as an integrator so that there is no static phase error in the PLL. For the higher frequency portion of the signal, the loop filter functions as either an attenuator or an amplifier, depending on which is necessary to achieve the desired loop bandwidth. The behavior of the loop filter is therefore essentially the same as that of an integral plus proportional controller from feedback theory. An example of how the loop filter might be implemented is shown in Figure 3.7(a). The functionality of this loop filter can be implemented in a much more efficient architecture that requires only one op amp as shown in Figure 3.7(b). It is also possible to implement the loop filter using only passive components, as shown in Figure 3.7(c), however, this design is seldom used. The DC gain of a passive loop filter, $F(0)$, can not exceed one. This means that, unlike the active filter designs which have high DC gains, this type of loop filter will not eliminate the static phase error caused by the VCO offset voltage. Also, a passive loop filter can only be used to attenuate at high frequencies. This is, however, the desired performance in most cases. On the other hand, in applications where a large static phase error can be tolerated, the advantage of a passive loop filter is that an op amp circuit does not have to be designed.

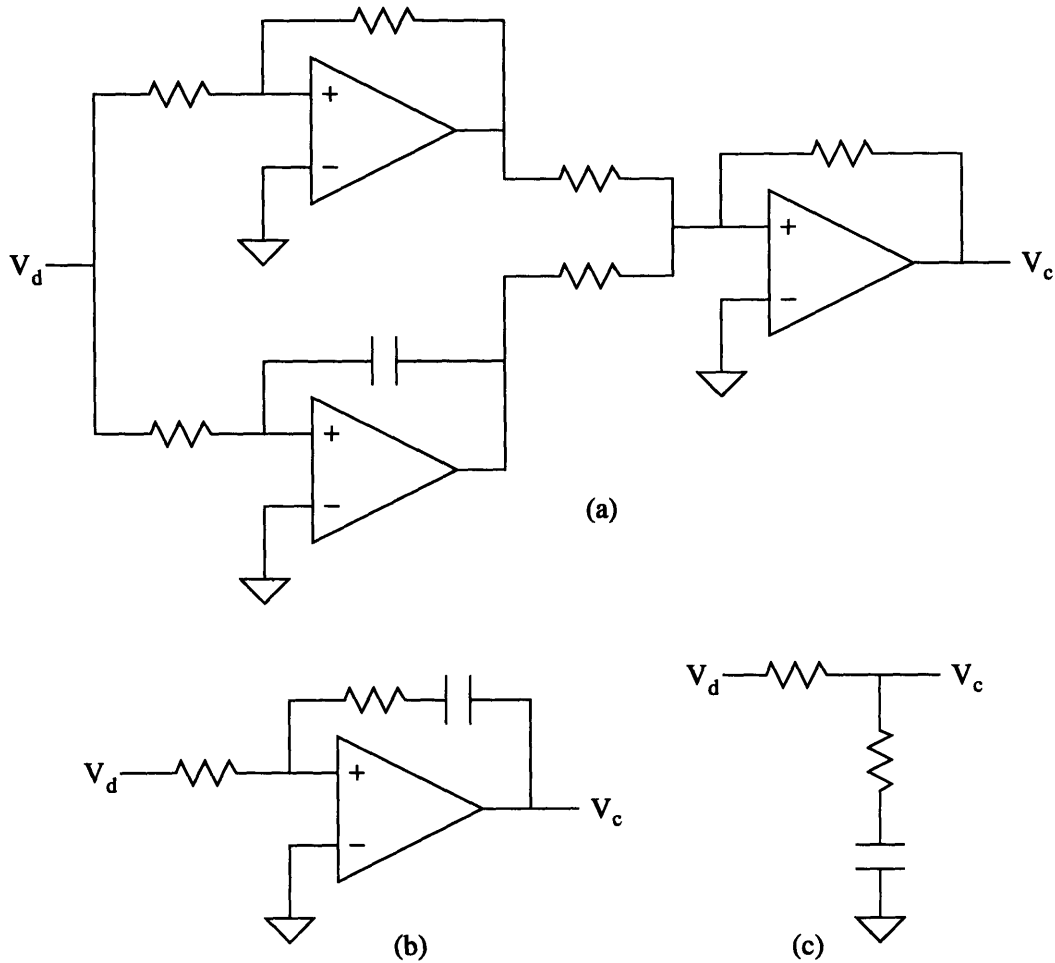


Figure 3.7 - Schematics for an integral plus proportional controller (a), an active loop filter (b), and a passive loop filter (c).

Supposing that the loop filter illustrated in Figure 3.7(b) is now included in the PLL, a new system function for the PLL needs to be found. The transfer function for the loop filter is

$$F(s) = K_h \frac{s + \omega_z}{s} \quad (3.16)$$

The loop filter alters the forward gain of the loop so that its new value is

$$G(s) = K_d K_h K_o \frac{s + \omega_z}{s^2} \quad (3.17)$$

To simplify this equation, a new variable called the loop gain is defined.




Figure 3.8 - Bode plots for PLL with loop filter

$$K = K_d K_h K_o \quad (3.18)$$

This simplifies the forward gain transfer function to

$$G(s) = K \frac{s + \omega_z}{s^2} \quad (3.19)$$

Using this new forward gain, the system function for the PLL becomes

$$H(s) = \frac{G(s)}{1 + G(s)} = \frac{Ks + K\omega_z}{s^2 + Ks + K\omega_z} \quad (3.20)$$

To illustrate the effect of the loop filter, the Bode plots for each of these equations are shown in Figure 3.8.

Looking at the new system function for the PLL and at its Bode plot, a new equation for the bandwidth of the loop can be derived. The new equation for the ω_{-3db} frequency is

$$\omega_{-3db} = K_d K_h K_o = K \quad (3.21)$$

What is important to note here is that by adding a filter circuit into the PLL, rather than simply an attenuation stage with a gain equal to K_h , the DC performance of the loop does not have to be compromised in order to control the loop's bandwidth. Some consideration must be given, however, to where the loop filter's zero is placed. If the zero is placed too close to the -3db frequency of the system function, it will cause unacceptable peaking of the system response. In control theory terminology, this situation corresponds to the system being underdamped. As a general rule, choosing the frequency for the loop filter's zero, ω_z , so that

$$\omega_z < \omega_{-3db}/4 \quad (3.22)$$

will produce acceptable performance. Choosing ω_z such that ω_z equals $\omega_{-3db}/4$ results in a system that is critically damped.

The notation which has been used up to this point, looks at the frequency response of a PLL in terms of K and ω_z . These parameters have the advantage of being easily associated with parameters of the components that make up the PLL. For those who are familiar with control theory, these parameters may be converted so that the behavior of the system may be expressed in terms of natural frequency, ω_n , and damping ratio, ζ . The equations for these conversions are

$$\zeta = \frac{1}{2} \sqrt{\frac{K}{\omega_z}} \quad (3.23)$$

$$\omega_n = \sqrt{K \omega_z} \quad (3.24)$$

As the previous references to the damping of the PLL imply, this conversion is sometimes useful to enable those familiar with control theory to draw on their background. In general, though, the notation used up to this point is considered clearer and more intuitive.

Looking again at the Bode plot of the system function for the PLL which includes an active loop filter, one sees that the gain basically follows the lesser of unity and the forward loop gain, $|G(s)|$. There is, however, some peaking of the system response where the gain exceeds unity. As was mentioned earlier, the extent of this peaking varies with the position of the loop filter's zero. While acceptable performance is generally guaranteed by simply sticking to choosing ω_z according to Equation (3.22), it is possible to calculate the peak value of the system response and the frequency at which it occurs. The equation for the peak value of the system response, H_p , is

$$H_p = [1 - 2\alpha - 2\alpha^2 + 2\alpha(2\alpha + \alpha^2)^{1/2}]^{-1/2}, \quad \alpha = \omega_z/K \quad (3.25)$$

The equation for the frequency at which the peak occurs, called the peaking frequency, ω_p , is

$$\omega_p = \omega_z \left[\left(\frac{2K}{\omega_z} + 1 \right)^{1/2} - 1 \right]^{1/2} \quad (3.26)$$

These equations are taken from Wolaver's book, which also provides a chart of approximations to these equations which might be more useful to those trying to gain insight. [5] This chart is shown in Table 3.1.

Table 3.1 - Approximations for peaking parameters

Damping	ω_z/K	ω_p	H_p
Over	<0.25	$1.2\omega_z^{3/4}K^{1/4}$	$1 + \omega_z/K$
Critical	0.25	$1.4\omega_z$	1.15
Under	>0.25	$\sqrt{K\omega_z}$	$\sqrt{\omega_z/K}$

A final note about the order of the PLL after the adding of the loop filter. Adding the loop filter to the PLL caused a s^2 term to appear in the denominator of the system function. This makes the PLL a second order system. A PLL without a loop filter is a first order system because of the integration that takes place in the VCO. In some PLL's a loop filter with two integrators is added to the PLL, making the PLL a third order system. The

advantage of such a loop filter is that the PLL is better able to track a ramp in frequency by its input signal. This type of loop filter is not used very often, however, because of the extra complexity involved, including that it is not inherently stable like a second order PLL. For these reasons, this type of PLL has not been discussed, but more information on it may be found in Gardner's book for those who are interested. [6]

3.5 Limits on Linear Behavior

This section will discuss the limitations on a PLL's ability to remain in lock while tracking a changing input signal. This section will look at the response of a PLL to various changes in its input signal, including a phase-step, a frequency-step, and a frequency-ramp. A PLL's response to sinusoidal modulation of its input signal will also be discussed briefly. Before covering these areas, however, it is important to point out the absolute limits on steady-state tracking by a PLL.

The fundamental limitation on a PLL's ability to track an input signal is the frequency range of its VCO. Clearly, the VCO of a PLL must be designed so that all expected inputs fall well within its frequency range. This limitation aside, however, since all practical PLL's include at least one integrator in the loop filter, a PLL should be able to track any signal within its VCO's range. The output of the phase detector can be integrated to whatever control voltage is necessary at the input of the VCO. Without this integrator, the PLL's tracking range would be further limited by the range of output voltage possible from the phase detector, as well as any attenuation by the loop filter.

A number of assumptions are made for the remainder of this section. First, it is assumed that all inputs to the PLL remain within its steady-state tracking range. It is also assumed that the PLL starts out fully locked to the input signal. Finally, the analysis focuses on a PLL which includes an active filter of the type shown in Figure 3.7(b). To further simplify the presentation, it is assumed that the zero for this loop filter was chosen to be equal to a quarter of the PLL's bandwidth, making the system critically damped. The error responses that are looked at below would be different for different dampings, however this case should be sufficient to explain the basic behavior.

The key in determining whether a change in the input signal causes the loop to lose lock is to look at the phase error this signal generates within the loop. If the change causes

a phase error which is greater than the linear range of the phase detector, the loop's behavior is no longer linear, and it is considered to have lost lock. To determine if this happens, the transfer function for phase error in the system must first be found, beginning with

$$H_e(s) = \frac{\theta_e(s)}{\theta_i(s)} \quad (3.27)$$

Recalling Equations (3.2) and (3.13),

$$\theta_e(s) = \theta_i(s) - \theta_o(s) = \theta_i(s) - H(s)\theta_i(s) \quad (3.28)$$

$H_e(s)$ may be rewritten as

$$H_e(s) = 1 - H(s) \quad (3.29)$$

Using this equation, the transfer function for phase error may also be rewritten in terms of the forward gain of the PLL.

$$H_e(s) = \frac{1}{1 + G(s)} \quad (3.30)$$

Using this equation, the phase error transfer function for the PLL under consideration may be determined as

$$H_e(s) = \frac{s^2}{s^2 + Ks + Kw_z} \quad (3.31)$$

This transfer function may now be used to find the error signal for any input of interest. Multiplying the Laplace transform of the input signal by this transfer function and then finding the inverse Laplace transform will produce a time domain description of the error signal in the PLL. One may then look at this error signal to see if it exceeds the linear range of the PLL.

The three major changes in the input signal which are considered when looking at a PLL's error response are a step in phase, a step in frequency, and a ramp in frequency. The Laplace transforms for these three input conditions are $\Delta\theta/s$, $\Delta\omega/s^2$, and $\Delta\omega/s^3$, respectively. The response of the example PLL, where $\omega_z = K/4$, has been calculated for

Table 3.2 - Error Response of a PLL

Input	Phase Step	Frequency Step	Frequency Ramp
$\theta_i (s)$	$\Delta\theta/s$	$\Delta\omega/s^2$	$\Delta\dot{\omega}/s^3$
$\theta_e (s)$	$\frac{\Delta\theta s}{s^2 + Ks + K\omega_z}$	$\frac{\Delta\omega}{s^2 + Ks + K\omega_z}$	$\frac{1}{s} \frac{\Delta\dot{\omega}}{s^2 + Ks + K\omega_z}$
$\theta_e (t)$	$\Delta\theta \left(1 - \frac{Kt}{2}\right) e^{-Kt/2}$	$\Delta\omega \cdot t e^{-Kt/2}$	$\frac{\Delta\dot{\omega}}{K^2/4} \left(1 - \left(1 + \frac{Kt}{2}\right) e^{-Kt/2}\right)$
$\theta_e (\infty)$	0	0	$\frac{\Delta\dot{\omega}}{K^2/4}$
$\theta_{e,max} (t)$	$\Delta\theta$	$0.74 \frac{\Delta\omega}{K}$	$\frac{\Delta\dot{\omega}}{K^2/4}$

each of these input conditions, and the results are shown in Table 3.2. Along with the time domain equations which describe the error signal for each of these inputs, this chart also shows the maximum error seen for each of these inputs and the steady state error for each of these inputs. The steady state errors were calculated using the final value theorem.

$$\lim_{t \rightarrow \infty} y(t) = \lim_{s \rightarrow 0} s Y(s) \quad (3.32)$$

As was stated before, in order for the PLL to stay in lock, the maximum phase error seen for a change in the input must not exceed the linear range of the phase detector.

The other input signal which is generally of interest is an input with a sinusoidally modulated frequency.

$$\Delta\omega_i = \Delta\omega \sin(\omega_m t) \quad (3.33)$$

For this input signal, it is useful to find a transfer function relating phase error directly to the change in frequency of the input. This may be done easily by modifying the phase error transfer function.

$$\frac{\theta_e}{\Delta\omega_i} = \frac{\theta_i}{\Delta\omega_i} \frac{\theta_e}{\theta_i} = \frac{1}{s} H_e \quad (3.34)$$




Figure 3.9 - Transfer function for sinusoidal modulation of input frequency.

Thus, for the PLL being considered here, the new transfer function is

$$H_{we}(s) = \frac{\theta_e(s)}{\Delta\omega_i(s)} = \frac{s}{s^2 + Ks + K\omega_z} \quad (3.35)$$

Assuming that K is much larger than ω_z , an approximation of this transfer function is shown in Figure 3.9. Looking at this Bode plot and remembering that the transform of the input will consist of an impulse at both ω_m and $-\omega_m$, one can see that the error signal will simply be a sinusoid whose amplitude is modulated by the gain of this new transfer frequency at ω_m .

$$\theta_e(t) = |H_{we}(\omega_m)| \Delta\omega \sin(\omega_m t) \quad (3.36)$$

In this case, the loop should stay in lock as long as $|H_{we}(\omega_m)| \Delta\omega$ is within the phase detector's linear range.

At this point, the behavior of the PLL within its linear range and the limits of this linear range should be fairly clear. The question now is how the PLL gets to this linear range, both initially and when it has been pushed out by some change in the input signal.

3.6 Frequency Acquisition

The process by which a PLL locks to an input signal occurs in two stages. First, the PLL matches the frequency of its output to the input, and then it matches its phase.

The first stage of this process is called frequency acquisition, and unlike the second stage, it is completely non-linear. While not going deep into the complexity of this topic, this section will describe the basic frequency acquisition process for a PLL which uses a 3-state phase detector and will point out the major design considerations regarding frequency acquisition for this kind of PLL.


Frequency acquisition for a PLL with a 3-state phase detector is considered for two reasons. First, the PLL designed for this thesis uses this kind of phase detector. Second, this type of phase detector, unlike most other phase detectors, is sensitive to frequency as well as phase. For this reason, a 3-state phase detector is also often referred to as a phase frequency detector. PLL's which use other types of phase detectors generally require some additional frequency acquisition circuitry which is not necessary here.

In order to look at the behavior of a 3-state phase detector, a sample architecture for this circuit and a state diagram illustrating its behavior are shown in Figure 3.10. The key feature of this phase detector, which differentiates it from other phase detectors, is that its output is edge-triggered. This feature causes its output to be discrete rather than continuous. Other phase detectors, such as a multiplier, produce a constant voltage that is proportional to the phase error. For the 3-state phase detector, the output for the signal which is being asserted will look like a square wave whose duty cycle is proportional to the phase error. This difference requires a rewriting of the linear equation for a phase detector. Now, it is the average voltage out of the phase detector, \tilde{v}_d , that is proportional to the phase error.


$$\tilde{v}_d = K_d (\theta_i - \theta_o) \quad (3.37)$$

An example set of outputs for this phase detector are shown in Figure 3.11.

As shown in the sample architecture, another feature of this type of phase detector is that there are two outputs out of the phase detector, one for an up signal and one for a down signal. This design therefore requires a different type of the connection to the loop filter, which in most cases means including a charge pump circuit between the phase detector and the loop filter. The charge pump consists essentially of two switches, one connecting the loop filter node to ground, and the other connecting it to the power supply. Assuming that the VCO output frequency increases for increasing control voltages, when



(a)



(b)

Figure 3.10 - Sample architecture (a) and state diagram (b) for 3-state phase detector.




Figure 3.11 - Example set of outputs for 3-state phase detector

the up signal from the PLL turns on, the switch connected to the power supply is turned on. This causes charge to be pumped into the loop filter circuit, charging its capacitor. Conversely, charge is pulled off the loop filter capacitor when the down signal is asserted by turning on the switch connected to ground. As the state diagram shows, the phase detector is designed so that both signals are never asserted at the same time.

Another advantage of this phase detector is its wide linear range. As the output characteristic shown in Figure 3.12 illustrates, this detector has a range of $\pm 2\pi$. The overlapping nature of the output characteristic points out another feature of this phase detector. This phase detector has memory of what the previous inputs were that affects the next output. Also, note that this output characteristic represents a mixed signal of the two outputs, $V_u - V_d$.

This phase detector's sensitive to frequency can be understood by considering again the sample architecture and the state diagram which are shown in Figure 3.10. Assume that asserting V_u causes the output frequency of the VCO to increase, and that V_u will be asserted for cases where the phase of the input is ahead of the phase of the output. Next, consider the case where the frequency of the input signal is higher than the frequency of the PLL's output signal. In this case, the phase detector will see significantly more rising edges from the PLL's input, R^\wedge , than rising edges from the PLL's output, V^\wedge . This means that even if the phase detector starts out in State 1, it will quickly be pushed so that it oscillates back and forth being States 2 and 3. By keeping the phase detector in these two states, only the V_u signal will ever be asserted. This will cause the output




Figure 3.12 - Output characteristic of a 3-state phase detector

frequency of the VCO to increase, pushing the PLL towards frequency lock. Since this system is symmetric, the reverse behavior is true when the output frequency is higher than that of the input signal.

In designing a PLL, there is another critical design consideration that must be kept in mind in regard to frequency acquisition, and that is acquisition time. The amount of time a designer can allow a PLL for this process depends on the application. Looking at a standard active loop filter, such as the one shown in Figure 3.7(b), one sees that there are two components to the VCO control voltage, the voltage across the resistor, V_r , and the voltage across the capacitor, V_c . The voltage across the resistor can change very rapidly, however, it has a limited range. On the other hand, the voltage across the capacitor changes slowly, but is capable of spanning the whole input range of the VCO. From this viewpoint, frequency acquisition may be considered complete when the voltage across the capacitor has been moved close enough to the required control voltage that the resistor can provide the rest. Thus, to determine an equation for acquisition time, one must look at two things, the limit on resistor voltage and the rate at which charge can be moved to or from the loop filter capacitor.

The frequency tuning limit for the loop filter resistor is determined by the phase detector's linear range and the PLL's bandwidth. The equation for this may be derived by stepping through the conditions required to generate the maximum frequency deviation possible from the resistor voltage. The starting equation for this maximum frequency deviation is

$$\Delta\omega_{rm} = K_o V_m \quad (3.38)$$

where V_m is the maximum voltage that may be seen across the loop filter resistor. Remembering that this voltage is proportional to the maximum phase detector output voltage, $V_m = K_h V_{dm}$, the equation becomes

$$\Delta\omega_{rm} = K_h K_o V_{dm} \quad (3.39)$$

where $V_{dm} = K_d \theta_{em}$, so that

$$\Delta\omega_{rm} = K_d K_h K_o \theta_{em} = K \theta_{em} \quad (3.40)$$

Thus for a PLL with a 3-state phase detector, the maximum frequency deviation possible from the resistor is

$$\Delta\omega_{rm} = 2\pi K \quad (3.41)$$

When calculating the acquisition time, this value may be subtracted from the initial frequency error. The remaining frequency error must be compensated for by adjusting the capacitor's voltage.

How fast the capacitor's voltage changes depends on the average current the capacitor sees during the tuning process. Within the PLL built for this thesis, the two output signals from the 3-state phase detector are mixed into a single signal, and this signal is connected directly to the input of the loop filter, rather than indirectly through a charge pump. This means that the current the loop filter capacitor sees is simply proportional to the voltage out of the phase detector divided by the input resistor of the loop filter. A simplified version of the architecture used is shown in Figure 3.13.

