Low-Noise Monolithic Frequency Synthesizers for Wireless Transceivers

by

Amr N. Hafez

A thesis
presented to the University of Waterloo
in fulfilment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Electrical Engineering

Waterloo, Ontario, Canada, 2000

National Library of Canada
Acquisitions and Bibliographic Services
395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque nationale du Canada
Acquisitions et services bibliographiques
395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

Our file Notre référence

0-612-51200-2

Canada

The University of Waterloo requires the signatures of all persons using or photocopying this thesis. Please sign below, and give address and date.

Acknowledgements

All praise is due to Allah, Most Gracious, Most Merciful, Whose help and bounties are ever dominating throughout my life.

I would like to express my gratitude to my supervisor, Prof. M.I. Elmasry, for his support and supervision during my program. My deep appreciation goes to my colleagues in our research group, in particular to Ayman ElSayed for very helpful and insightful discussions. I would also like to thank Dr. Peter Edmonson of Research in Motion for his help and for bringing an industrial perspective to the work.

This work has been funded in part by NSERC and Research In Motion, Waterloo. This support is greatly appreciated.

My deep gratitude to my mother and my wife. This work would not have been possible without their continuous encouragement, patience, support and assistance.

Abstract

The wireless market is currently experiencing a huge growth in terms of both the number of users and the range of services offered. With this growth, there is an acute demand for reducing the cost and size of the device and increasing the lifetime of its battery. In achieving this, a high level of performance must be maintained in order to meet the stringent requirements of the wireless system.

A major and critical component of a wireless transceiver is the frequency synthesizer. Its performance affects many of the requirements imposed on the transceiver, such as speed, sensitivity, and adjacent channel rejection. Being active for a large proportion of the time, it has a significant effect on the battery life. In addition, it is prone to interference with other components such as the power amplifier. The objective of this thesis is to enhance the performance of the wireless frequency synthesizer through architectural innovations. PLL architectures that attempt to improve the phase-noise while meeting the spurious requirements are presented.

A key improvement that this work tries to achieve is the ability to meet the stringent phase-noise requirements using monolithic integrated VCOs. Not only does this lead to a significant reduction in the cost and size of the radio, but also to much less interference and leakage problems. Furthermore, the use of an on-chip VCO eliminates the power needed to drive off-chip 50Ω impedances.

In this thesis, we present three different phase-locked loop architectures that attempt to reduce the VCO phase-noise and enhance the synthesizer performance. The first architecture is a nested-loop PLL that achieves very wide BW while maintaining the required frequency resolution and spur rejection. The wide-BW loop, including the loop filter, is integrated on a single chip in a 25GHz bipolar process. The PLL achieves a phase-noise of -100dBc/Hz at 10kHz offset from 1GHz and consumes 9.9mA from a 3.3V supply.

The second architecture is a frequency synthesizer based on a novel closed-loop voltage-to-frequency converter. A VCO is configured in a wide-bandwidth voltage-locked feedback loop that suppresses the VCO phase-noise. By configuring this architecture in a PLL, low phase-noise frequency synthesis is possible using integrated VCOs. The proposed architecture is designed in a 0.8μm BiCMOS technology and achieves up to 20dB reduction in the integrated VCO phase-noise.

Finally, we investigate the use of subsampling in frequency synthesis. A novel PLL architecture that employs subsampling in the feedback path is presented. The loop eliminates the prescaler and greatly reduces the division ratio. The architecture is particularly suitable for DDS-driven PLLs since it relaxes the requirements on the DDS thus further reducing the power consumption. The advantages of the architecture are highlighted and the design considerations discussed. System-level simulations are also presented.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The wireless market is currently booming. The convenience that wireless devices offer is creating an increasing demand for them. Market surveys show that more than 20,000 people join the cellular phone system in the United States daily. The wireless revolution is not only confined to cellular phones, but also to broader applications such as wireless LANs and wireless mutlimedia devices. This huge demand is motivating manufacturers to provide increasingly competitive devices. The key figures-of-merit of a wireless device are its price, size and battery life. The latter is of particular importance, since it is important for the user to have a device that will continue to function for extended periods of time. Manufacturers are continuously trying to improve these three parameters. The present goal is to reduce both power consumption and the price of cellular phones by 30% each year. From the electronics point of view, price could be decreased by using low-cost technologies (tend to be CMOS), size could be reduced by minimizing the number of chips and off-chip components and battery life could be extended by decreasing

the power consumption. Another glorious goal is to build a programmable device that can be used in different parts of the world over different frequency bands and standards.

One of the major blocks of the wireless radio is the frequency synthesizer. The frequency synthesizer is responsible for translating the desired signal from the RF band down to baseband in the receive path, and in the other direction in the transmit path. Sometimes it is used for modulation as well. In most wireless transceivers, the frequency synthesizer is also responsible for selecting the desired channel. Given the fact that the desired channel lies at a high frequency and is usually surrounded by much stronger and very close adjacent channels, it is no surprise that the frequency synthesizer has to meet very stringent requirements.