The question now becomes what voltages will be seen from the phase detector. For a 3-state phase detector, it is possible to place a conservative bound on what the average output voltage will be during this process. Consider a case where the frequency of the output is slightly higher than that of the input. If these two signals were to stay at the same frequency, the phase error between the input and output would slowly, but repeatedly grow from 0 to -2π . The average phase error would simply be $-\pi$. Thus, the average voltage out of the phase detector would be

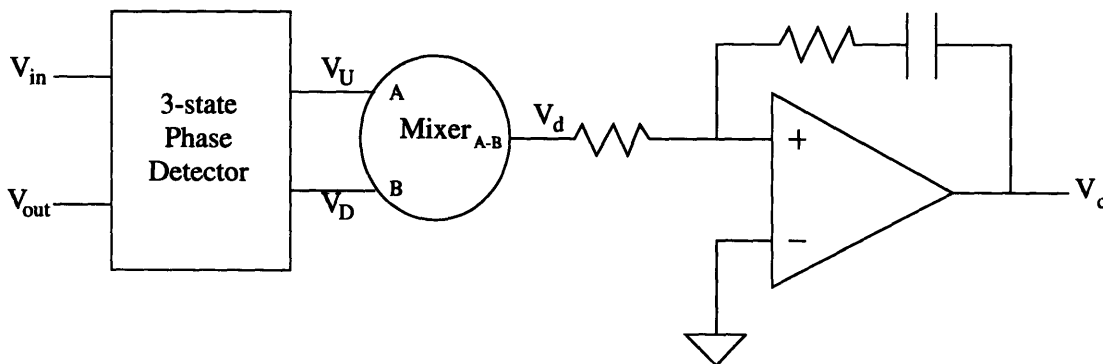


Figure 3.13 - Simplified illustration of phase detector connection to loop filter.

$$\tilde{v}_d = -\pi K_d = \frac{-V_{dm}}{2} \quad (3.42)$$

It can be shown that as the frequency difference between the input and output grows, the average voltage out of the phase detector also grows. To get an idea of how this works, consider an output signal that is more than twice the frequency of the input signal. Since V^{\wedge} edges come at more than twice the rate of R^{\wedge} edges, the phase detector will have to spend at least half its time in State 3 of the state diagram shown in Figure 3.10. The average voltage out of the phase detector works out to be $-3V_{dm}/4$. On the other hand, whatever the frequency error is, the average voltage never goes below $-V_{dm}/2$. Thus, for the analysis here, this value serves as a conservative bound that is sufficient.

The equation for acquisition time may now be derived. First, frequency error within the system is defined via the equation

$$\omega_e = |\omega_o - \omega_i| \quad (3.43)$$

with the initial frequency error in the system being labeled ω_{eo} . The equation for the component of the output frequency due to the capacitor's voltage is

$$\omega_c = K_o V_c \quad (3.44)$$

Assuming that the output frequency starts out above that of the input, ω_c is initially equal to $\omega_i + \omega_{eo}$, while at the end of frequency acquisition it is equal to $\omega_i + \omega_{em}$. Remembering that $dV_c/dt = i/C$ and $i = \tilde{v}_d/R_1$,

$$\frac{d\omega_c}{dt} = K_o \frac{dV_c}{dt} = \frac{K_o \tilde{v}_d}{R_1 C} \quad (3.45)$$

Now substituting for \tilde{v}_d , and remembering that $\omega_z = 1/R_2 C$ and $K_h = R_2/R_1$,

$$\frac{d\omega_c}{dt} = \frac{-\pi K_d K_o}{R_1 C} = -\pi \omega_z K_d K_h K_o \quad (3.46)$$

This simplifies to

$$d\omega_c = -\pi \omega_z K dt \quad (3.47)$$

which may be integrated to

$$\omega_c = -\pi\omega_z Kt + \omega_{e0} + \omega_i \quad (3.48)$$

In this equation, $\omega_{e0} + \omega_i$ is the constant of integration, which was found by looking at what ω_c should equal for $t = 0$. The equation for acquisition time is now found by setting ω_c equal to $\omega_{e0} + \omega_{rm}$, and solving for t .

$$T_{acq} = \frac{\omega_{e0} - \omega_{rm}}{\pi\omega_z K} \quad (3.49)$$

Substituting for ω_{rm} this becomes

$$T_{acq} = \frac{(\omega_{e0}/K) - 2\pi}{\pi\omega_z} \quad (3.50)$$

This equation for acquisition time reveals another trade off in PLL design. A designer interested in designing a PLL with a narrow bandwidth for the purpose of noise rejection must sometimes allow for a wider bandwidth in order to get reasonable acquisition times.

3.7 Final PLL Architecture

In covering static phase error, bandwidth, and acquisition time, this section has highlighted the major system design parameters which must be considered in the design of a PLL. There is one issue which has not been discussed here, and that is the noise performance of the PLL. Rather than including this topic here, the consideration that was given to minimizing noise within each block of the PLL shall be discussed in the sections that describe the circuits which were designed. At the system level, the one main consideration given to noise was to use differential signals as much as possible. Using differential signals minimizes the effect of common-mode noise within the PLL. The final PLL architecture is shown in Figure 3.14. The details of the decision process that went into selecting the circuits for each of the blocks of this PLL are outlined in the following chapters.

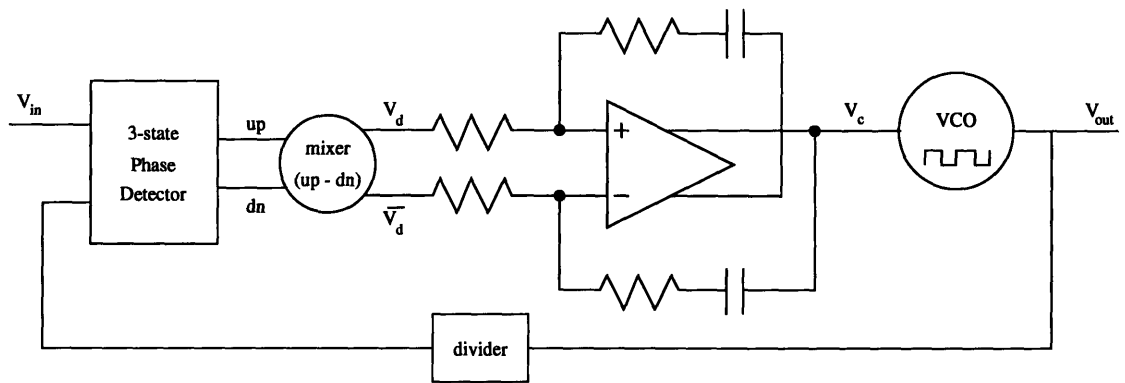


Figure 3.14 - Final PLL architecture

References

- [1] D. H. Wolaver, Phase-Locked Loop Circuit Design, Prentice Hall: Englewood Cliffs, NJ, 1991.
- [2] F. M. Gardner, Phaselock Techniques, John Wiley & Sons: New York, 1979.
- [3] D. H. Wolaver, Phase-Locked Loop Circuit Design, Section 2.8.
- [4] D. H. Wolaver, Phase-Locked Loop Circuit Design, Section 3.3.
- [5] D. H. Wolaver, Phase-Locked Loop Circuit Design, Section 3.5.
- [6] F. M. Gardner, Phaselock Techniques, Section 2.4.

Chapter 4 - The Phase Detector

4.1 Introduction

This chapter examines the details of the phase detector that was built for this PLL. It begins by looking at the alternative architectures that were considered. The reasons for choosing the architecture used here are explained. It then focuses on the details of this architecture including the circuits that compose it. Some of the stages in the evolution of the architecture used are presented here. Finally, the simulated performance of the phase detector is discussed. Some suggestions for improving its performance are made.

4.2 Design Alternatives

In choosing what type of phase detector to use, a designer has a wide variety of choices. A number of fundamentally different architectures exist with widely varying properties. This large collection of circuits has grown from the number of different environments with different demands in which PLLs are used. For one particular application, however, one flavor of phase detector is usually predominant. For example, for frequency synthesis circuits operating in the Gigahertz range, sampling phase detectors and sampling phase frequency detectors have significant advantages. The sampled nature of their output prevents spurious modulation of the VCO frequency [1], while their high operating frequencies allow the PLL to be built without divider circuits that lead to additional phase noise [2]. For applications which do not require such high performance in terms of frequency range, such as the tracking of a FM radio signal, a more traditional circuit such as a multiplier or an XOR gate is generally used for phase comparison. [3,4] These circuits are fairly simple to build and can operate in noisy environments without dramatic degradation of performance.

In clock buffering applications, the architecture of choice is the 3-state phase detector. This phase detector is also frequently referred to as a phase frequency detector because of its most distinguishing feature, its sensitivity to frequency as well as to phase. In contrast, most other phase detectors have very limited sensitivity to frequency, and therefore require that additional circuitry be built with the PLL for frequency acquisition. This has the drawbacks of increased design complexity, additional power consumption,

and larger required die areas. The one other frequency sensitive circuit found in the literature on GaAs phase detectors was the sampling phase frequency detector mentioned above. [5] This circuit, however, had a die area significantly larger than that allowed for the phase detector in the PLL designed here. In addition, the power consumption for this type of circuit, while not stated in the reference article, can be expected to be much too high if it is similar to that of a normal sampling phase detector. [6] Some of these trade-offs can be expected considering that its frequency performance is also over one hundred times greater than that required for this application. One additional advantage of a 3-state phase detector that can be recognized from the chapter on PLLs is its wide linear range of $\pm 2\pi$ that increases the PLL's ability to stay in lock while tracking changes in the input.

The drawbacks of a 3-state phase detector must also be considered though. Unlike other phase detectors, a 3-state phase detector is not sensitive to the actual frequency spectrum content of its inputs, but instead just to the transitions of these signals. These transitions must be representative of this content. Therefore, this phase detector requires a periodic signal. It will not work in an environment such as clock regeneration where the reference signal might not transition for every clock period. 3-state phase detectors also function poorly in noisy environments where noise might cause a false transition to be registered. In clock buffering applications, however, the input and feedback signals will clearly be periodic, and because both of these signals are essentially digital waveforms, there will be a large tolerance for noise. Therefore, the wide linear range, low power consumption, and frequency sensitivity of a 3-state phase detector make it the predominant choice for this application.

There are a number of different architectures to choose from for building this phase detector. One standard architecture that has been in use for over two decades is shown in Figure 4.1(a). This architecture, which is shown here implemented with all NOR gates, can also be implemented using all NAND gates with a similar topology. The all NOR gate version of this architecture consists of two "input-signal" NOR gates, two 3-input "output-signal" NOR gates, a 4-input "reset-signal" NOR gate, and two RS latches. To understand its operation consider the case where the *Ref* signal is leading the *Fdbk* signal. In the initial state, both output signals are low, the internal reset signal is low, the outputs of both latches have been reset, and both input signals are low. When the *Ref* signal transitions




Figure 4.1 - Phase detector architectures: standard combinatorial (a), 4 RS latch (b), dual flip-flop (c), and thesis architecture (d).

high, the *Up* signal will be asserted, but the state of the internal latches will remain the same, and the reset signal will not be triggered. When the *Fdbk* signal transitions high, the *Dn* signal will assert momentarily, but then both output signals will transition low because the reset signal will have been triggered. The reset signal will also cause both RS latches to be set. This will force both output signals to remain low until the latches are reset. Each latch will be reset when the corresponding input signal de-asserts. Once both input signals have de-asserted so that both latches have been reset, the circuit will be back in its

initial state. The difference in the widths of the pulses generated in the *Up* and *Dn* output signals corresponds to the phase difference between the two input signals.

Two additional architectures that were found in the literature regarding this topic are shown in Figures 4.1(b) and 4.1(c). The first architecture is built primarily around four RS latches, while the second architecture shows a simpler method of building this type of phase detector using two flip-flops and an AND gate. The architecture of the phase detector built for this thesis is shown in Figure 4.1(d). The development of the subsequent phase detector architectures shown here has been primarily in response to performance issues surrounding the original all NOR or all NAND architectures. The most significant of these performance issues has been crossover distortion.

Crossover distortion refers to degradation of or anomalies in the output signal of the phase detector for small phase errors. Performance of the phase detector is critical here because it will affect the PLL's performance when it is in lock. The most frequently mentioned crossover distortion is deadband. Deadband refers to a region around zero phase error for which the phase detector produces no output signal. Deadband is detrimental to the performance of the PLL because the phase error is not constrained within this region and will therefore fluctuate freely. This results in increased phase jitter from the VCO output. In practice, the phase detector gain does not have to go completely to zero for this effect to start to play a role, but instead, any significant drop in the phase detector's gain near zero phase error will allow the output's phase to wander more freely.

The old solution to this problem was to simply place a high value resistor between the loop filter node and ground. The resulting voltage leakage off the loop filter node would cause the phase error to be constantly pinned against one edge of the deadband. The drawback of this technique was that it added a constant phase error to the system. This technique has generally not been used in the design of new high performance PLLs. Thus the crossover distortion found in 3-state phase detectors has received more direct attention.

In particular, it has been shown by Gavin and Hickling that crossover distortion in a 3-state phase detector is linked primarily to the rise and fall times of the detector's internal circuits. [7] To illustrate this property, consider the behavior of the traditional

phase detector shown in Figure 4.1(a). Suppose that the reference input is leading the feedback input by some fixed time increment. When the reference signal transitions high, the *Up* output will be asserted two gate delays later. Examining this circuit architecture closely, one sees that when the feedback signal transitions high, the *Dn* signal will be asserted two gate delays later. Three gate delays after the feedback signal transitions high, both output signals will be reset low. The result of this behavior, however, is a pulse the length of one gate delay on the *Dn* signal. The length of the *Up* signal pulse will equal the length of the phase error plus one gate delay. Based on this behavior, the equation for the average signal out of an ideal version of this phase detector is

$$\tilde{v}_d = K_d \frac{T_{up} - T_{dn}}{T_{per}} \quad (4.1)$$

where T_{up} and T_{dn} are the lengths of the pulses seen on the *Up* and *Dn* signals, respectively. For this equation it is assumed that the two input signals are at the same frequency so that T_{per} represents the period of both signals.

Depending on the lengths of the rise and fall times within the circuit relative to the length of the gate delay, this equation may not provide an accurate description of the phase detector's output. First, consider the case where the rise and fall times of the circuit are short compared to the gate delay: $T_r = T_f = T_d/2$. This case is illustrated in Figure 4.2(a) for an arbitrary phase error. Even though the shapes of the *Up* and *Dn* pulses have been distorted in this case, the area under the pulses is the same. This will be the case even for very small phase errors. Thus, the average output voltage under these conditions will be the same as that described by Equation (4.1).

Now consider the case where the rise and fall times are greater than the gate delay. This case is illustrated in Figure 4.2(b) for $T_r = T_f = 2T_d$. Begin by examining the pulse on the secondary output, where the secondary output is defined to be the output which is asserted only in order to reset the phase detector's state. Because the length of this secondary pulse, one T_d , is less than the rise time in the circuit, the maximum amplitude of this pulse will be clipped. At the same time, the length of this pulse at its 50% voltage points will continue to be one T_d . The net result here is that the average voltage content of this pulse has been distorted by the rise and fall times of the circuit. The effect of the rise




Figure 4.2 - Output pulses from two standard combinatorial phase detectors with different rise and fall times.

and fall times on the primary output signal, however, depends on the size of the phase error. If the phase error is large enough, the added length to the primary output pulse will allow sufficient time for the pulse to rise to its maximum value so that its area is not distorted. For small phase errors though, this pulse will be increasingly distorted. Gavin and Hickling have shown that this property causes the average output voltage to have a “square law” characteristic near zero phase error. [8] This characteristic is illustrated in Figure 4.3. The extent of this distortion increases as the ratio of the gate delay to the rise and fall times gets smaller.

It should be noted, however, that the phase detector is not the only possible source of crossover distortion within the PLL. For example, if a charge pump is used, asymmetries in its response can also be a major source of crossover distortion. [9]

The key property of the phase detector design shown in Figure 4.1(b) is that it deliberately extends the length of the dual pulse by its outputs before resetting. This longer minimum length pulse ensures that the rise and fall times do not distort the phase detector’s output signals. In this architecture, there is a two T_d delay after an input signal is asserted before the appropriate output signal is asserted. Once both outputs have been asserted, there is a five T_d delay until they are both reset. Tests of this architecture have shown no detectable crossover distortion. [10] The trade-off, however, is a lower




Figure 4.3 - Illustration of square-law distortion of a phase detector output characteristic.

maximum operating frequency. In order for this phase detector to operate linearly, neither input should be asserted a second time until the phase detector has been fully reset from the previous cycle. If the leading input does assert a second time, it will not be detected, and a 2π shift in the phase detector's output signal will result because the other input signal will now appear to be leading. The limiting case on this behavior will be when the two input signals are 180° out of phase. The seven T_d delay from when the second input asserts to when the phase detector is fully reset leads to a maximum operating frequency of $1/14T_d$. In contrast, in the original architecture, it can be shown that the maximum operating frequency is $1/10T_d$. The literature on this topic also mentions phase detectors with $1/8T_d$ maximum operating frequencies, but does not reveal the associated architectures. With the introduction of "quasi-combinatorial" architectures, such as the one in Figure 4.1(c), this $1/8T_d$ figure may be misleading, not representing a pure combinatorial architecture.

The primary advantages of the phase detector architectures shown in Figure 4.1(c) are its size and simplicity. Unlike the classic architecture and the four RS latch architecture which require 9 and 10 standard cells, respectively, to construct, this architecture requires only 5 cells, two for each flip-flop and one for the AND gate. In

addition, the behavior of this phase detector is conceptually simpler. Consider the case where the reference signal is leading the feedback signal. When the reference signal transitions high, the associated flip-flop is clocked, causing the high signal on its data input to be passed through to its output. This corresponds with the *Up* output signal being asserted. The delay before the output transitions is equal to the flip-flop *clock-to-Q* delay. This signal will now stay high regardless of what the reference signal does until the phase detector's state is reset. Next, the feedback signal clocks the second flip-flop, causing the *Dn* output to be asserted. When the *Dn* output asserts, the AND gate will cause both flip-flops to be reset, thus resetting the state of the whole phase detector. The length of time for which both outputs will be asserted is equal to one gate delay plus the flip-flop *reset-to-Q* delay. Assuming the same limiting case as for the previous architectures, the maximum operating frequency for this phase detector will be

$$F_{\max} = \frac{1}{2(T_{\text{clk-q}} + T_d + T_{\text{rz-q}})} \quad (4.2)$$

The flip-flops used with this design must accept an asynchronous reset signal. Also, the type of flip-flop used will generally have an all-overriding reset, meaning that its output can not be asserted as long as the reset signal is high.

The phase detector which was built for this thesis represents an attempt at a new architecture for phase detection with minimal deadband. This architecture was developed by Ray Sundstrom, a member of the BiCMOS design team at Motorola that was working to develop a comparable PLL design. The key feature of this new architecture is that the designer can easily control the length of the dual output pulse by adjusting the delay through the reset signal delay paths found in the circuit. Choosing to build this phase detector was essentially an experiment to see what kind of performance could be achieved with it. The decision was backed up by the fact that another member of the GaAs design team was building a second phase detector using the standard all NOR gate architecture.

To explain this phase detector's behavior, the output signals for several ranges of phase errors are discussed here. First, consider the case where there is zero phase error, meaning that the PLL is in both frequency and phase lock. In this case, this phase detector will produce two synchronous output pulses of equal length. The input signals will

simultaneously clock the two respective flip-flops, causing the flip-flop data inputs to be feed through to the circuit outputs. Then, rather than being reset by a feedback signal from the outputs, the two input signals, after a set delay, will appear at the opposite flip-flop from the one they clocked and reset the output signals. Assuming that the *clock-to-Q* delay and *reset-to-Q* delay are roughly the same, the lengths of the output pulses will be equal to the length of the reset signal delay, T_{delay} . The exact length will be

$$T_{\text{pulse}} = T_{\text{delay}} + T_{\text{rz-q}} - T_{\text{clk-q}} \quad (4.3)$$

If the reference signal now starts to move in front of the feedback signal, both edges of the *Up* signal will move out, while both edges of the *Dn* output pulse move in. The *Up* signal will start before and end after the *Dn* signal. The length of the new *Up* pulse will equal the reset signal delay plus the length of the phase error, while the *Dn* pulse length will equal the reset signal delay minus the length of the phase error. When the length of the phase error exceeds the flip-flop *clock-to-Q* delay, however, the *Dn* signal will disappear completely. This is because the associated flip-flop's data input will have transitions low before the feedback signal clocks it. Illustrations of the phase detector's outputs for the three ranges described here, zero phase error, phase error $> T_{\text{clk-q}}$, and phase error $< T_{\text{clk-q}}$, are shown in Figure 4.4.

Now consider the issue of crossover distortion in this phase detector. For zero phase error, the length of the reset signal delay is chosen deliberately to eliminate any crossover distortion due to rise and fall times. When a phase error appears in the system, one pulse's length will increase, while the other pulse either shrinks or disappears. For the pulse whose length increases, rise and fall times will continue not to produce any distortion. The minimum length of the other pulse before it disappears is $T_{\text{delay}} - T_{\text{clk-q}}$. Thus, if the designer wants, the length of the reset signal delay may be chosen so that rise and fall times do not produce unacceptable distortion of this minimum length pulse either.

The disappearance of this pulse for phase errors greater than $T_{\text{clk-q}}$, however, functions as a second source of crossover distortion not found in the other phase detector architectures. In this case, when this signal disappears, there will be a jump in the average voltage output of the phase detector. Above this jump, the gain of the phase detector is also reduced to roughly half of its original value. The gain is smaller because there is only

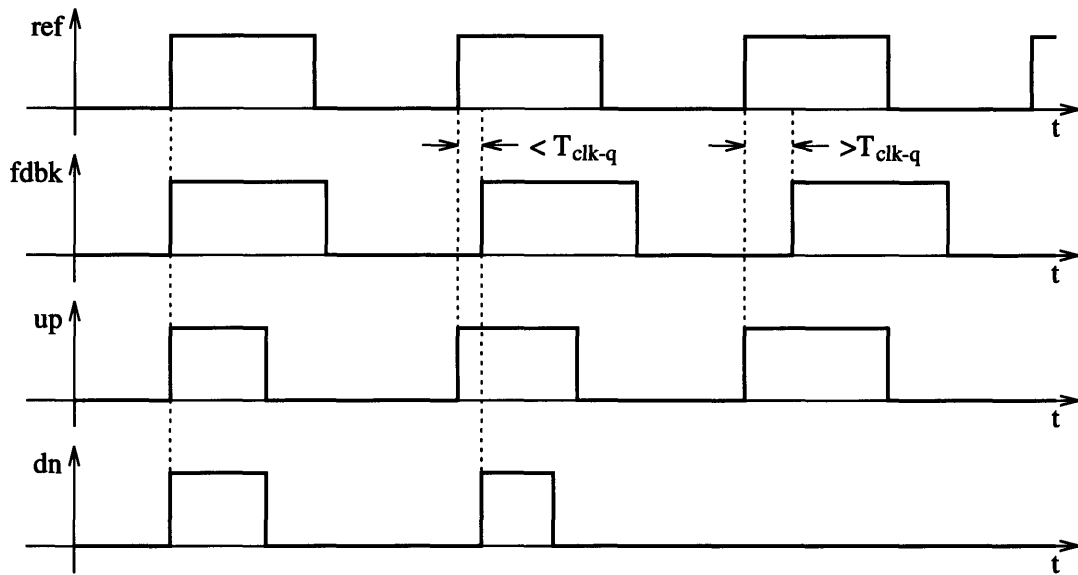


Figure 4.4 - Illustration of phase detector outputs for three phase error regions: zero phase error, phase error less than Clk-to-Q delay, and phase error greater than Clk-to-Q delay.

one pulse changing size now. In order to extend the phase error range over which the jump in the average output voltage takes place, it may be desirable to choose a reset delay length that is short enough that the pulse on the secondary output is deliberately compressed by the rise and fall times as it approaches its minimum length before disappearing. This way, when the pulse does disappear, the resulting jump in the average output voltage will not be as large. The shape of the crossover distortion which this behavior produces is shown in Figure 4.5. This figure is a close-up of the simulated output characteristic of the actual circuit designed for this thesis. It is important to note that this crossover distortion is symmetric about the zero phase error point. Asymmetries about this point are even more detrimental to the performance of a PLL because of small signal stability issues. [11]

4.3 Phase Detector Circuits

This section will look at the details of the circuits which compose the phase detector built for this thesis. A more detailed top-level schematic of the architecture that was used is shown in Figure 4.6. First, the function and circuitry of the initial input signal buffers are described. Next the motivations behind the delay line topology will be explained, and the delay buffer circuitry will be shown. Then the requirements on the flip-




Figure 4.5 - Graph of small phase error crossover distortion for thesis phase detector architecture.

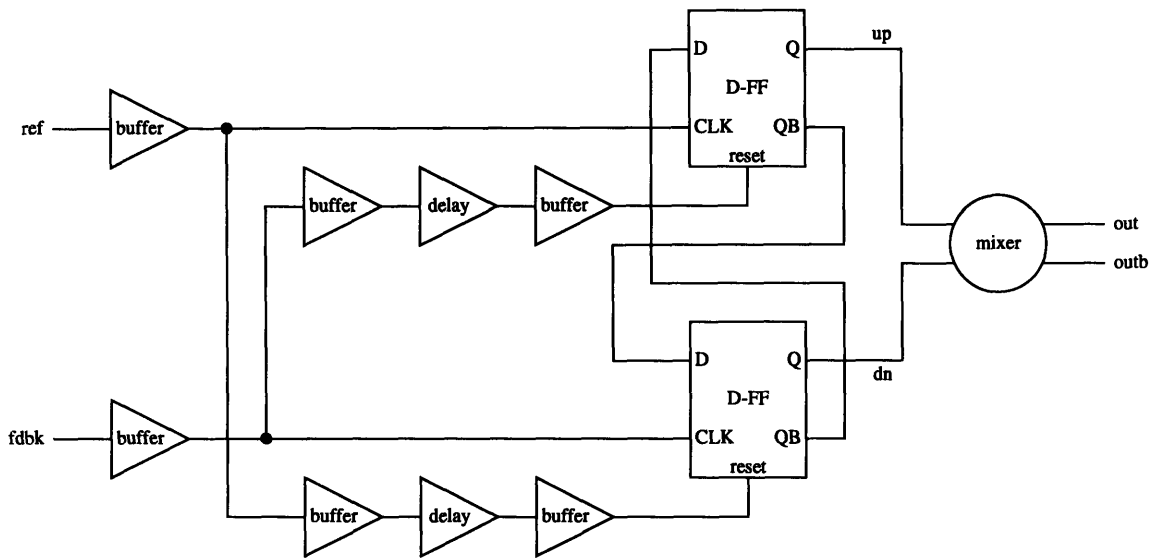


Figure 4.6 - Full architecture of final phase detector built for thesis.

flops' behavior will be described, along with the circuit used. Finally, this section will look at the functionality of the mixer circuit that was added to this architecture and the motivations for adding it.

The function of the initial input buffer that each input signal passes through is to ensure the integrity of the signals passed to the rest of the circuit. These buffers will clean up the input signals and provide output signals at the right voltage levels and with acceptable rise and fall times. The circuit used here closely resembles the standard SCFL gate shown in Chapter 2. The schematic for this circuit is shown in Figure A.1 in Appendix A. The single resistor, R6, included above the two normal load resistors is placed there to ensure that the top transistors in the source follower stages always stay in saturation. The diodes in the source follower stages are used for level shifting of the output signals. Otherwise, this circuit has the same properties as the standard SCFL cell described back in Chapter 2.

The next block of interest within the circuit is the delay line. The delay line is composed of a delay gate surrounded on both sides by buffer gates. These buffer gates, which are identical to the buffers on the input signals, are used to isolate the delay gate from the rest of the circuit. The front buffer is used so that the signals out of the input buffers are not degraded by the loading of the delay gate. The buffers on the delay gate outputs ensure that the reset signals running to the flip-flops have adequate rise and fall times. The topology of the delay gate circuit is identical to that of the buffers except for two modifications of the input connections. This topology is shown in Figure A.2 in Appendix A. The first modification to this circuit is the addition of capacitors between the gates of the input transistors and the switching nodes. These capacitors act as clamps on the switching nodes and effectively increase the gate-drain capacitance of the input transistors. The result is that more charge must be moved either to or from the input gate nodes before the circuit can switch states. The second modification of the circuit is the placing of resistors in the input signal paths. These resistors further restrict the current flow to the input gate nodes. The net result of these modifications is the addition of an external RC delay to the switching time of the circuit. By varying the values of these components, the designer can now adjust the switching time of the circuit to whatever value is required. The resistor and capacitor values that are used in this circuit were chosen, based on simulations of the whole delay line, to fix the reset delay at 500ps.

The flip-flops are the next major block of interest within the circuit. Two main requirements are made on the flip-flops in order for this phase detector to function

properly. The first is that they have an asynchronous reset signal, since the reset signal must not depend on the clock signal in this phase detector architecture. The second requirement is that the reset signal not be all-overriding. Once the reset signal has caused the output to de-assert, it should not prevent the output from being asserted again if the circuit is clocked again before the reset de-asserts. In an all-overriding circuit, the output can not be asserted until the reset signal de-asserts. Unfortunately, this second requirement was not understood at the time the flip-flops were designed, so that the circuit that was built does have an all-overriding reset. As the next section on the simulated performance of the phase detector will show, using this all-overriding reset design limits the linear range of the phase detector to $\pm\pi$. It also interferes dramatically with the frequency sensitivity of the phase detector. It was shown in behavioral simulations that a PLL using this phase detector with the wrong type of flip-flop will not always acquire the input signal's frequency. Instead, the output signal sometimes becomes trapped within a small range of frequencies where it oscillates, so that it never reaches the frequency of the input signal. These results are discussed in more detail in Chapter 7, which discusses the behavioral modeling work.

The circuit topology of the flip-flop that was used is fairly standard for SCFL design, and is shown in Figure A.3 in Appendix A. Like the buffer and delay gates, this design also uses diodes in the source follower stages for level shifting and an additional resistor above the normal load resistors to ensure that the top source-follower transistors stay in saturation. The major design challenge for this gate, which required several design iterations, was to ensure that all transistors would stay in saturation for a power supply voltage of 3V. Also, this circuit design went through one revision where the currents within the circuit were doubled to improve the rise and fall times of the circuit. The final rise time for the circuit in simulations was 500ps for a 420mV output swing, while the final fall time was 350ps for the same voltage swing.

The last major block of this phase detector is the mixer. This block was not included in the original architecture of this phase detector, but was added because of concerns that became apparent later. The major concern prompting the adding of this mixer circuits resulted from problems with the connection of this phase detector to the loop filter. At this point, a decision had already been made to use an active loop filter,

rather than a passive loop filter together with a charge pump. The first problem with the connection between these two blocks is that the active loop filter only allows for two input signals. This meant forfeiting the differential output signals from the flip-flops, which would open up this connection to common-mode noise within the PLL. In addition, any DC offset between these two signals, due to normal processing related variations, would require a static phase error within the PLL to prevent the loop filter from integrating the offset signal. Adding the mixer eliminates both of these problems. The output signals from the mixer depend on relative comparisons of Up to Upb and Dn to Dnb , not the absolute DC levels of the two sets of signals. Also, the inputs and outputs of this circuit are both differential, preserving the associated common mode noise rejection properties for this connection to the loop filter.

The topology for the mixer circuit is shown in Figure A.4 in Appendix A. Looking at this topology, one sees two sets of switching gates connected to the same set of switching nodes. The currents pulled through the two sets of switching gates are identical. One set of switching gates is controlled by the Up and Upb signals, while the other set is controlled by the Dn and Dnb signals. The output of this mixer is essentially a subtraction of the Dn signal from the Up signal. The equation describing this behavior is

$$V_{out} = K_m (V_{up} - V_{dn}) \quad (4.4)$$

where K_m is the gain of the mixer. In order to explain this behavior, the details of the circuit topology must be examined. If the Up signal is higher than the Upb signal, more current will flow through the switching node that sets the $Outb$ signal. Unless the Dn signal counters this action by pulling the same increased ratio of current through the other switching node, the voltage of the Out signal will increase, while the voltage of the $Outb$ signal decreases. This output condition, where the Out signal is above the balance point for the Out and $Outb$ signals, can be interpreted as a positive result from the subtraction. Conversely, if the Dn signal is pulling more current through its switching node than the Up signals is pulling through its associated node, then the $Outb$ signal will increase while Out decreases. This output condition parallels a negative result from the subtraction.

As Equation (4.4) implies, the addition of this mixer to the phase detector produces drastic changes in the attributes of the phase detector's output signals. When in lock, this

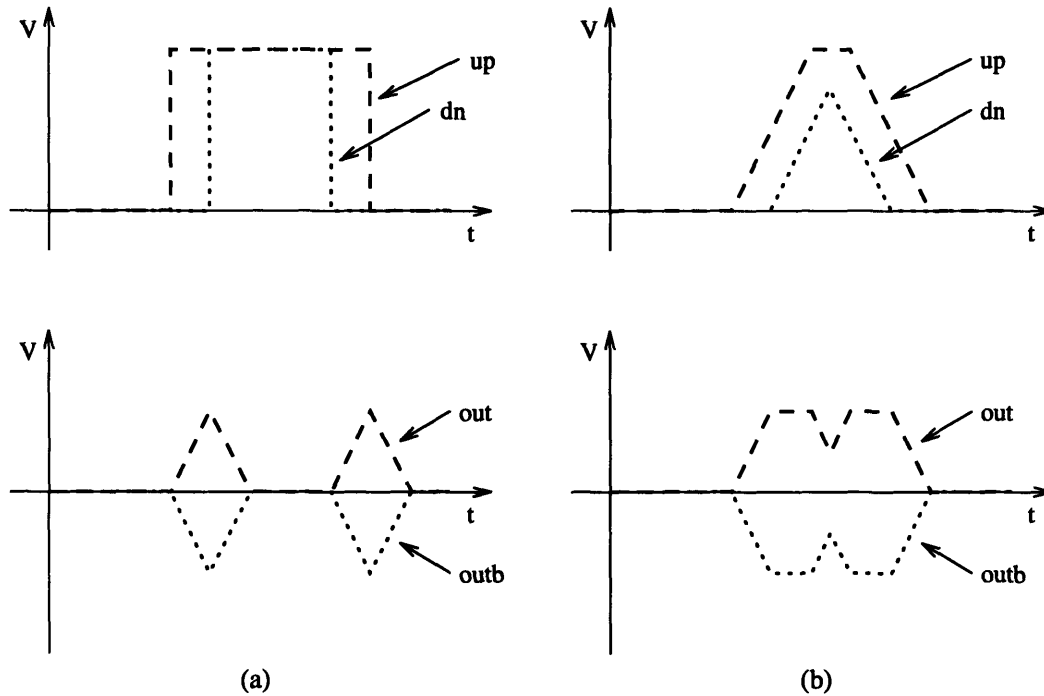


Figure 4.7 - Illustration of mixer outputs with rise and fall times taken into account for idealized (a) and actual (b) versions of up and dn signals from original portion of phase detector.

phase detector will now not produce any output signal. For large phase errors, where the pulse on the second output signal has disappeared, two symmetric, opposing pulses will be produced. For small phase errors, however, where the second pulse still is present, each mixer output will be composed of two short, abrupt pulses. This behavior results because the mixer only produces an output signal when there is a difference between the two input signals, which in this case will be limited to a short region before and after the secondary output pulse. Since these output pulses will be short, the issue of whether the rise and fall times might distort the output signals arises again. Fortunately, however, when considering the actual output of the original portion of the phase detector, this turns out not to be a problem.

Figure 4.7 illustrates an idealized version of the output from the original portion of the phase detector and the actual output for the circuit which was designed. Looking at the idealized output one sees that rise and fall times would produce a serious distortion of the desired output from the mixer. Looking at the actual output from the original portion of the circuit, however, one sees that the effect of rise and fall times on this initial output

prevent the mixer output from being distorted. For the actual output, at no point is there a large difference between the two signals. Instead, the difference between the two signals is spread across the length of almost the entire pulse. The result is that the voltage swing required by the output of the mixer is not as large, and there is also more time for the desired levels to be reached. Because of these properties, the rise and fall times do not produce an unacceptable amount of additional distortion. Also, it should be noted that the diverging output pulses from the mixer are better than the original Up and Dn signals at moving the voltage of the loop filter. The addition of the mixer thus produces several improvements in the performance of the phase detector.

4.4 Simulated Performance

The circuits which constitute the phase detector built here were all designed with the aid of HSPICE for simulating their performance. Once all the component circuits were designed and functioning well, the whole phase detector was also simulated with HSPICE. The primary characteristic that was looked at in these full circuit simulations was the average output voltage versus the phase error. These simulations consisted of feeding two pairs of differential square signals, at the proper voltage levels and with adequate rise and fall times, into the phase detector's two inputs. The two sets of square wave signals were offset by a delay that was varied to cover the whole range of phase errors that were of interest. The outputs of the phase detector were observed for each delay increment and the average voltage signal out of the phase detector was calculated using an HSPICE averaging function. The results of these simulations are shown in Figure 4.8. Based on these results, the average gain of this phase detector over the majority of its linear range is approximately 0.18 V/rad.

Unfortunately, the full circuit simulations revealed a problem with the circuit which had already been designed and laid out. As Figure 4.8 shows, the range of this phase detector is only from $-\pi$ to π , while the normal range for a 3-state phase detector should be $\pm 2\pi$. A thorough analysis of the phase detector's behavior revealed that this problem was originating from the type of flip-flop which was being used in the circuit. The particular problem with this flip-flop was that it had an all-overriding reset signal. To understand how this property limits the phase error range to $\pm\pi$ consider the case where

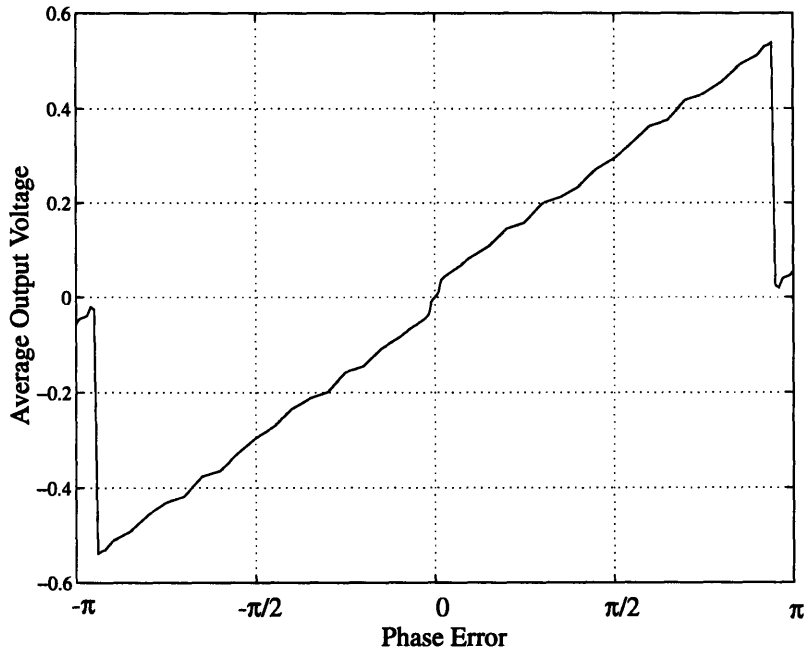


Figure 4.8 - Final simulated phase detector output characteristic (50MHz inputs).




Figure 4.9 - Illustration of how the all-overriding flip-flop will restrict of phase detector's linear range.

the reference signal is leading the feedback signal by an increasing phase margin. Also, for now, assume that the reset signal delay is insignificant compared to the periods of the two sets of input signals. In this case, the reset signals correspond almost exactly with the clock signals seen by the opposite flip-flops. An illustration of this situation is shown in Figure 4.9.

Initially, the feedback signal de-asserts before the next rising edge of the reference signal. As the phase error between the two inputs continues to increase, however, the feedback signal stays asserted through the rising edge of the reference signal. This corresponds to the “*Up*-output” flip-flop’s reset signal staying asserted while the reference signal tries to clock it. Because the flip-flop has an all-overriding reset, this will prevent the flip-flop’s output from being asserted. On the next feedback signal rising edge, the “*Dn*-output” flip-flop will still see an asserted data input because the *Up* signal was held low, thus leading to a 2π shift in phase detector’s output signal. While this problem was identified, a new flip-flop design was not completed because of time constraints.

Looking closely at the phase detector output characteristic shown in Figure 4.8, one sees another distortion in the characteristic at the edges of the linear range. This deadzone in the phase detector gain is caused by the reset delay. While this delay was neglected in the analysis above because it is small compared to the periods of the input signals that were considered, it has an important effect on the phase detector’s performance. Assuming the same input conditions as above, consider how this delay modifies the phase detector’s behavior. The reset signal to the “*Up*-signal” flip-flop trails the feedback signal by the length of the reset delay. This means that there must be an additional delay between when the feedback signal de-asserts and when the reference signal transitions high for the “*Up*-signal” flip-flop to be clocked without the reset signal overriding its input signal. On the other hand, the *Dn* signal can not be asserted under these conditions either because, regardless of the length of the reset delay, the reset signal to the “*Dn*-signal” flip-flop will still be asserted when the feedback signal transitions high. The waveforms for this situation are illustrated in Figure 4.10.

It is important to note that this deadzone region is the result of a fixed circuit parameter. It will cover a fixed amount of time around $\pm\pi$ phase error, as opposed to an amount which is relative to the periods of the inputs. This means that as the frequency of the input signals increases, this deadzone will become more problematic. In any case, this deadzone is intolerable because of the problems it can cause during the frequency acquisition process. The details of this situation will be described in the chapter on behavioral modeling. Fortunately, though, using the correct type of flip-flop will remove this problem. With this change, the only drawback of the reset delay is that the edges of




Figure 4.10 - Illustration of deadzone caused by the reset signal delay.

the phase detector's linear range, which are now ideally at $\pm 2\pi$, will be pushed in by the amount of the reset delay. This behavior is acceptable, as none of the architectures considered here has a linear range that actually extends all the way out to $\pm 2\pi$.

The final performance concern for the phase detector was its frequency range. The range of output frequencies for the PLL being built was 200-500MHz, with this signal being divided down by eight before being fed back into the phase detector. This meant that the frequency of the input signals to the phase detector would be between 25MHz and 62.5MHz. The majority of the tests of the phase detector's behavior were done at 50MHz. To ensure robustness, however, the phase detector and all of its internal circuits were tested to 80MHz with no degradation of the detector's or any of the subcircuits' performance. The maximum frequency range of the phase detector was not tested. To provide some examples of the phase detectors actual simulated output, the results of several simulations are included in Appendix B. Specifically, the *Up*, *Dn*, *Out*, and *Outb* signals are shown for 50MHz inputs at a number of different phase offsets.

References

- [1] D. H. Wolaver, *Phase-Locked Loop Circuit Design*, Prentice Hall: Englewood Cliffs, NJ, 1991.
- [2] J. Andrews, A. Podell, J. Mogri, C. Karmel, K. Lee, "GaAs MMIC Phase Locked Source," *IEEE MTT-S Digest*, pp 815-818, 1990.

- [3] F. M. Gardner, *Phaselock Techniques*, John Wiley & Sons: New York, 1979.
- [4] D. H. Wolaver, *Phase-Locked Loop Circuit Design*, Sections 4.1-4.6.
- [5] T. Nakagawa, T. Hirota, and T. Ohira, "A Novel MMIC Approach to Sampling Phase/Frequency Detection for Microwave Quartz-Lock Oscillators," *IEEE MTT-S Digest*, pp 1301-1304, 1992.
- [6] J. Andrews, et al., "GaAs MMIC Phase Locked Source."
- [7] D. T. Gavin and R. M. Hickling, "A PLL Synthesizer Utilizing a New GaAs Phase Frequency Comparator," *Proceedings of RF Technology Expo 88*, Cardiff Publishing: Englewood, CO., pp 457-71, 1988.
- [8] Gavin and Hickling, "A PLL Synthesizer."
- [9] Gavin and Hickling, "A PLL Synthesizer."
- [10] Gavin and Hickling, "A PLL Synthesizer."
- [11] Gavin and Hickling, "A PLL Synthesizer."

Chapter 5 - The Voltage-Controlled Oscillator

5.1 Introduction

This chapter focuses on the voltage-controlled oscillator (VCO) circuit built for this PLL. It begins by describing the performance criteria that are associated with VCOs. This is followed by an explanation of the reasons, based on these criteria, for choosing to use a ring oscillator from among the possible design alternatives. Then, in a slight departure, the effects of the divider circuit placed in the feedback path from the VCO are discussed. Next, the details of the ring oscillator circuits that were designed are presented. The following section describes the performance attained in simulations for the architecture selected. A final section will look at the special attention given to eliminating noise in the design and layout of these circuits.