The frequency synthesizer affects almost all the performance aspects of the radio. Not only does it constitute a significant portion of the cost and size of the radio, but it also contributes to the spurious emission of the transmitter and is thus subject to governmental regulations. It affects the sensitivity of the receiver and its ability to function properly in the presence of large blocking signals. In addition, it is one of the main power drains in the radio. Hence, improving the performance of the frequency synthesizer has a great impact.

The issues related to the performance of the synthesizer can mostly be attributed to the component at its heart: the voltage-controlled oscillator. Due to conflicting requirements on the synthesizer, it is usually necessary to use an off-chip discrete VCO. This leads to extra board space, hence size, extra cost due to the additional component, extra power consumption and more interference problems. It is an

ambitious goal to be able to use monolithic integrated VCOs.

The main obstacle in integrating the VCO is the poor quality factor of monolithic inductors that results in poor phase-noise performance. The obvious way to integrate the VCO is thus to attempt to improve the technology so as to boost the inductor quality factor. Indeed IC technologies geared towards RF application can achieve today quality factors up to 10-12 as opposed to 4-6 in ordinary technologies several years ago. The improvement however is small and in most cases insufficient. Even bondwire inductors that suffer from integration problems do not provide a much larger quality factor, and are not able to meet the specifications in most cases.

An alternative method to integrate the VCO is through architectural-level innovations. A wide-bandwidth phase-locked loop can track the noise of the VCO and attenuate it within its bandwidth. The key to this method is developing a creative wide-bandwidth loop that attenuates the VCO phase-noise while not sacrificing the other parameters such as the spur level and input phase-noise. This has to be done in a power-efficient manner and with minimal extra hardware.

This latter method is the focus of this work. We present three distinct architectures for frequency synthesis that achieve wide BW with low phase-noise and spur levels. The emphasis of the work is on the architectural-level rather than the circuit and implementation level.

In chapter 4 we present a novel nested-loop PLL architecture that uses a low-frequency secondary VCO to drive a high-frequency PLL. The architecture is simple and achieves a substantially high BW without degrading the spur performance or settling time. The extra hardware is small, and could be eliminated by slight mod-

ifications to the superheterodyne transceiver architecture. A completely integrated implementation achieves the lowest reported phase-noise of an LC-VCO at 10kHz offset.

In chapter 5, we propose an improvement over the nested-loop architecture in which the secondary VCO is eliminated through a novel voltage-locked loop that uses frequency-to-voltage converters. The result is a closed-loop feedback system that converts the voltage signal into a frequency. When implemented, close to 20dB reduction in the phase-noise of a ring oscillator was achieved.

A third synthesizer architecture is presented in chapter 6, where subsampling is used in the feedback path of a PLL to downconvert the VCO signal. The result is a great reduction in the division ratio and hence the reference spurs and input phase-noise. This architecture is particularly suitable for DDFS-driven PLLs as it greatly relaxes the requirements on the DDFS and DAC.

The thesis starts with a brief review of wireless transceiver architecture in chapter 2, followed by an overview of frequency synthesis in chapter 3. In chapter 7 we summarize the conclusions of our work and suggest future work.

Chapter 2

The Wireless Transceiver

2.1 Overview

2.1.1 Architecture

Figure 2.1 shows a conventional wireless architecture. In the receiver, the signal is amplified and mixed down to the IF stage where the desired channel is selected using an off-chip highly selective SAW filter. The signal is then mixed down to baseband where it is converted to the digital form and processed by the DSP.

In the transmitter, the digital signal is converted to the analog domain via a DAC and upconverted to RF, where it is amplified by the power amplifier and fed to the antenna. The LO signal needed is usually synthesized from a crystal oscillator by a PLL.

The above architecture is the classical superheterodyne receiver that has been widely used in FM radio receivers for a long time. It has proved to be a robust