5.2 VCO Performance Criteria

There are a number of important performance measurements of VCOs, including frequency range, linearity, gain, modulation bandwidth, spectral purity, and sensitivity to external interference. The frequency range specifies the minimum and maximum frequency output signals a VCO can produce. In general, a VCO will have some maximum frequency and possibly some minimum frequency at which it stops oscillating. The frequency range, however, will normally be a subset of this range where the gain and linearity of the VCO are acceptable. One measure of linearity is how much the VCO gain deviates over a specified frequency range. Unlike the VCO characteristic shown in Chapter 3, the output frequency versus input voltage characteristic of a real VCO circuit will not be perfectly linear. Linearity is usually specified as a percentage deviation from a specific VCO gain. Linearity is desirable because it makes the performance of the PLL more predictable and, thus, simplifies designing the PLL. As specified before, gain refers to the change in the VCO's output frequency that results from a change in its input voltage. Gain is thus reflected in the slope of the VCO transfer characteristic.

Modulation bandwidth refers to a property of VCO gain that was not described in the section on VCOs in the earlier chapter on general PLL behavior. If the input signal to a VCO is modulated at an increasing frequency, but with the same magnitude, the




Figure 5.1 - Effect of modulation bandwidth of VCO gain.

response of the VCO, in terms of the amount the output frequency changes, will start to decline above a certain frequency of modulation. Modulation bandwidth describes the frequency at which the VCO gain starts to decline. Taking this behavior into account, the new equation for VCO gain becomes

$$\frac{\omega_o}{V_c} = \frac{K_o}{1 + \frac{s}{\omega_{mb}}} \quad (5.1)$$

where K_o is the low frequency VCO gain and ω_{mb} is the modulation bandwidth. The effect of ω_{mb} on the VCO gain is shown in Figure 5.1. Since this behavior adds another pole to the PLL transfer function, it is important that this pole be located significantly above the PLL's unity gain frequency, or it may cause instability problems. A general rule to avoid problems with this pole is to make sure that ω_{mb} is at least four times greater than the PLL's bandwidth, K .

Spectral purity is a description of the quality of the output signal from the VCO. Ideally, the frequency spectrum of the output of a VCO would be a δ -function distribution at the desired frequency. In a real VCO, however, this distribution will have some finite width. The wider this distribution is, the more phase noise that will appear in the VCO's output signal. A common measure of the spectral purity of an oscillator is its quality factor, Q . Q is a measure of the output pulse's width, relative to its center frequency, at 3db below the maximum amplitude of the pulse.

$$Q = \frac{\Delta\omega_{-3\text{db}}}{\omega_c} \quad (5.2)$$

The higher the Q of an oscillator, the better its spectral purity is.

The last criteria for VCOs is susceptibility to external interference. This characteristic is used to describe two parameters of a VCO: its immunity to noise and its susceptibility to injection locking. These parameters are grouped together because they are generally linked. A VCO with good immunity to noise will generally not be very susceptible to injection locking and vice versa. In the scope of the work for this thesis, consideration of injection locking was neglected, though it is an important characteristic and should be considered in a fully optimized VCO circuit. Attention was paid, however, to making the VCO less susceptible to noise. The specific steps which were taken to prevent noise in the VCO are described later in this chapter.

In most cases, there are trade-offs between these different performance criteria. For example, there is always a trade-off between frequency range and linearity. The percentage of the VCO's total frequency range which is considered usable depends on the linearity requirement. Loosening this requirement will extend the frequency range of the VCO. Another trade-off is between gain and sensitivity to external interference, such as noise in the input voltage signal. Noise on this node will obviously produce more phase jitter in a VCO with high gain. In addition, because the gain of a VCO is tied closely to its frequency range, larger frequency ranges will also generally lead to higher sensitivity to external interference. This effect can be countered by increasing the range of the input voltage, but the extent to which this can be done is significantly limited by what input voltage ranges are possible when operating at the available supply voltage.

5.3 Design Alternatives

In the process of deciding what type of VCO circuit to design for this PLL, five major design alternatives were considered: crystal oscillators, resonant oscillators, multivibrators, ring oscillators, and delay lines. The last alternative listed here, a delay line, represents a deviation from the normal concept of a PLL. In this approach, the PLL does not generate its own output signal. Instead, the input signal is simply fed through a series of delay gates that buffer the signal. Then, before the signal is output, its phase is

compared to the phase of the input. If there is a phase difference, the delay through the delay line is changed in order to align the phase of the two signals. The advantage of this design is improved loop stability. This approach only works, however, when the input signal is precisely the same signal one wants to see at the output.

In order to select which alternative to pursue, the characteristics of each were compared with the performance requirements of the PLL being designed. While crystal oscillators and resonant oscillators are capable of very high frequency ranges and both have better spectral purity than the other alternatives, their limited frequency tuning ranges rule them out for this application. Delay lines were ruled out because of their requirement that the output signal frequency be the same as that of the input. The specifications for the circuit being designed required that output signals be available at both twice and four times the frequency of the input signal. This type of frequency synthesis is only possible using a regular VCO circuit together with a divider in the feedback path to the phase detector. Both multivibrators and ring oscillators, however, are capable of the performance required for this application. The decision to pursue a ring oscillator design was made due to the fact that another member of the design team was pursuing the multivibrator alternative.

5.4 Effects of Using Divider Circuits

As the previous section indicated, the PLL built for this thesis includes a divider in the feedback path from the VCO to the phase detector. There is also a divider that the main PLL output signal runs through before being made available externally. This second divider can be set to divide either by two or by four, while the divider in the feedback path is fixed at divide by eight. These two dividers have different functions. The divider on the output signal is used to buffer the output signal from the VCO, while the divider placed in the feedback path is required for the PLL to do frequency synthesis at a multiple of the input signal frequency. The position of these dividers within the overall PLL architecture is illustrated in Figure 5.2. In addition, one simple method for implementing these dividers is also shown in Figure 5.2.

The divider used on the output signal is important because the original VCO output signal generally has asymmetries which would be unacceptable. In particular,

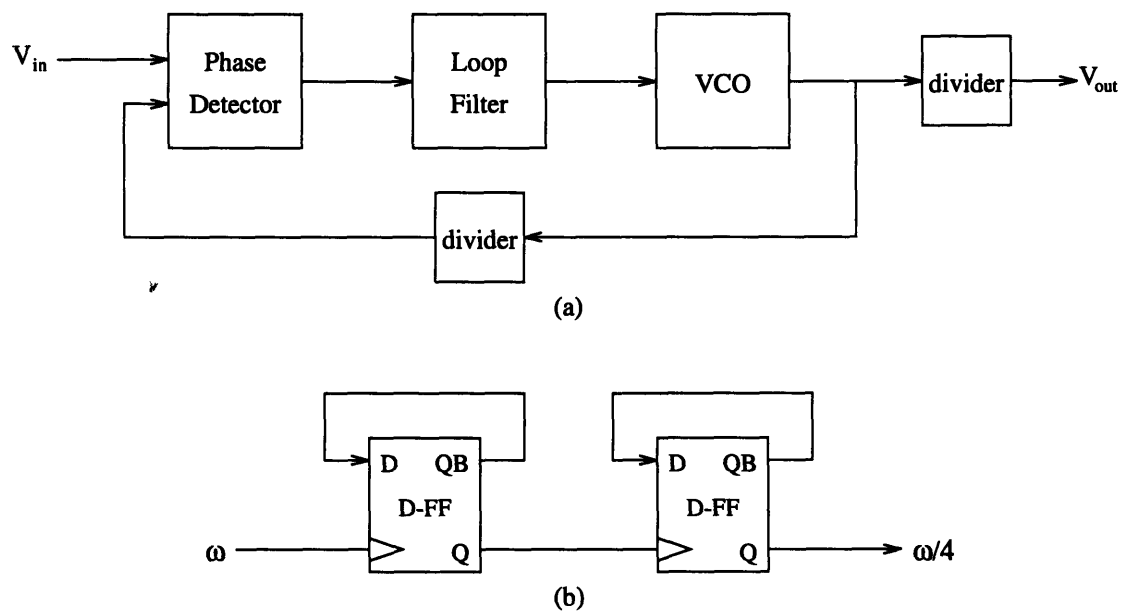


Figure 5.2 - PLL architecture with dividers (a) and an example of divide by four divider (b).

dividing the VCO output signal by two guarantees that the final output signal has a 50% duty cycle. Designing a VCO circuit which has a symmetric output signal with a 50% duty cycle would be a substantially more difficult design challenge. The trade-off, however, is that the VCO must generate its output signal at twice the frequency of the final output signal. The option of dividing by four is included to broaden the application range of the chip to lower frequency applications.

As was stated above, the feedback path divider is necessary for frequency synthesis. There are several motivations for adding this functionality to this PLL. First, the output signal conditioning process just described requires the VCO output signal to be at twice the desired final output frequency. Generating a VCO output signal at greater than twice the frequency of the input signal, however, has an additional advantage; it allows the main clock signal, which is distributed globally within the electronic system, to be at a lower frequency, since it can be boosted to the desired frequency locally. This allows a lower cost crystal to be used for clock generation, reduces RF radiation, and simplifies board design. [1] To understand the frequency synthesis process, consider that both inputs to the phase detector must be at the same frequency for the PLL to be in lock. Thus, the output frequency of the VCO must be N times the frequency of the reference

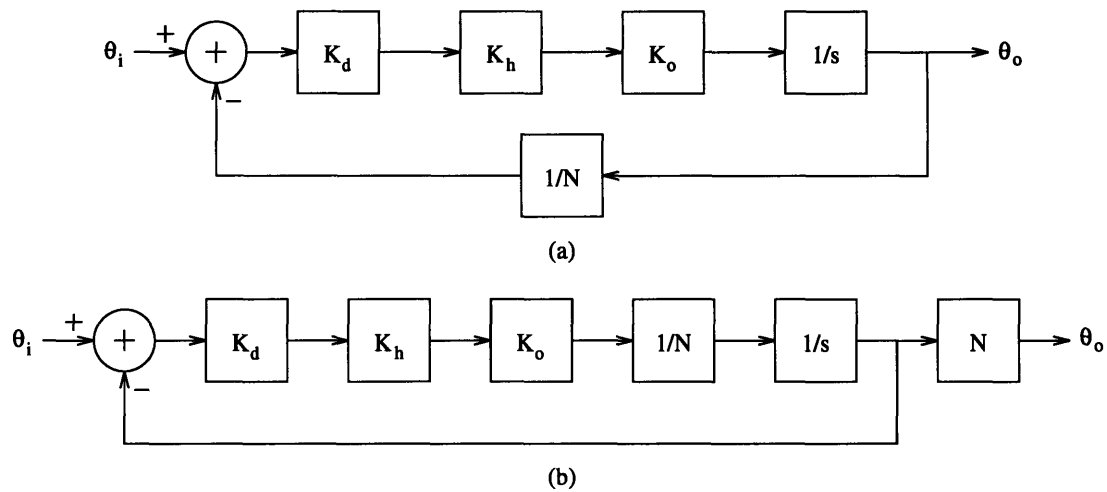


Figure 5.3 - Linear model for PLL with feedback path divider (a) and a manipulation of this model (b).

signal, where N is the division factor of the divider placed in the feedback path. In this PLL, this means that the VCO output frequency will be eight times the frequency of the reference signal. Depending on the output signal division factor selected, this leads to a final output signal which is at either two or four times the frequency of the reference signal.

The addition of this divider in the feedback path, however, also has an effect on how the performance characteristics of the PLL should be calculated. The primary effect of the divider is to reduce the bandwidth of the PLL. The easiest way to illustrate this is with a simple manipulation of the linear model for this new PLL. The basic linear model for this PLL which includes the divider is shown in Figure 5.3(a). This linear model is identical to the AC linear model present in Chapter 3, except that the division block has been added to the feedback path. Now, following the standard rules for the manipulation of these diagrams, an equivalent new linear model is shown in Figure 5.3(b). Remembering that the bandwidth of the original PLL shown in Chapter 3 was

$$BW = K_d K_h K_o \quad (5.3)$$

it should be apparent that the new bandwidth will be

$$BW = K_d K_h K_o / N \quad (5.4)$$

One factor to keep in mind, however, is that the addition of the divider will invoke a

change in the VCO. The frequency range of the VCO has to be modified so that both its minimum and maximum frequencies are moved to N times their original values. Without this modification, the PLL would not be able to continue to lock to the same range of input frequencies. This change results in a VCO gain which is N times the original value. Thus the actual bandwidth value for a particular PLL will stay the same.

One other parameter which requires additional consideration when calculated for this PLL architecture is the acquisition time. The initial frequency error can now be measured in one of two places: at the VCO output or at the input to the phase detector. If the frequency error is measured at the VCO output, the equation for acquisition time stays the same as Equation 3.50.

$$T_{acq} = \frac{\frac{\omega_{vco, eo}}{K} - 2\pi}{\pi\omega_z} \quad (5.5)$$

If one measures the initial frequency error at the phase detector inputs, however, the correct equation will be

$$T_{acq} = \frac{\frac{N\omega_{pd, eo}}{K} - 2\pi}{\pi\omega_z} \quad (5.6)$$

These equations are essentially the same since the frequency error at the VCO output will be N times larger than the frequency error at the phase detector inputs.

One last concern when adding a divider to the feedback path is the delay which will be added. Preferably, this delay should be matched to the delay of the input buffers through which the reference signal passes. [2] Otherwise, a static phase error will be introduced between the actual clock input and output signals of the clock buffering chip, even though it will not appear at the phase detector.

5.5 Ring Oscillator Circuits

A ring oscillator circuit is essentially composed of an odd number of inverter circuits connected in a ring. If the output of one of these inverter circuits is tapped to an external output, it will be observed to oscillate at a frequency which is determined by the number

of inverters in the chain, N_{inv} , and the delay through each inverter, T_d . Specifically, the frequency of oscillation will be

$$f_{out} = \frac{1}{2N_{inv}T_d} \quad (5.7)$$

Normally, the minimum number of inverters which can be used is three. If the inverter circuits have differential outputs, as with SCFL designs, however, this minimum may be reduced to two by wiring one circuit so that it is non-inverting. These two ring oscillator architectures are illustrated in Figure 5.4.

One common method for making a ring oscillator voltage-controlled is to use a “current-starved” inverter circuit. By increasing the current through these circuits, the delay through each inverter will be decreased. This will cause the frequency of oscillation to increase. A sample schematic for this type of circuit is shown in Figure 5.5. In this particular circuit, the amount of current drawn by the current sink transistor will increase if V_{cm} is increased. Another important characteristic of this design is that it uses diode loads. While resistor loads can be used, the result is a VCO with a much lower gain. Increasing the current, when resistor loads are used, results in a larger output voltage swing which counters the effect of the increased current.

The one drawback to building a ring oscillator out of current-starved inverters is that the range the bias current is varied over must be restrained in order to preserve linearity. The bias current can not be varied by an order of magnitude while still maintaining reasonable linearity. This in turn limits the frequency range which is possible with this type of VCO circuit. The PLL being constructed here requires an output range which extends at least from 200MHz to 500MHz. Simulations showed that this large frequency range was not possible using current-starved inverter circuits for the ring




Figure 5.4 - Basic ring oscillator architectures.




Figure 5.5 - Current-starved inverter circuit.

oscillator. Because of this linearity limitation on the standard current-starved inverter, however, a number of alternate circuits have been developed.

The circuits designed for this thesis are based on a Motorola patented circuit designed by Mavin Swapp. [3] This circuit consists of a normal inverting input gate with a latch gate attached to its outputs. The function of the latch gate is to slow down the rate at which the inverting gate's outputs change. In this circuit, rather than varying the current through the inverting input gate, the current through the latch gate is varied. The result is improved linearity over an increased frequency range. Two versions of this circuit were designed for this thesis. Both designs are described here to point out the design considerations which were taken into account. The schematics for these two circuits are shown in Figure C.1 and Figure C.2 in Appendix C.

The major difference between these two circuits is the mechanism used to control the current which flows through the latch gate. In the version of the circuit designed first, another switching gate was used. The current sink attached to this switching gate will draw a fixed current. By adjusting the voltages on the V_f and V_r inputs to this circuit, the portion of this sink current drawn through the latch gate, as opposed to directly from the

Table 5.1 - Output frequency for input voltages.

Input Voltage (V)	Output Frequency (MHz)	Input Voltage (V)	Output Frequency (MHz)
-1.3	574	-2.0	308
-1.4	541	-2.1	270
-1.5	509	-2.2	243
-1.6	475	-2.3	224
-1.7	438	-2.4	209
-1.8	398	-2.5	196
-1.9	354	-2.6	186

supply voltage, may be varied. Originally the V_r signal was intended to be a reference voltage, but, if desired, a pair of differential inputs could be connected to the V_r and V_f inputs. This would, however, reduce the input voltage range of the oscillator circuit, since the current would switch faster if both signals were changing. For a single-ended control voltage signal, the range of input voltages over which the change in output frequency will be linear is limited to approximately 0.8V by the inherent properties of the MESFET devices used. In addition, variations in the gain or threshold of the D-FETs used for this gate will have a significant impact on the performance of this circuit.

The second version of this circuit uses a different approach to varying the latch current that allows a wider input voltage range and which should be more robust to normal variations of device parameters. First, a larger than minimal D-FET is used in the current sink for the latch gate so that its V_{gs} drop will not vary significantly for the range of currents which will be flowing through the device. This causes the voltage on the drain of this device to stay fixed at V_{cs} minus its V_{gs} , so that the current sink resistor always requires the same current. This current must now either be supplied by the current-supplementing mechanism controlled by V_f or drawn through the latch gate. By changing the size of the resistors in the supplement current path, the input voltage range for V_f can be set by the designer. Because of the advantages offered by this circuit, subsequent work in this thesis, including the work matching the VCO input range to the loop filter output

range and the behavioral modeling work, dealt with this circuit. For this reason, the section on VCO performance will focus on the simulation results for this second version of the circuit.

One other important feature of both these designs is the diode-clamped resistor loads. This type of load was preferred over a simple diode-load because it provides more control over the output voltage swing. The resistor portion of the load helps to pull the logic voltage swing high by forcing the voltage across the diode to zero when there is no current through the load. The voltage across a simple diode-load would stay close to the diode threshold voltage. On the other hand, the diode portion of this load keeps the voltage swing fairly constant regardless of what the latch current is. A fairly large resistor value was used so that only a small current is required to cause this diode to start conducting.

5.6 Simulated VCO Performance

The final ring oscillator was constructed using the second version of the “latched” inverter circuit described in the previous section. This ring oscillator consists of three of these inverters in a ring. The circuit which was tested included an output buffer that represented the normal load that would be expected on the ring oscillator’s outputs. Once the loop filter had been constructed, the input voltage range of this ring oscillator was tuned to match the output voltage range of the loop filter op amp. This tuning was done by adjusting the value of the resistor in the supplement current path of each inverter.

The primary performance characteristic tested via simulations was output frequency versus input voltage. Based on these simulations, several additional changes were made to the “latched” inverter circuit. Because the initial frequency range of the circuit was well above the desired frequency range, the currents through the inverting gate and through the follower stages were reduced. In addition, the widths of the devices used in the inverting gate and the latch gate were increased to add extra capacitance to the circuit. The sink current for the latch gate was also adjusted on the basis of these simulations, so that it was large enough to allow the output frequency to be varied across the whole frequency range that was required.

The simulation results showing output frequency versus input voltage for the final version




Figure 5.6 - Output frequency versus input voltage.

of the circuit are shown in Figure 5.6. Table 5.1 lists the output frequencies for the specific input voltages that were simulated. Based on these results, the mean gain value for this ring oscillator is 320MHz/V. The maximum positive deviation from this value is +40%, while the maximum negative deviation is -68%. Another definition of this VCO's gain, for which the deviation is symmetric, is 280MHz/V \pm 61%.

While the linearity of this ring oscillator is fairly poor, this trade-off of linearity for frequency range is expected when designing a VCO with an output range as large as was required here. The main problem introduced by this non-linear behavior is that it makes predicting the PLL's performance more difficult. In particular, the equations used for calculating the PLL's bandwidth and acquisition time are based on an assumption that the VCO's gain is linear. This points to one of the advantages, however, of performing behavioral modeling, as was done for this thesis. The behavioral modeling allows the PLL's performance to be verified, while taking into account this non-linearity, as well as some other non-idealities of the circuits.

Simulations of this ring oscillator also looked at the time-domain characteristics of its output signals across its frequency range. Views of the ring oscillator's output

waveforms are shown at several different frequencies in Appendix D. As these waveforms show, this circuit has increasing asymmetries in its output as the frequency decreases. This is due to the increased effect of the latch, which acts as a feedback mechanism. These asymmetries are acceptable, however, because the outputs are fed through dividers, which remove these asymmetries, before being made available externally.

One parameter which was not characterized for this circuit was its modulation bandwidth. This testing was not performed because of a lack of familiarity with this parameter at the time the thesis work was done. Any subsequent work with this circuit should include determining its modulation bandwidth to ensure that it will not interfere with the stability of the PLL.

5.7 Noise Considerations

The VCO is one of the major potential sources of phase jitter in a PLL. The VCO input is particularly sensitive since any noise on this node will be translated directly into phase jitter. This is the motivation for trying to keep the VCO gain low if possible. The finite input resistance of the VCO input nodes is another potential source of noise. The forward-bias gate-source conduction of the input transistors will lead to a leakage current from the loop filter capacitor. In CMOS circuits, this effect is minimized by the high input resistance of MOSFETs. With MESFETs, however, this current will be much larger, making this effect more significant. Another potential source of noise in the VCO is $1/f$ noise from the devices. This effect should be negligible, however, at the high VCO frequencies seen in this PLL. [4] One noise advantage in this PLL is that, due to the output divider, the phase noise in the VCO will be either cut in half or to a quarter at the final output. Another noise advantage is that, since designing SCFL circuits using enhancement/depletion MESFETs leads to constant-current current-switching circuit designs, the switching noise associated with CMOS-type voltage-switching circuits will not be present.

A number of the steps which should be taken to minimize noise in the VCO and throughout the PLL involve considerations in the layout of the circuits. First, the routing of signals is important. Careful layout practices should be followed which minimize noise

coupling through parasitic capacitances. Next, if there are a large number of other circuits on the chip, separate power and ground pins should be used for the PLL. Finally, the PLL section of the chip is frequently guard-banded from the rest of the circuits on the chip. [5]

References

- [1] D. Chen, "Designing On-Chip Clock Generators," *IEEE Circuits and Devices*, vol. 8, pp 32-36, July 1992.
- [2] D. Chen, "Designing On-Chip Clock Generators."
- [3] M. Swapp, "Programmable Time Delay Circuit for Digital Logic Circuits," Patent 5,063,311, owned by Motorola, Inc., November 5, 1991.
- [4] D. Jeong, G. Borriello, D. Hodges, R. Katz, "Design of PLL-Based Clock Generation Circuits," *IEEE Journal of Solid-State Circuits*, vol. SC-22, pp 255-261, April 1987.
- [5] D. Chen, "Designing On-Chip Clock Generators."

Chapter 6 - The Loop Filter

6.1 Introduction

The primary decision to be made regarding the loop filter for this type of PLL, which uses a 3-state phase detector, is whether to use an active loop filter or a charge pump together with a passive loop filter. This chapter will begin by discussing the basic features of a charge pump and some of the advantages this design offers. Next, however, the reasons why an active loop filter was used will be explained. The following section will examine the details of the loop filter architecture which was implemented. Then, the design of the op amp circuit used within the loop filter will be described. The final section will discuss the performance of the op amp, and its affects on the overall performance of the PLL.

6.2 Charge Pumps

The primary feature of a charge pump is that it allows the discrete pulses from a 3-state phase detector to be utilized in such a way that a passive filter can be used within the PLL without a static phase error resulting. The charge pump essentially disconnects the loop filter from the phase detector. In this configuration, rather than acting as a filter for the phase detector output signal, the loop filter acts as a reservoir for charge. The signals from the phase detector direct the charge pump to either add or remove charge from this reservoir. To understand how this configuration eliminates the static phase error normally seen when a passive loop filter is used, first recall from Chapter 3 that a constant offset voltage, V_{co} , which depends on the frequency of the input signal, is required at the VCO input to produce the appropriate output frequency. In the original PLL configuration where the loop filter output voltage, V_c , is a filtered version of the phase detector output voltage, V_d , a static phase error was required within the system in order to produce V_{co} . The magnitude of this static phase error depended on the DC gain of the loop. Because a passive loop filter has a maximum DC gain of one, an active loop filter was required to make the static phase error negligible. With this new configuration using the charge pump, however, the loop filter output voltage, V_c , does not depend directly on the phase detector output voltage. Instead, V_c stays fixed unless a signal from the phase detector

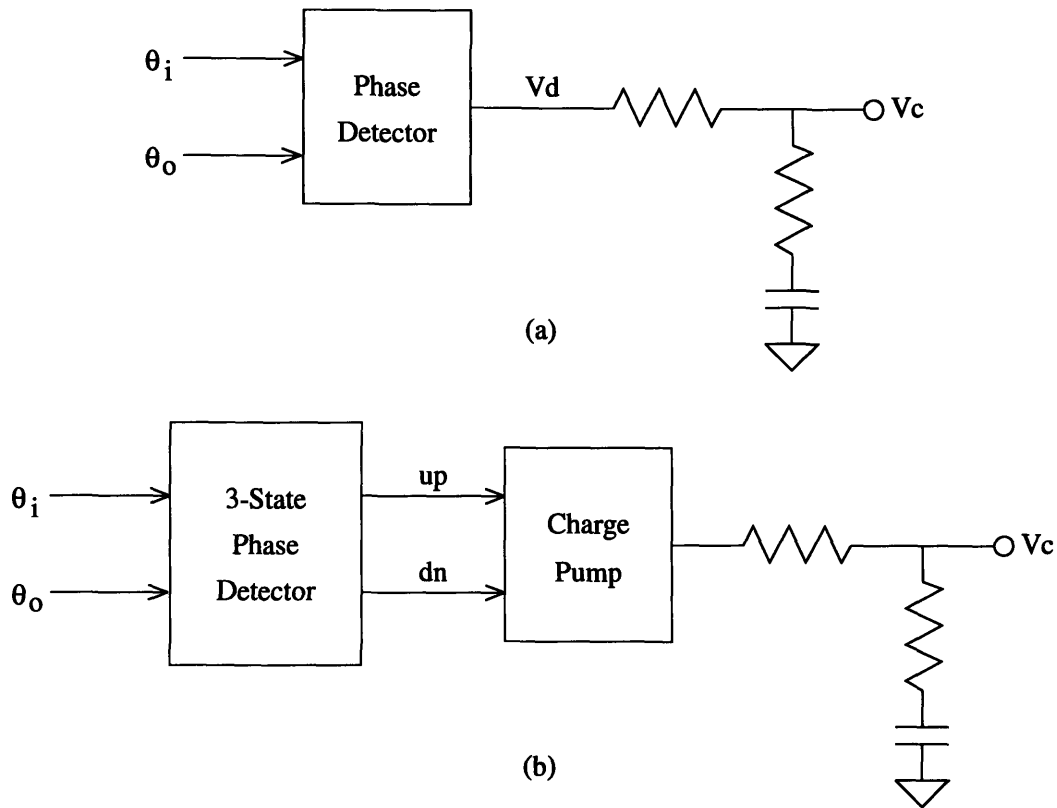


Figure 6.1 - Comparison of phase detector to loop filter connection with and without charge pump.

triggers the charge pump to add or remove charge from the loop filter. This means that once V_c has been tuned to compensate for V_{co} , a static phase error is not required within the system to maintain this voltage. An illustration of these two different configurations is shown in Figure 6.1.

There are two types of charge pumps: voltage pumps and current pumps. The simpler of these two types is the voltage pump. A possible implementation of a voltage pump in CMOS is shown in Figure 6.2. The transistors in this architecture act simply as switches. When the Up signal is asserted, the p-channel FET will become conductive, allowing a current, called the source current, to flow into the loop filter. When the Dn signal is asserted, the n-channel FET will become conductive, causing a current, called the sink current, to flow out of the loop filter. The magnitude of these currents will be determined by the voltage difference between the loop filter output node, V_c , and the appropriate supply voltage rail, divided by the value of the resistor, $R1$. The designer's primary control over the size of these currents is through adjusting the value of the




Figure 6.2 - Voltage pump type charge pump.




Figure 6.3 - Concept behind a current pump.

resistor, R_1 . It is important to note that for most values of V_c , the source and sink currents will not be equal. In many PLL configurations, this behavior would be unacceptable. If the phase detector used produces a pulse on both outputs when the PLL is in lock, the unequal source and sink currents would change the control voltage seen by the VCO. This would cause the PLL to lose lock on the input signal. In addition, this behavior will cause the magnitude of the PLL's response to a phase error to depend on both the current control voltage, V_c , and the sign of the phase error. This asymmetric response can lead to problems with small signal stability within the PLL. [1]

In a PLL where the source and sink currents must be matched, a current pump is the appropriate choice. Figure 6.3 illustrates the basic functionality of a current pump. As this figure shows, the magnitude of the source and sink currents are independent of the

loop filter output voltage for this type of circuit. In order to achieve this functionality, a feedback path that compensates for the present loop filter voltage may be necessary. [2] Another advantage of this type of charge pump is that it is possible to reduce the size of the loop filter components required by reducing the magnitude of the source and sink currents. There is, however, a limit to the extent that this can be done. At some point, the magnitudes of these currents become difficult to control, as they become overly sensitive to noise. [3] In addition, the magnitude of these currents is now an important parameter for calculating the overall performance of the PLL, since the relationship between phase detector voltage and current into the loop filter has been made arbitrary. Because a charge pump was not used for this thesis, this derivation is not presented here. For those readers who are interested, however, Gardner has written a detailed article describing the effects of charge pumps on PLL performance. [4]

6.3 Loop Filter Design Trade-offs

This section explains the reasons why the active filter approach was chosen for this PLL. This decision is interesting partially because the charge pump approach is the standard approach in PLLs built for clock distribution applications. In order to get high performance from the charge pump approach, a current pump must be used. It is a difficult design task, however, to design a current pump so that the charge and discharge currents stay symmetric across the full range of the loop filter output voltage. Thus, when a current pump is used, this circuit block is generally the critical section of the PLL. For this reason, designs for these circuits are seldom found in the literature on this topic.

The lack of available designs in other technologies is one reason why a current pump was not used here. By choosing the active loop filter approach, this design problem was eliminated. At the same time, another member of the design team, with several years of circuit design experience, did pursue the current pump approach. Within the scope of this project for Motorola, this was advantageous since it meant that both alternatives would be explored and compared. The trade-off to the active filter approach, however, is that it requires designing an op amp circuit.

In addition, there are some specific problems associated with this approach that are introduced. In particular, the input offset voltage and input offset current of the op amp




Figure 6.4 - Final loop filter architecture.

circuit used within the active loop filter will both contribute to the static phase error within the PLL. The details of these effects are described in the section on op amp performance. Another requirement with this approach is that the absolute voltage levels of the signals coming from the phase detector must be well matched or an additional static phase error will be introduced. In a charge pump, depending on the specifics of the circuit architecture used, it may be possible to treat these signals in a digital manner, so that the signal levels are not as critical. A final consideration with the active filter approach is the normal temperature sensitivity of op amp circuits.

On the other hand, it is much simpler to balance the response of an active loop filter. Also the active filter approach tends to be more sensitive to the phase detector output signal than the charge pump approach. This leads to a smaller phase deadband.

6.4 Active Loop Filter Architecture

This section examines the important details of the loop filter architecture used in this thesis. A diagram of this architecture is shown in Figure 6.4. The most noticeable feature of this architecture is the dual feedback paths. This symmetry serves two purposes. First, it simplifies balancing the response of the loop filter, as compared with using a charge pump. Second, it helps to minimize the phase detector offset voltage. The phase detector offset voltage, as was explained in Chapter 3, is the voltage seen by the loop filter from the phase detector when there is no phase error. This voltage introduces a static phase error in the system which can not be eliminated by increasing the DC loop gain. Instead, either the phase detector must be designed to output zero volts for zero

phase error, or some technique must be used to minimize the offset voltage seen by the loop filter. This architecture is one of the best approaches to minimizing this offset voltage. The equations for calculating the loop performance parameters are the same for this architecture.

One drawback of this architecture is that it requires twice as many resistors and capacitors as the standard loop filter architecture. This is a concern since one of the specifications set by the product engineers as a goal for this project was to eliminate the necessity of any external components. In general, there are a number of trade-offs between using on-chip or external components. Implementing all components on-chip will reduce costs for the end user, save board space, and improve reliability. [5] On the other hand, on-chip components make it difficult to adjust the loop parameters to fit different applications and to accommodate process variations. It is unlikely, however, that the capacitors required for this PLL could be implemented on-chip for any of the loop filter architectures considered. The process used for this thesis lacked sufficiently high value capacitors to implement the loop filter capacitors within an acceptable die area. In general, it is possible to reduce the required capacitor values by increasing the resistor values, however, at some point large resistor values become equally difficult to implement.

A second important feature of this loop filter is the ripple suppression capacitor, C3, which has been added between the input connections. The pulse shape of the inputs from the phase detector adds a substantial high frequency component to the input signal. The op amp selected for the active loop filter may not be able to handle these signal components. When the PLL is in lock, these high frequency components may produce distortions in the control voltage produced by the loop filter, and these distortions would cause phase jitter within the PLL. Capacitor C3 smooths out the control voltage signal by filtering out high frequency components. This additional capacitor is generally also found in charge pump type PLLs for the same reason.

The essential effect of this capacitor is to add a high frequency pole to the PLL. The equation for this pole is

$$\omega_{hf} = \frac{4}{R_1 C_3} \quad (6.1)$$

At the same time, this capacitor's position is carefully chosen so that the equations for the PLL's other poles are not modified from those shown in Chapter 3 for the standard active loop filter. The introduction of this pole makes the PLL a third order system, and thus, it is no longer inherently stable. The effects of this pole on the performance of the PLL can be minimized, however, by simply following a design rule that ω_{hf} is always kept greater than 4K. If a designer is not careful, though, the VCO and the loop filter op amp both introduce poles that could affect the PLL's performance in a similar, but less controllable manner.

The last architectural element of this loop filter which should be mentioned is that its output is directly connected to the input nodes of the VCO. This means that the input range of the VCO is limited to the output range of the op amp used for this loop filter. There is actually one simple change which could be made to improve this connection. Currently, the loop filter is connected to each of the three inverter stages within the VCO. If an intermediate source follower stage was added, the leakage current from the loop filter due to forward-bias conduction currents of these connections could be reduced.

Once this loop filter architecture was decided upon, the next step was to determine what resistor and capacitor values to use. These values were chosen based on a combination of performance requirements and loop stability considerations. The most critical performance requirement was that the maximum acquisition time allowed for this PLL was 10ms. At the same time, however, the bandwidth of the PLL should be kept as narrow as possibly in order to minimize phase jitter. Recalling Equation 5.5 that describes the acquisition time for this PLL,

$$T_{acq} = \frac{(\omega_{eo,vco}/K) - 2\pi}{\pi\omega_z} \quad (6.2)$$

one can see that the maximum acquisition time will depend on the maximum initial frequency error within the system at the VCO output. Based on simulations of the loop filter and of the VCO, the free-running frequency was determined to be 240MHz. Since the required VCO range is 200-500MHz, this equates to a maximum initial frequency error of 260MHz. This definition is sufficient when the VCO is expected to move only

once from its initial state to a fixed frequency, as in clock distribution applications. In an application, where the input frequency is not fixed, however, the maximum increment that the input frequency can move at one time would also have to be considered.

In order to calculate the required loop filter component values, the phase detector gain and the VCO gain also had to be extracted from simulations. In Chapter 5, the phase detector gain was shown to be 0.18V/rad. In Chapter 6, the VCO gain was specified as either 280MHz/V or 320MHz/V. The higher value was chosen here since it is more reflective of the gain across the region of the VCO output characteristic which must be traversed in the case of the maximum frequency error. Converting to the proper units, the VCO gain is 2010Mrad/s/V.

Next, in order to prevent unreasonable peaking of the system response, ω_z was defined to equal $K/4$. This results in a system which is critically damped. By specifying this relationship, it is possible to define R_1 in terms of the values for R_2 and C .

$$\omega_z = \frac{1}{R_2 C} = \frac{K_d K_o R_2}{4 R_1} = \frac{K}{4} \quad (6.3)$$

$$R_1 = \frac{K_d K_o (R_2)^2 C}{4} \quad (6.4)$$

Between this relationship and the 10ms lock requirement, an equation for C in terms of just R_2 can be derived. This is useful since it means that choosing R_2 will defined the other two major component values. First recalling Equation 6.2 in a slightly different form,

$$T_{acq} = \frac{\omega_{eo,vco}}{\pi K \omega_z} - \frac{2}{\omega_z} = 0.01 \quad (6.5)$$

Next, substituting the equations for K and ω_z , as well as the relationship just specified for R_1 ,

$$\frac{\omega_{eo,vco} (R_2)^2 C^2}{4\pi} - 2R_2 C = 0.01 \quad (6.6)$$

Finally, neglecting the $2R_2 C$ term since it will be insignificant for reasonable component

values, the equation for C is

$$C = \sqrt{\frac{4\pi(0.01)}{\omega_{e0,vco}(R_2)^2}} \quad (6.7)$$

The last loop parameter of interest, ω_{hf} , was set equal to 4K in order to ensure loop stability while performing the maximum amount of ripple filtering. This fixes the value of C3 in relationship to the other parameters which have already been determined.

$$C_3 = \frac{1}{R_1 K} \quad (6.8)$$

These equations were then placed in a spreadsheet, and values of R2 were tried until reasonable values for the other components resulted.

The values which were selected with this procedure were intended to be a reasonable first guess. The behavioral modeling was then intended to be used to verify these values, since the non-linearity of the VCO gain and the small-signal behavior of the phase detector would be taken into account there. The values chosen here, however, were $R_1 = 694K\Omega$, $R_2 = 1K\Omega$, $C = 8.8nF$, and $C_3 = 3.1pF$. The loop bandwidth for these values is 456Krad.

6.5 Op Amp Design

A fairly straight forward amplifier circuit was used for the op amp in this active loop filter design. The schematic for this circuit is shown in Figure 6.5. As the schematic shows, this circuit consists of three identical gain stages. Large resistor loads were used with each stage, because they actually produced better gain than using active loads as shown by simulation.

A compensation capacitor was included across the switching nodes of the first gain stage. This capacitor cuts off the gain of the op amp at higher frequencies to ensure a good phase margin for stability, since the op amp will be used in a feedback circuit. The value chosen here, 200pF, was selected based on simulations where it provided a phase margin of 47 degrees. As will be discussed below, some consideration could be given to increasing this phase margin further.

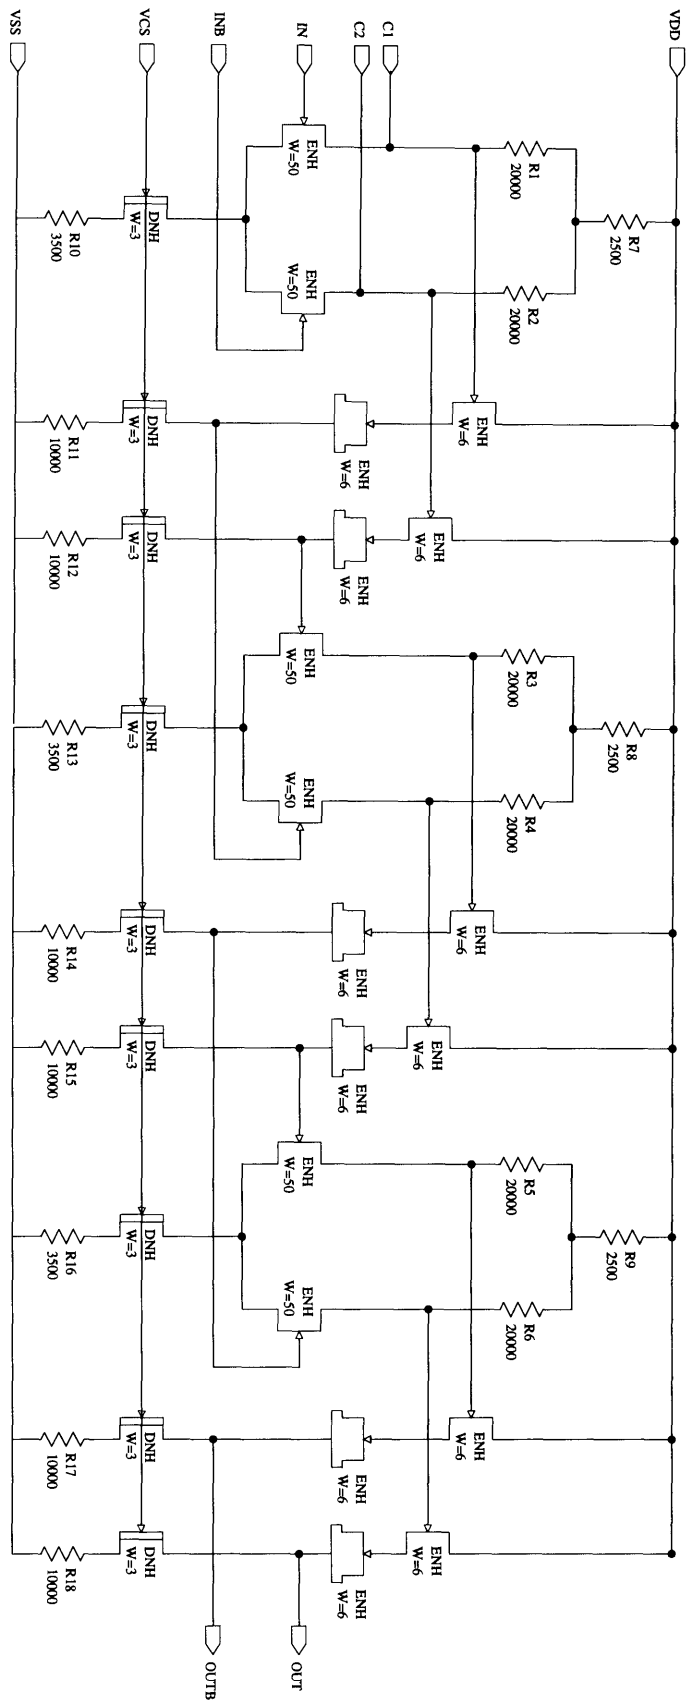


Figure 6.5 - Op amp circuit.

6.6 Op Amp Performance and Effects on PLL Performance

The most important performance characteristics for an op amp used in this configuration are gain, bandwidth, output voltage range, input offset voltage, and input offset current. In this circuit, the widths of the switching transistors were increased until the open loop gain for this three stage configuration was approximately 1000. This value was considered sufficient for this application.

The next important parameter is the bandwidth, which is defined here in terms of the gain-bandwidth product (GBP). The GBP reflects the fact that the bandwidth of the op amp is a function of the gain required from the op amp. Dividing the GBP by the gain required yields the bandwidth of the op amp for that application. The value of this parameter for this amplifier circuit was 47MHz. In a PLL, the GBP is used to determine where the high frequency pole introduced by the op amp is located. The equation for this pole's location is [6]

$$\omega_{hf} = 2\pi \frac{R_1}{R_1 + R_2} \text{GBP} = \frac{2\pi}{1 + K_h} \text{GBP} \quad (6.9)$$

If a separate high frequency pole has already been added to the PLL, as was done here with the ripple suppression capacitor, a designer should make sure that this pole is located sufficiently above the deliberate pole not to affect the PLL's stability. In this circuit, where K_h will be very small compared to one, the pole introduced by the op amp will be at 295Mrad. This value, which is more than 160 times greater than the location defined for ω_{hf} , will not present a problem. In fact, it is this large value that suggests that the compensation of the op amp circuit could be increased if desired.

The next op amp performance characteristic that requires attention in this loop configuration is the op amp output voltage range. The output voltage range possible for this circuit was -1.3V to -2.6V. As was mentioned above, the one output from the op amp is connected directly to the inputs of the VCO. Therefore, improving this range would allow the gain of the VCO to be lowered. One option for improving this range would be to redesign the current circuit. A second option, however, would be to design an intermediate circuit that takes advantage of both outputs from the loop filter to provide a wider input voltage range to the VCO.

The final characteristics which should be considered for the op amp circuit are the input offset voltage, V_{os} , and the input offset current, I_{os} . These parameters are important because they will affect the static phase error in the PLL in the same manner that the phase detector offset voltage does. For this reason, their effects are modeled as an addition to the phase detector offset voltage so that equation for this contribution to the static phase error becomes [7]

$$V_{do,new} = V_{do} + V_{os} + I_{os}R_1 \quad (6.10)$$

Thus, Equation 3.7 becomes

$$\theta_{eo} = \frac{-V_{do,new}}{K_d} + \frac{V_{co}}{K_d F(0)} \quad (6.11)$$

These parameters can not be determined from simulations, however, since they are the result of device variations which occur only with the actual process. All devices in a simulation are identical since the same model parameters are used to calculate each device's performance. A comparison can be made, however, between different process technologies. In CMOS and BiCMOS processes, MOSFETs are used for the input gates so that the input bias currents into the op amp are very small. This leads to a very small input offset current. In GaAs, where MESFETs are the only devices available, the input bias currents will be significantly larger. Thus, a higher input offset current must be tolerated in GaAs designs.

References

- [1] D. T. Gavin and R. M. Hickling, "A PLL Synthesizer Utilizing a New GaAs Phase Frequency Comparator," *Proceedings of RF Technology Expo 88*, Cardiff Publishing: Englewood, CO., pp 457-71, 1988.
- [2] D. Chen, "Designing On-Chip Clock Generators," *IEEE Circuits and Devices*, vol. 8, pp 32-36, July 1992.
- [3] D. Chen, "Designing On-Chip Clock Generators."
- [4] F. M. Gardner, "Charge-Pump Phase-Lock Loops," *IEEE Trans. Commun.*, vol. COM-28, pp 1849-1858, Nov. 1980.
- [5] D. Chen, "Designing On-Chip Clock Generators."
- [6] D. H. Wolaver, *Phase-Locked Loop Circuit Design*, Prentice Hall: Englewood Cliffs, NJ, Sections 3.7, 1991.
- [7] D. H. Wolaver, *Phase-Locked Loop Circuit Design*, Sections 3.4.

Chapter 7 - Behavioral Modeling

7.1 Introduction

This chapter describes the work which was done developing a behavioral model for a PLL composed of the circuits discussed in the last three chapters. The chapter begins by explaining the motivations for developing this kind of model. Next, the steps taken to model each block of the PLL are examined. Finally, the performance observed from this model is discussed, along with the implications of these results.

7.2 Motivations for Behavioral Model

As was discussed in the previous chapter, the linear equations presented in Chapter 3 for predicting a PLL's behavior have a limited accuracy. They necessarily abstract away from the details of any specific implementation. In particular, these equations will not account for any non-linearity in the phase detector gain or VCO gain. On the other hand, trying to simulate a PLL's performance using a circuit simulator such as HSPICE also presents problems. The acquisition process of a PLL can take anywhere from several microseconds to several milliseconds. At the same time, a simulation of this process must have nanosecond resolution because of the frequency of the signal to which the PLL is trying to lock. This requirement, along with the large number of devices which must be simulated, leads to simulations which can take several days on a mini-computer or workstation. [1] Using a behavioral model, the performance of a specific PLL can be simulated while reducing the simulation time to minutes. The trade-off is that the designer must spend additional time constructing accurate models for each block of the PLL.

7.3 Construction of Behavioral Models

In the behavioral model built here, the majority of the PLL circuits, including the VCO, the loop filter, and the feedback path divider, were constructed out of standard components provided with the behavioral modelling software. Both the VCO and the feedback path divider were modelled using single blocks designed to simulate their functions. The model used for the VCO allowed the VCO's gain to be specified as a piecewise-linear function by specifying voltages and the associated output frequency. The data used here to specify this gain was taken from Table 6.1. It is important to note that

while the input to this block was analog, the output waveform was digital. Because the divider block was completely digital, the only parameters which had to be specified were the delay and the divide ratio. For this model, the divider was set to divide by eight, while the propagation delay was left at zero.

The model for the loop filter was constructed using a differential op amp model together with standard resistor and capacitor components. This portion of the behavioral model was completely analog. Using the differential op amp model required specifying a number of parameters including differential gain, common-mode gain, slew rate, output voltage range, and bandwidth. The values for all of these parameters were extracted from the HSPICE simulations of the op amp circuit. In order to check the accuracy of the resulting model, the frequency response of a sample RC feedback network constructed around this op amp model was compared with the response obtained from an HSPICE simulation of the identical circuit network. The characteristics of the behavioral model were tuned slightly so that these responses were closely matched.

The implementation of the model for the phase detector was more complex. First, a decision was made not to attempt the detailed characterization of the phase detector's performance that would have been required to accurately model the phase detector in a single block. Instead, the approach selected was to reconstruct the phase detector by modeling its individual blocks. Besides being simpler, this approach was considered more likely to reproduce subtleties in the phase detector's behavior which might be missed in the characterization process.

An additional complication which had to be addressed in this model was that, while the inputs were digital, the signals fed to the mixer and then to the loop filter needed to be analog. Within the SABER behavioral modeling package used, this required specifying a "hyper-model". This model allowed the appropriate signal characteristics to be defined for converting digital signals to analog signals. Conversion points were placed so that the up and dn signals fed to the mixer would be converted into analog signals with the appropriate rise and fall times, as well as accurate voltage levels.

Once this "hyper-model" was in place, the majority of the phase detector architecture was modeled using the standard digital models provided for delay buffers,

inverters, and D-type flip-flops. The propagation delay specifications required for each of these circuits were extracted from HSPICE simulations of the corresponding circuit.

A new behavioral model, however, had to be developed, in order to model the mixer block. The basic function of this circuit block is described by the equation

$$\Delta V_{\text{out}} = G \cdot \Delta V_{\text{in}} \quad (7.1)$$

where ΔV_{out} and ΔV_{in} represent, respectively, the voltage difference between the circuit's outputs and the voltage difference between its inputs, and G represents the gain of the mixer circuit. The gain was the major parameter which needed to be extracted from HSPICE simulations.

A number of secondary parameters had to be determined first, however, in order to be able to correctly ascertain and model the gain. These secondary characteristics included the quiescent operating points of the input and output nodes and the input and output voltage swings. The output related values mentioned here were required to produce output waveforms with the correct voltage levels, while the input related values were important because they must be taken into account when determining the gain of the mixer. These values were all obtained from very straightforward HSPICE simulations of the mixer circuit.

Next, the output characteristic of the mixer was simulated in HSPICE for a controlled set of input conditions where the common-mode bias point of the input signals was kept centered at the quiescent operating point of the input nodes. The gain of the mixer under these conditions was calculated from this output characteristic by dividing the output voltage difference by the corresponding input voltage difference. In order to reproduce this gain in the behavioral model, a fourth order polynomial was fit to the data collected from this simulation.

$$G = 3.078 - 2.716\Delta V_{\text{in}}^2 - 2.528\Delta V_{\text{in}}^4 \quad (7.2)$$

As this equation shows, the gain of the mixer varies with the size of the input voltage difference. Figure 7.1 shows a plot of the gain data collected from HSPICE with the gain curve produced by the polynomial derived here overlaid. By determining the gain for a uniform common-mode bias point, as was done here, the dependence of the gain on this

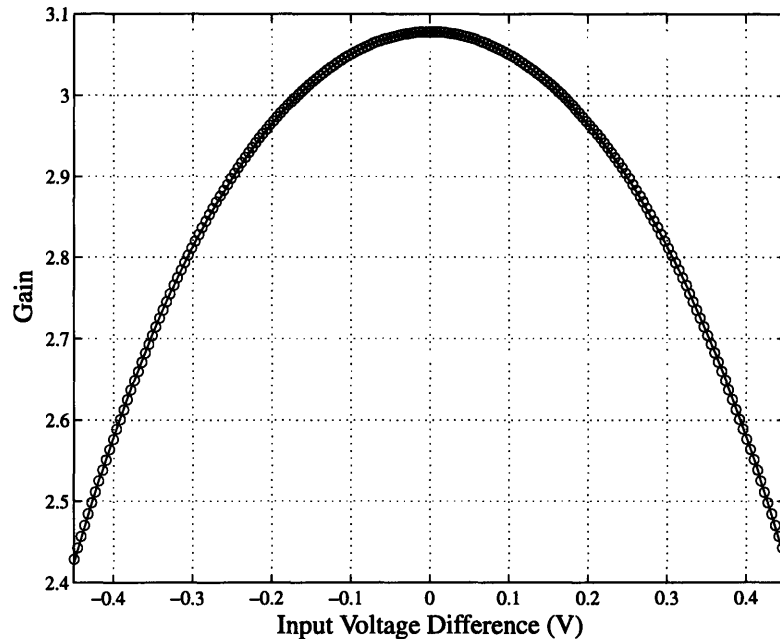


Figure 7.1 - Mixer gain versus size of input voltage difference.

parameter can be factored out and modelled separately.

In order to test the effect of the common-mode bias point of the input signals on the mixer's gain, a second set of HSPICE simulations were performed. In these simulations, the common-mode bias point of a fixed input voltage difference was varied across the expected range of voltage based upon the mixer's input voltage swing. In general, these simulations showed that the mixer's gain decreased as the common-mode bias point moved away from the quiescent operating point in either direction. The rate at which the gain decreased, however, was found to depend on the size of the input voltage difference. The larger the input voltage difference is though, the smaller the expected voltage range for the common-mode bias point is. This is because there is less space for the DC levels of the input signals to move within boundaries on the input voltage swing when there is a large voltage difference between these signals. For this circuit where the range of the input voltage swing was 450mV, the common-mode bias point of a 100mV voltage difference can move 350mV, while the common-mode bias point of a 400mV swing can only move 50mV. In addition, emphasis was placed on accurately modeling the behavior of the phase detector for small phase errors where the input voltage differences

seen by the mixer would be smaller. Due to these factors, the effect of the common-mode bias point was modelled based on the simulation results produced for a 100mV input voltage difference. This selection provided more accuracy for smaller input signals, while still providing reasonable accuracy for large input signals.

The effect of the common-mode bias point on the mixer's gain was modelled as a constant that the original gain was multiplied by to determine the final gain for a particular input signal.

$$G_f = M \cdot G \quad (7.3)$$

where G_f is the final gain of the mixer, and M is the constant determined by the common-mode voltage. The equation for M , which was found by fitting a fourth order polynomial the results of the simulation mentioned above, is

$$M = 1 - \frac{45}{3.05} V_{\text{cmo}}^2 - \frac{200}{3.05} V_{\text{cmo}}^4 \quad (7.4)$$

In this equation, V_{cmo} represents the common-mode offset voltage from the quiescent operating point of the mixer's inputs. The characteristic produced by this equation is shown in Figure 7.2.




Figure 7.2 - Illustration of gain multiplication constant for common-mode voltage effect.

One last factor that was considered in modeling the mixer's gain was its frequency dependence. HSPICE simulations testing this parameter of the mixer showed the pole of the mixer's frequency response to be at over 7GHz. This value was sufficiently high that this effect was neglected.

Once all the essential parameters for the mixer were accurately modelled, the output from the behavioral model for the full phase detector was compared to the output seen in HSPICE simulations for several different phase errors. In all cases, including for very small phase errors, the behavioral model closely approximated the performance seen in HSPICE. The code written to model the mixer circuit can be found in Appendix F. This code was written with the help of another Motorola engineer, Raymond Garcia, who was familiar with the details of implementing behavioral models in the SABER behavioral modeling software used. A diagram of the final behavioral model for the full PLL is shown in Figure 7.3.

7.4 PLL Performance in the Behavioral Model

The primary goal of the simulation work done with the behavioral model developed here was to test the acquisition process of the PLL which had been designed. Extensive simulations were run to test this performance. The results of these simulations, however, showed a sporadic, but reoccurring phenomena where the PLL output, after initially starting to tune towards the frequency of the input signal, would become trapped within a range of frequencies where it would oscillate without being able to escape to finish tuning to the frequency of the input. Conversely, in the majority of the simulations, the PLL did acquire frequency and phase lock. Initially, this erratic behavior was believed to be linked to a problem with the PLL's bandwidth. After extensive investigation, though, it was these results that pointed to a problem with the phase detector, which had been overlooked originally. This flaw in the phase detector was found by careful examination of the phase detector output signals, which showed the outputs to be cycling in such a manner that the average output voltage to the loop filter was zero. This behavior was more a result of the reduced linear range problem described in Chapter 5 than the deadband problem also discussed there.

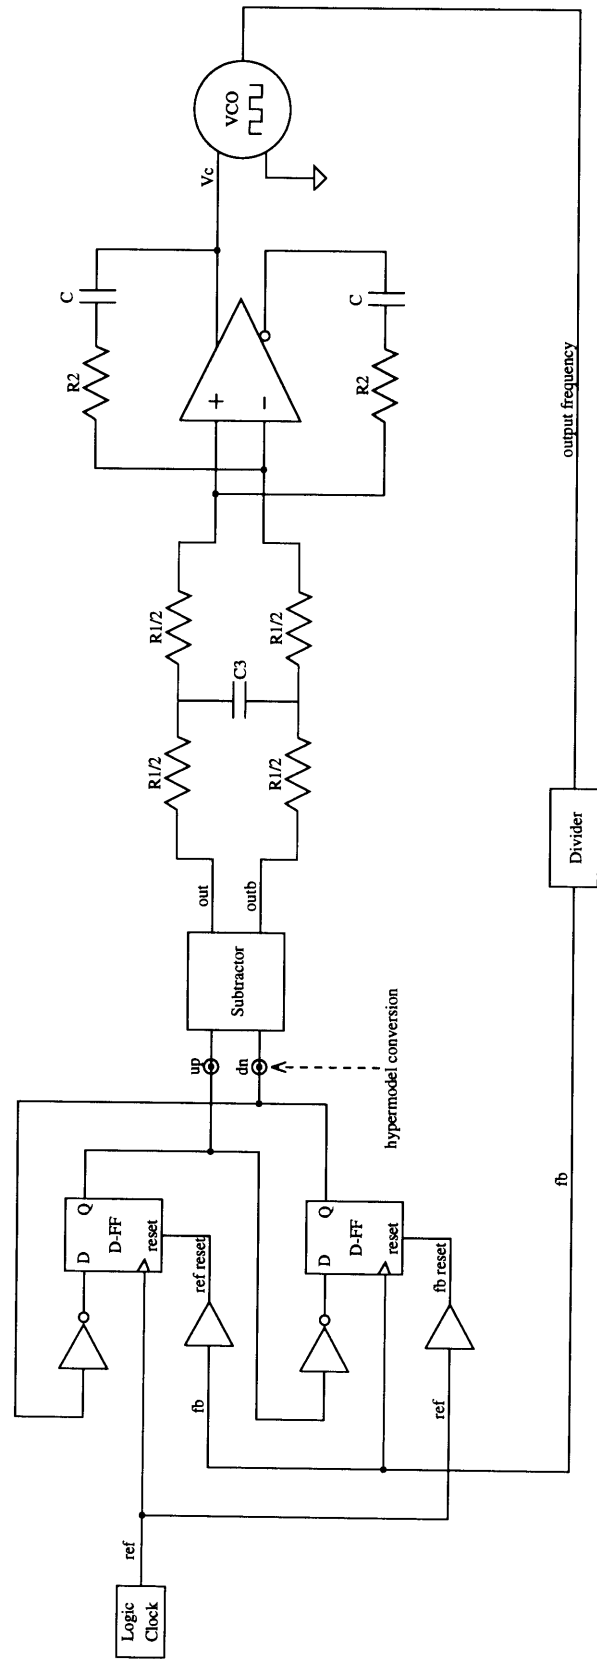


Figure 7.3 - Behavioral model for full PLL.

Thus, while the original goal of this simulation work had to be temporarily laid aside, the behavioral model proved to be a very useful test of the PLL's functionality. To understand the full advantage of the behavioral model in this regard, consider how few signals would have been available for inspection if these results had been observed in a fabricated version of the PLL. With the behavioral model, all of the internal signals were available for inspection. In addition, once the necessary changes have been made to correct the design flaw found in the phase detector circuit, updating this behavioral model and using it for its original purpose is an easy task.

References

- [1] D. Chen, "Designing On-Chip Clock Generators," *IEEE Circuits and Devices*, vol. 8, pp 32-36, July 1992.

Chapter 8 - Conclusions

8.1 Summary of Work

The primary work completed in this thesis was the design of a phase detector circuit, a VCO circuit, and a loop filter circuit, such that a full PLL could be constructed. A significant part of the work done within this thesis, however, was simply a compilation and comparison of the approaches to designing these components. While circuit implementations will change, the general design issues and design parameters discussed within this thesis are likely to remain the same. The other important element of the work completed here is the behavioral model which was developed. Not only was this model useful for simulating the PLL designed in this thesis and as a potential base for future models, but it also served as an introduction to the advantages of performing this type of simulation: advantages which were clearly seen in examining the functionality of the PLL built here.

8.2 Suggestions for Continuing Work

Work could be continued on this thesis in a number of areas. First, the circuit block which requires the most attention is the phase detector. A redesign of the flip-flop used in this architecture would be required in order to make the circuit function properly. The alternative, however, would be to choose one of the other phase detector architectures considered in Chapter 5, such as the dual flip-flop plus AND gate implementation. This design might be easier to realize since the current flip-flop design would be acceptable for this architecture. The only new design required would be a standard AND gate. The trade-off would be between using a more common, well-tested design and testing a new approach.

Several circuit design suggestions were also made in the chapters on the VCO and on the loop filter. Expanding the control voltage range of the VCO by redesigning the op amp circuit or by adding an intermediate circuit should be considered seriously because of the noise sensitivity of the input node. Another suggestion discussed an approach for reducing the loop filter leakage current into the VCO inputs. This improvement may be less critical, but would not require much effort to implement.

A second facet of the work which needs to be continued is the layout and testing of actual circuits. While individual circuit blocks such as the VCO, dividers, and phase detector were developed through layout, fabrication, and, in some cases, test analysis, no attempt has been made to layout and fabricate a full PLL based on these designs. Once fabricated, these circuits would require testing to verify their performance. In addition, the behavioral modeling work for simulating the full PLL's performance could be continued.

A last area of suggestions for continuing the work begun for this thesis would be to make substantial improvements or changes in the design. One suggestion along these lines would be to consider adding a frequency acquisition circuit to the PLL. Adding such a circuit would allow the bandwidth of the PLL to be significantly reduced, while still acquiring lock in less than 10ms. The result of this lower bandwidth would be less phase jitter when in lock. A trade-off here is that lowering the bandwidth would require using larger capacitors in the loop filter feedback network, as well as additional die area for the frequency acquisition circuitry itself. In general, there are many avenues along which this work could be continued, but a solid foundation has been provided here for pursuing any of them.

Bibliography

- D. H. Wolaver, *Phase-Locked Loop Circuit Design*, Prentice Hall: Englewood Cliffs, NJ, 1991.
- F. M. Gardner, *Phaselock Techniques*, John Wiley & Sons: New York, 1979.
- D. Chen, "Designing On-Chip Clock Generators," *IEEE Circuits and Devices*, vol. 8, pp 32-36, July 1992.
- D. Jeong, G. Borriello, D. Hodges, R. Katz, "Design of PLL-Based Clock Generation Circuits," *IEEE Journal of Solid-State Circuits*, vol. SC-22, pp 255-261, April 1987.
- P. R. Gray and R. G. Meyer, *Analysis and Design of Analog Integrated Circuits*, John Wiley & Sons: New York, Chapter 10, 1977.
- S. L. Long and S. E. Butner, *Gallium Arsenide Digital Integrated Circuit Design*, McGraw-Hill Publishing Company, 1990.
- R. F. Pierret, *Field Effect Devices*, Addison-Wesley Publishing Co.: Reading, MA, Chapter 1, 1990.
- D. T. Gavin and R. M. Hickling, "A PLL Synthesizer Utilizing a New GaAs Phase Frequency Comparator," *Proceedings of RF Technology Expo 88*, Cardiff Publishing: Englewood, CO., pp 457-71, 1988.
- I. Shahriary, G. Des Brisay, S. Avery, P. Gibson, "GaAs Monolithic Digital Phase/Frequency Discriminator," *IEEE GaAs IC Symposium Technical Digest*, pp 183-186, 1985.
- J. Andrews, A. Podell, J. Mogri, C. Karmel, K. Lee, "GaAs MMIC Phase Locked Source," *IEEE MTT-S Digest*, pp 815-818, 1990.
- T. Nakagawa, T. Hirota, and T. Ohira, "A Novel MMIC Approach to Sampling Phase/Frequency Detection for Microwave Quartz-Lock Oscillators," *IEEE MTT-S Digest*, pp 1301-1304, 1992.
- F. M. Gardner, "Charge-Pump Phase-Lock Loops," *IEEE Trans. Commun.*, vol. COM-28, pp 1849-1858, Nov. 1980.
- M. Swapp, "Programmable Time Delay Circuit for Digital Logic Circuits," Patent 5,063,311, owned by Motorola, Inc., November 5, 1991.
- R. R. Cordell and W. G. Garrett, "A Highly Stable VCO for Application in Monolithic Phase-Locked Loops," *IEEE International Solid-State Circuits Conference*, pp 106-107, Feb. 1990.
- Jieh-T. Wu, "A Bipolar 1GHz Multi-Decade Monolithic Variable-Frequency Oscillator," *IEEE Journal of Solid-State Circuits*, vol. 27, pp 668-670, April 1992.

- M. Soyuer and J. D. Warnock, "Multigigahertz Voltage-Controlled Oscillators in Advanced Silicon Bipolar Technology," *IEEE Journal of Solid-State Circuits*, vol. 27, pp 668-670, April 1992.
- R. C. Walker, "Fully Integrated High-Speed Voltage Controlled Ring Oscillator," Patent 4,884,041, owned by Hewlett-Packard, Co., Nov. 28, 1989.
- R. Sodhi, "title of thesis," Thesis for Bachelor of Science, MIT, Cambridge, MA, 1995.

Acknowledgments

I would like to thank Eric Herald, Kurt Sakamoto, Peter Wennekers, and Tom Bushey for their advice regarding the design of the circuits found within this thesis. I would also like to thank Glenn Howell and Derek Everett for their help laying out the circuits for this thesis and Ray Garcia for his help in developing the behavioral models for the PLL. I would like to thank Ray Sundstrom and Dwight Esgar for the many discussions concerning their BiCMOS PLL designs and Phil Jeffery for the discussions regarding the VCO. I would like to thank Neil Tracht for his advice on technical matters and his help digging up files that were needed to write this thesis and Clif Fonstad for his help with the revisions of this text. Finally, I would like to thank Peter Zdebel for his general support and guidance, without which this project could not have been a success.

On the lighter side, I would like to offer a few more acknowledgments. I would like to thank Glenn and Derek for playing punchout with me and thereby helping to make my time at Motorola entertaining, as well as educational. By the way guys, don't even think about beating my Xbombs scores. I'd like to thank Ray for the precise work he does or whatever. I'd like to thank Greg Hansell for always pulling some new stunt to give us something to talk about. I'd like to thank Eric Herald for always talking about it or having another story to tell. I'm now well versed in the corporate politics (bs) at Motorola. I'd like to thank Neil for ragging on Fred, because, well, Fred is a weasel and deserves it (or even if he doesn't, it was fun). Also, thanks for the stick time there flyboy. I'd like to thank Clif for helping push me to new limits by catching me unaware with his two month trip to Columbia (the country) which started just two weeks before the original due date for my thesis. The story of writing my thesis wouldn't be half as entertaining without the little complications this added (though it would be a lot easier to tell). Besides, unexpected challenges bring unexpected rewards. So if you want to send something to Bogata, and it absolutely has to get there, even in spite of an earth quake a day prior, stick to UPS and not express mail. Allow at least two days. I'd like to thank Peter for his lessons on German culture and the differences between German and American cultures. Of course, one can also learn a lot just from observing. In any case, I look forward to reading your book someday. I'd like to thank Chuck who works at Toscanini's and

doesn't even ask anymore - large French Roast with a little milk. I'd like to thank my roommate, Tom Schaub, for dealing with my six weeks of completely nocturnal work habits. It was a fun game of tag team thesis writing. Good luck with the rest of your thesis. Before I forget, thanks Mom and Dad for, well, the usual stuff, but mainly for your support and understanding. Finally, I'd like to thank my friends for serving as a constant reminder that my thesis wasn't done yet. Their hundreds of inquiries in hopes of my success were all meant in just the right way. Thanks.

Appendices

Appendix A: Schematics for phase detector circuits

Appendix B: Simulated output waveforms for phase detector

Appendix C: Schematics for VCOs

Appendix D: Simulated output waveforms for VCO

Appendix E: Schematics for loop filter

Appendix F: Behavioral model code for mixer

Appendix A: Schematics for Phase Detector Circuits

Figure A.1 - Buffer circuit.

Figure A.2 - Delay circuit.

Figure A.3 - Flip-flop circuit.

Figure A.4 - Mixer circuit.

Figure A.5 - Full phase detector, top level view.

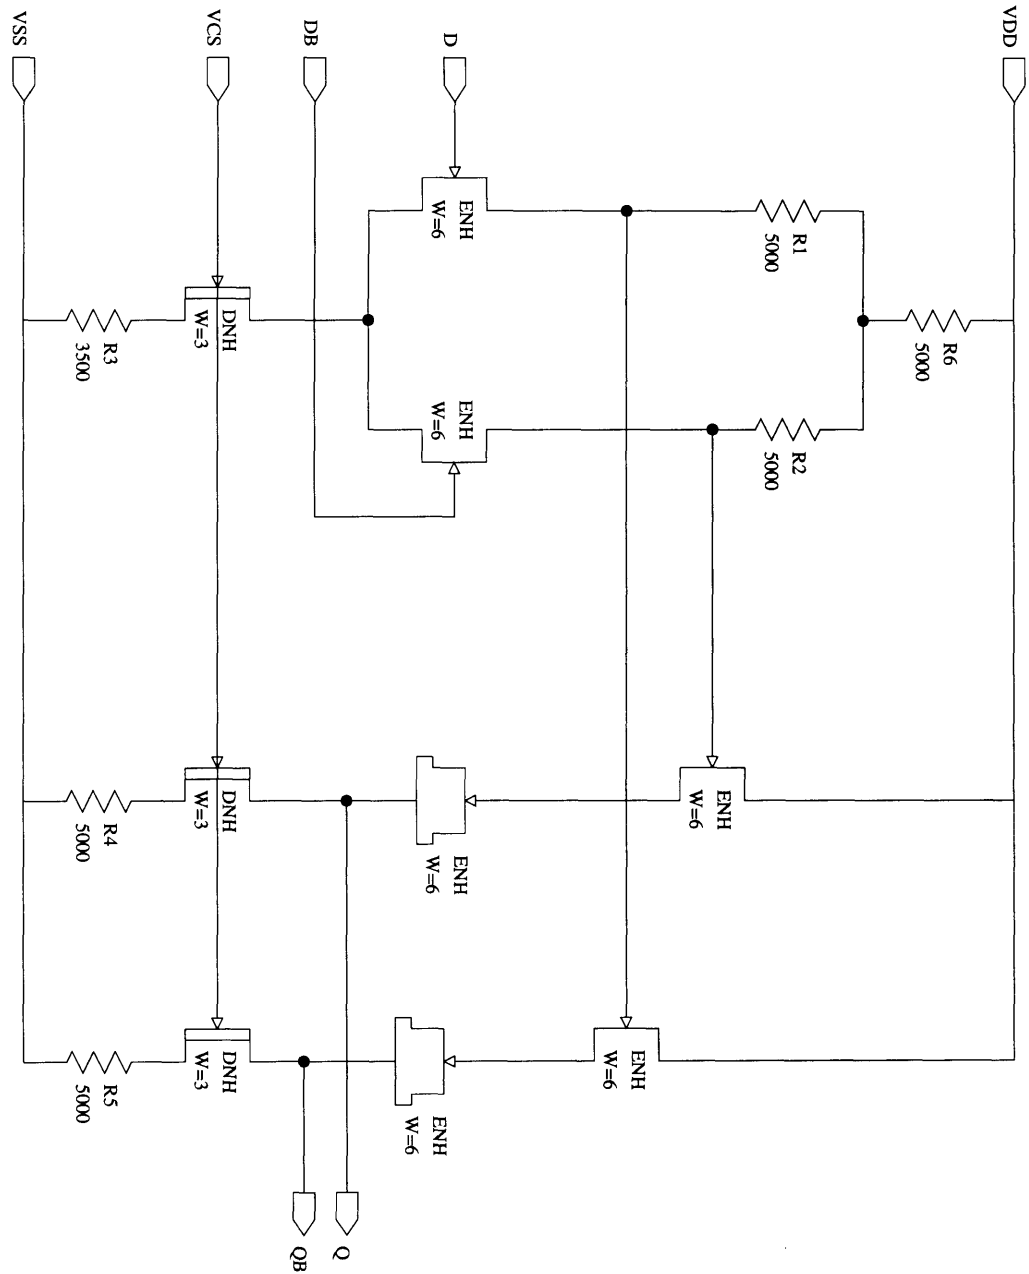


Figure A.1 - Buffer circuit.

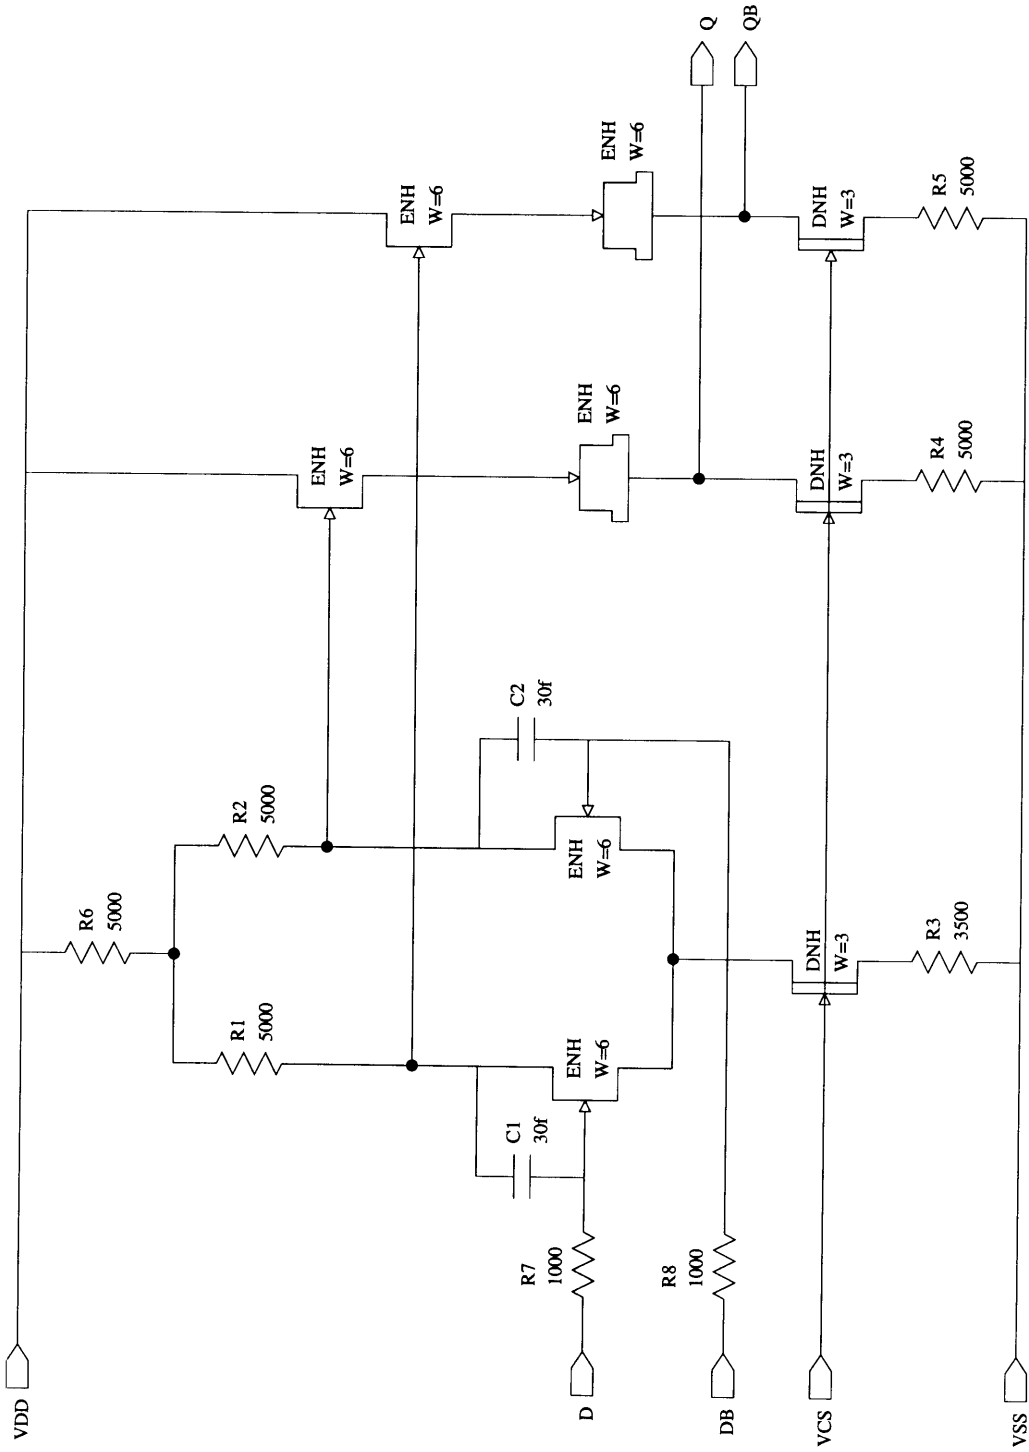


Figure A.2 - Delay circuit.

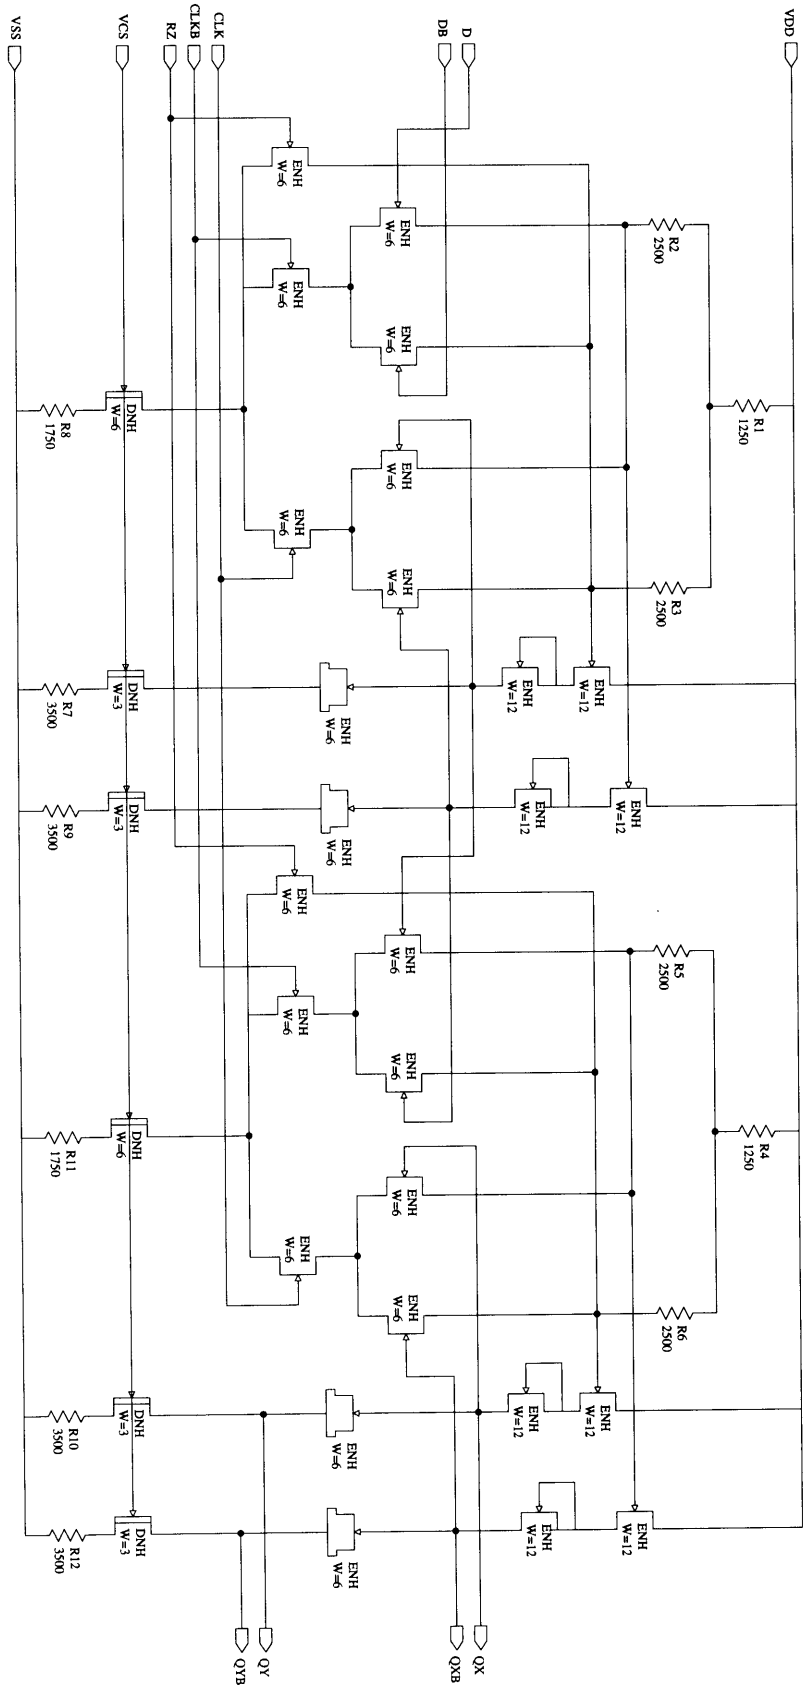


Figure A.3 - Flip-flop circuit.

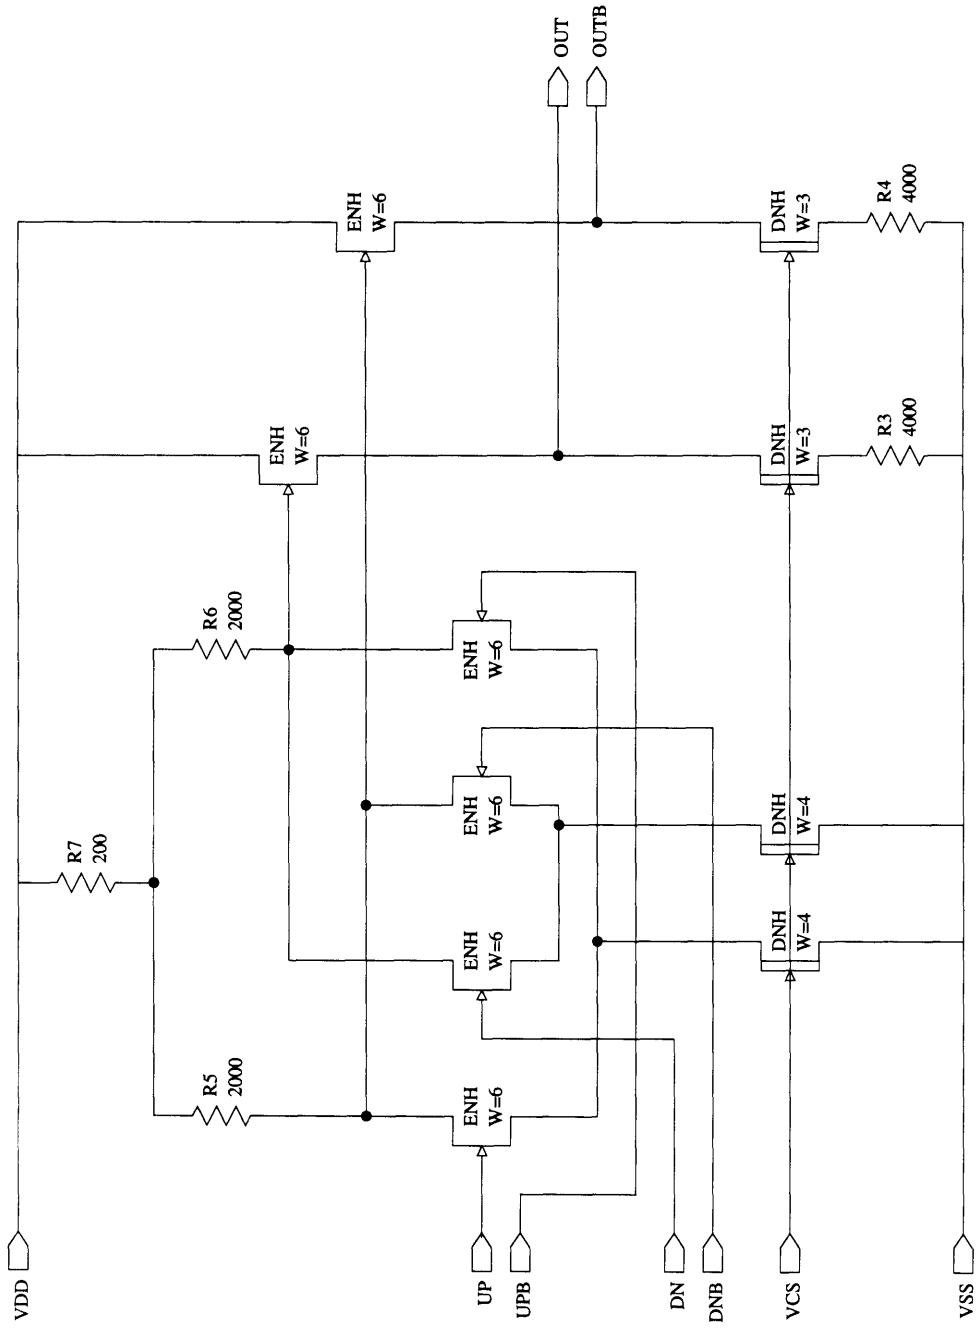


Figure A.4 - Mixer circuit.

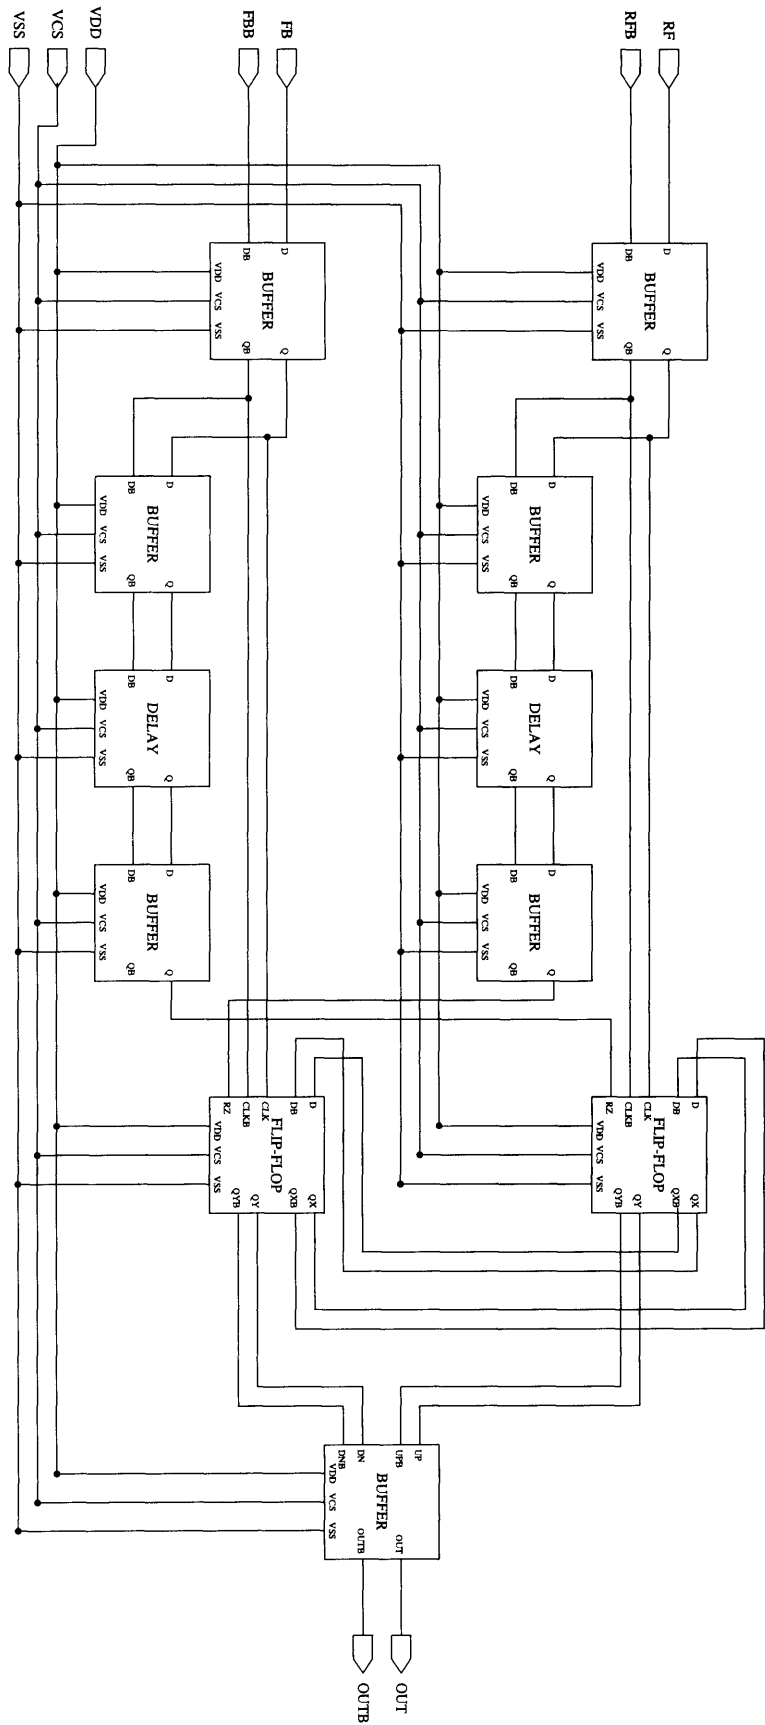


Figure A.5 - Full phase detector, top level view.

Appendix B: Simulated Output Waveforms for Phase Detector

Waveform B.1 - Phase error = 20ps.


Waveform B.2 - Phase error = 50ps.

Waveform B.3 - Phase error = 100ps.


Waveform B.4 - Phase error = 200ps.

Waveform B.5 - Phase error = 500ps.


Waveform B.1 - 50MHz inputs, Phase error = 20ps.




Waveform B.2 - 50MHz inputs, Phase error = 50ps.




Waveform B.3 - 50MHz inputs, Phase error = 100ps.



Waveform B.4 - 50MHz inputs, Phase error = 200ps.



Waveform B.5 - 50MHz inputs, Phase error = 500ps.



Appendix C: Schematics for VCOs

Figure C.1 - First version of latched-inverter circuit.

Figure C.2 - Second version of latched-inverter circuit.

Figure C.3 - Full ring oscillator, top level view.

Figure C.4 - Buffer circuit used to load oscillator.

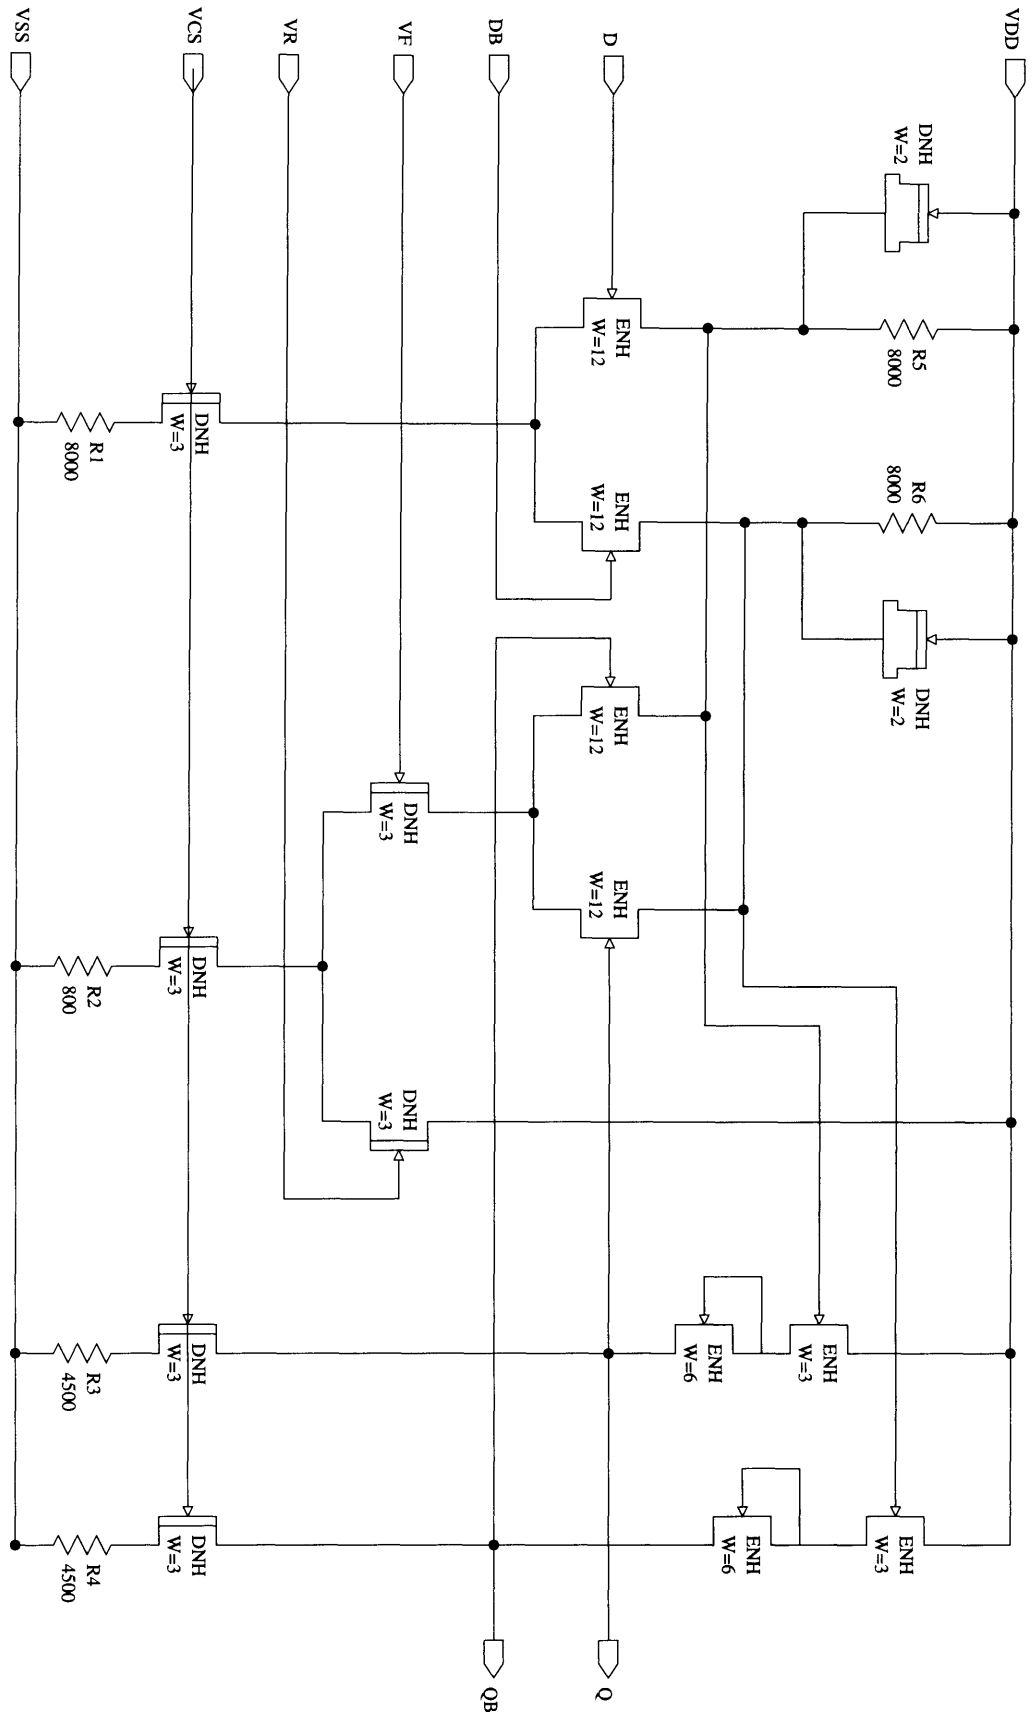


Figure C.1 - First version of latched-inverter circuit.

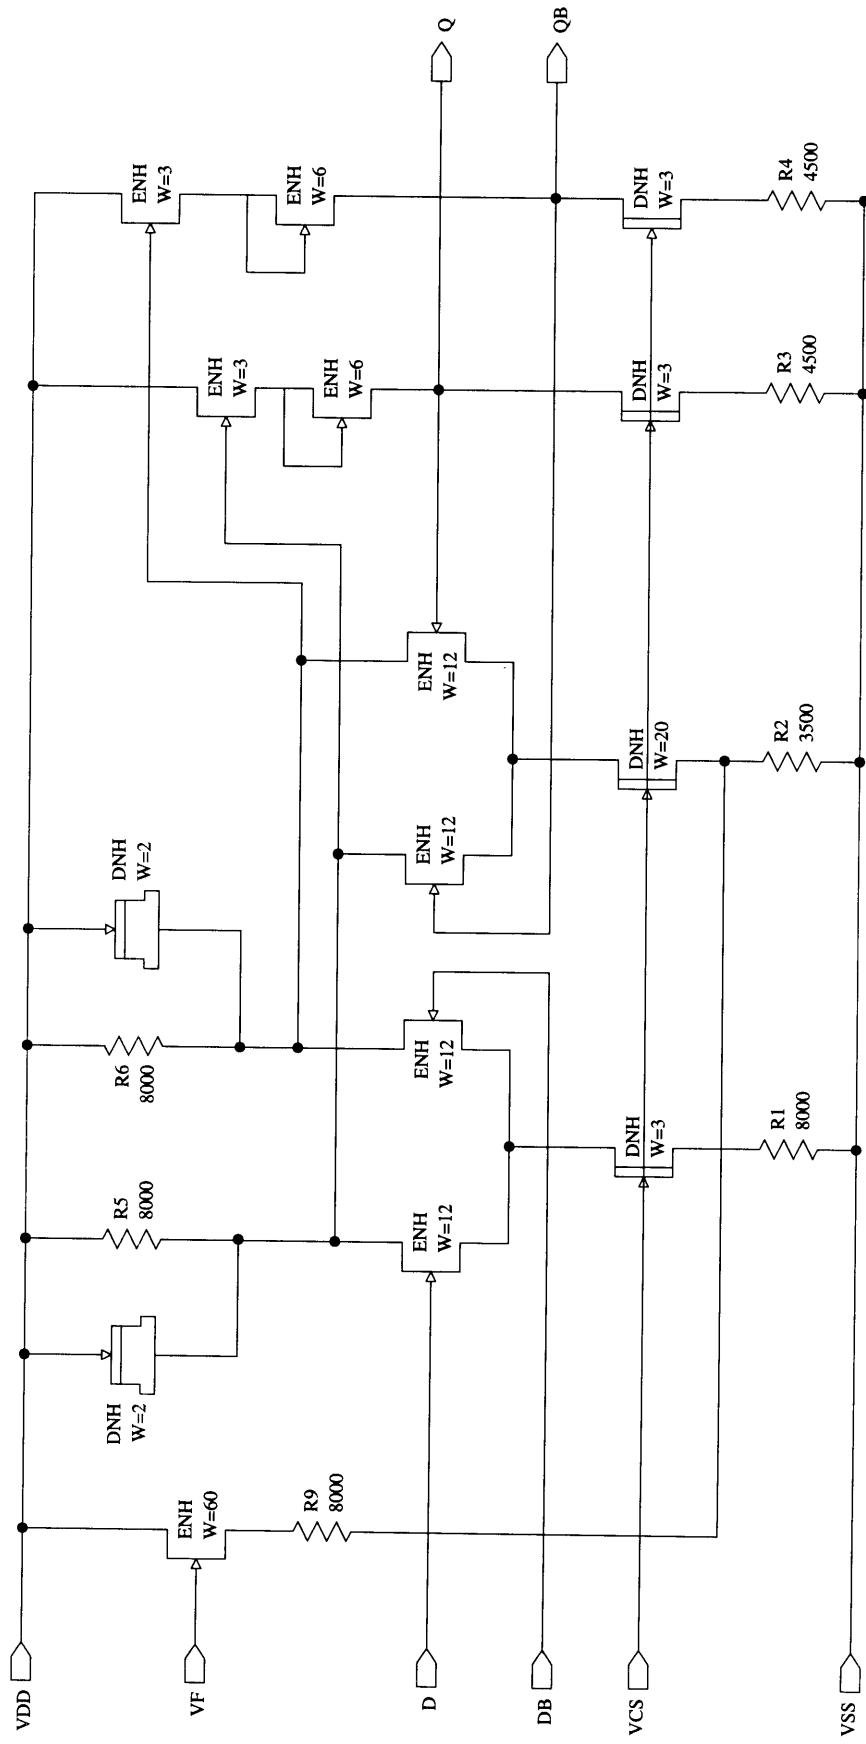


Figure C.2 - Second version of latched-inverter circuit.

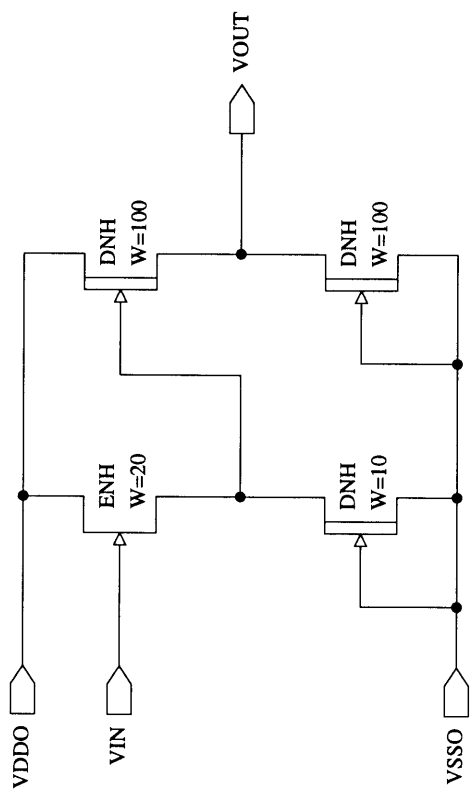


Figure C.4 - Buffer circuit used to load oscillator.

Appendix D: Simulated Output Waveforms for VCO


Waveform D.1 - Control voltage = -1.3V.

Waveform D.2 - Control voltage = -1.7V.


Waveform D.3 - Control voltage = -2.1V.

Waveform D.4 - Control voltage = -2.6V.


Waveform D.1 - Control Voltage = -1.3V.




Waveform D.2 - Control Voltage = -1.7V.



Waveform D.3 - Control Voltage = -2.1 V.



Waveform D.4 - Control Voltage = -2.6V.



Appendix E: Schematics for Loop Filter

Figure E.1 - Op amp circuit.

Figure E.2 - Full loop filter, top level view.

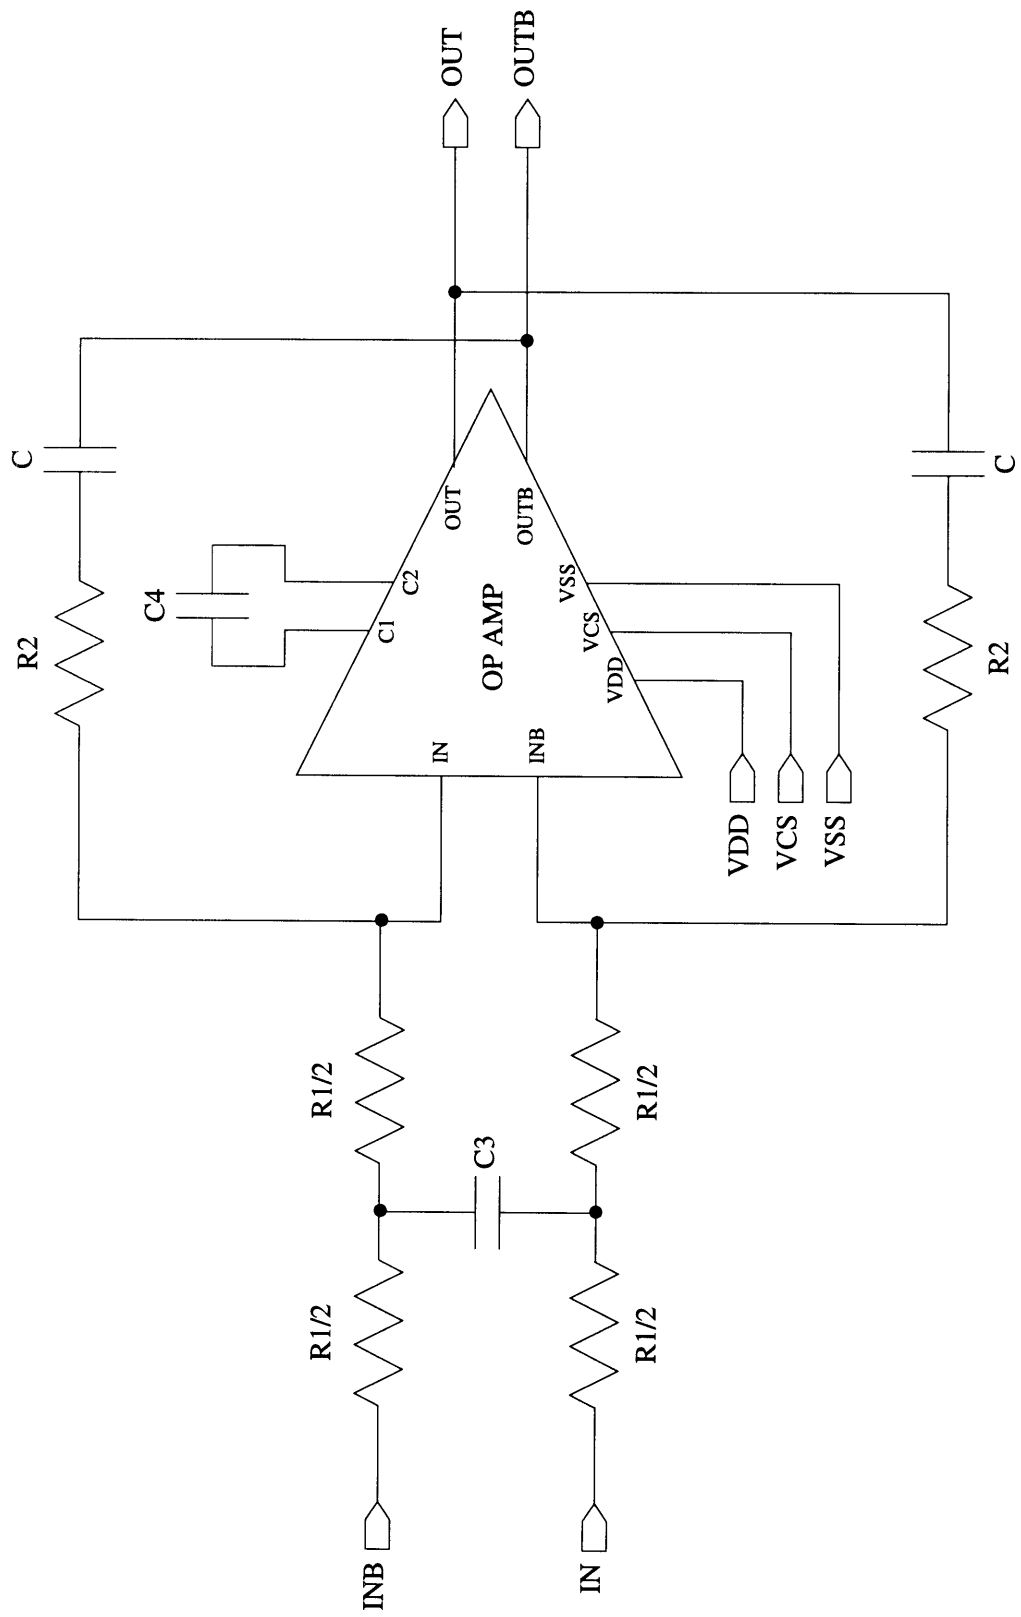


Figure E.2 - Full loop filter, top level view.

Appendix F: Behavioral Model Code for Mixer

```

# template developed for mixer circuit of PLL

element template mixerpol up dn outp outm gnd = vupdc vdowndc, voutdc,
voutdcb, fc
electrical up, dn, outp, outm, gnd

number vupdc = -1.3,      # maximum dc output voltage
      vdowndc = -2.3, # minimum dc output voltage
      voutdc = -1.8,    # output dc bias of noninverting output
      voutdcb = -1.8,   # output dc bias of inverting output
      fc      = 7.05g   # frequency of cutoff pole

{
<consts.sin
val v voutp, voutm, vout, voutb, vin, a, b, vgain, vgainb
val v vmiddc, vup, vdn, vmos, dvout
val nu gain, y1, y2, Gf, Gfb, Gx, Gxb
var i Iout, Ioutb
var nu svout, svoutb
number wc

parameters{
      wc = fc*2*math_pi
}

values{
      vup = v(up)
      vdn = v(dn)
      vin = vup - vdn
      voutp = v(outp)
      voutm = v(outm)
      gain = -2.528*(vin**4) - 2.716*(vin**2) + 3.078
      dvout = gain*vin
      vmiddc = (vup+vdn)/2
      vmos = vmiddc + 1.875
      y1 = 3.05 - 45*(vmos**2) - 200*(vmos**4)
      y2 = 3.05
      Gx = (gain/2)
      Gxb = (-gain/2)
      Gf = (y1/y2)*Gx
      Gfb = (y1/y2)*Gxb
      if (dc_domain){
          a = 1
          b = 0
          vout = dvout/2 + voutdc
          voutb = -dvout/2 + voutdc
      }
}

```

```

else if (freq_domain){
    a = 0
    b = 1
}
else if (time_domain){
    a = 1
    b = 0
    vgain = Gf*vin
    vgainb = Gfb*vin
    vout = vgain + voutdc
    voutb = vgainb + voutdcb
    if (vout > vupdc) vout = vupdc
    else if (vout < vdowndc) vout = vdowndc
    else vout = vout
    if (voutb > vupdc) voutb = vupdc
    else if (voutb < vdowndc) voutb = vdowndc
    else voutb = voutb
}
}

equations{
    i(gnd->outp) += Iout
    i(gnd->outm) += Ioutb
    Iout: a*voutp + b*(svout) + b*wc*voutp = a*vout + b*wc*Gf*vin
    Ioutb: a*voutm + b*(svoutb) + b*wc*voutm = a*voutb + b*wc*Gf*vin
    svout: svout = d_by_dt(b*outp)
    svoutb: svoutb = d_by_dt(b*outm)
}
}

```