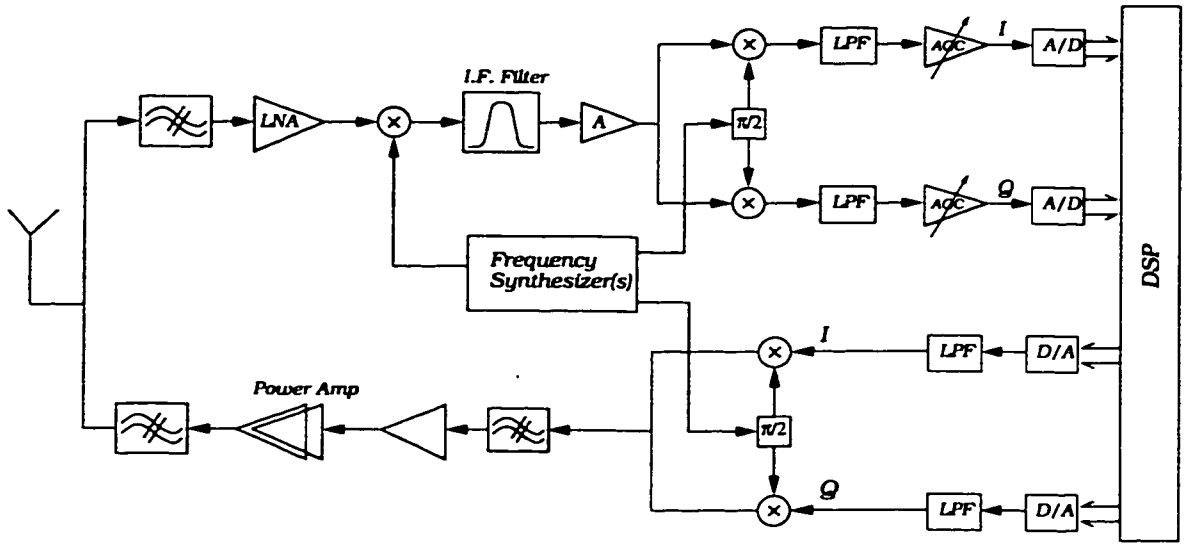

Figure 2.1: The superheterodyne Transceiver.

architecture that is able to meet very stringent standards. However, while this architecture is well suited for FM radio receivers, it incurs many disadvantages when employed in a portable device, some of which present a bottle neck for further size and power reduction. This is explained in section 2.2.2.

2.1.2 State of the Art

Table 2.1 shows some of the wireless transceivers in the market. None of these chipsets contains baseband processing. From this table, several things can be deduced. First, the superheterodyne architecture, shown in Fig. 2.1, is used in all ICs except in [6] where the ISM band has relaxed requirements. This confirms the robustness of this architecture. On the other hand however, the number of chips and off-chip components is also relatively high. A state-of- the-art radio will have 2-3 chips (one for baseband processing) in addition to off-chip power amplifiers,

Ref.	Company	Standard	Architecture	No. of chips	Tech.	Off-chip Components
[1]	AT&T	GSM	Super- heterodyne	1	12GHz Bipolar	LNA, PA, IF-SAWs, PLL filter
[2]	Philips	GSM	Super- heterodyne	2	13GHz Bipolar	PA, LOs, IF-SAWs
[3]	Hitachi	GSM	Super-het (double- IF)	1	15GHz bipolar	PA, VCOs, PLL, IF SAWs
[4]	National Semi.	DECT	Super- heterodyne	1	0.5μm BiCMOS	LNA, PA, IF-SAWs, VCO
[5]	Mitsubishi	PHS	Super- heterodyne	2	20GHz BiCMOS	VCO, RF & IF SAWs
[6]	Rockwell	ISM, cordless	Direct Conversion	1 Rx, 1 Tx, 1 FS	25GHz Bipolar	BB filter, PLL filter
[7]	Siemens	0.8-1.2GHz	Super- heterodyne	2	25GHz Bipolar	PA, IF-SAWs

Table 2.1: Commercial wireless transceivers.

SAW filers, oscillators and PLL loop filters. Another interesting aspect in table 2.1 is the technology used for implementing the wireless radio.

2.1.3 Technology

The digital part of the transceiver (the DSP) is implemented using the CMOS technology. This is due to the fact that the CMOS process is a low-cost process, it achieves higher integration levels and because CMOS has a low standby current (a very important factor in battery-powered portable devices). Due to this trend, it is very beneficial to implement the rest of the transceiver in the CMOS technology

to take advantage of its low cost and, more importantly, to be able to integrate as much of the analog circuit as possible on the same chip with the digital part, thereby reducing the chip count and size.

When partitioned according to the operating frequency, the analog part of the transceiver can be considered to consist of two distinct parts, the low-frequency baseband part and the high frequency front end (LNA, first mixer and possibly the IF stage). In the baseband part, CMOS technology has been used in market products [8] and this trend is continually increasing.

The use of CMOS in the high frequency part, however, has been mainly confined to research effort [9–14]. In addition to a few GaAs products [15,16], bipolar is the dominant technology used at high frequencies [1,2,4,6,17–19]. The reason for this is the higher f_t and transconductance (g_m) of the bipolar transistor with respect to the NMOS for a given bias current.

Nevertheless, with the continuous CMOS technology advancement and scaling of the minimum length along with the large research in this field, it seems that it will not be long until an all-CMOS radio becomes a reality.

As shown in Fig. 2.1, the wireless transceiver can be divided into three main parts: the receiver, the transmitter and the frequency synthesizer. Following is a discussion of the design considerations and available architectures of each part.

2.2 The Receiver

2.2.1 Design Considerations

One of the most important parameters of the wireless receiver is its sensitivity. The sensitivity of the device is defined as

$$S = 10log(V_{min}^2/50\Omega) \tag{2.1}$$

where V_{min} is the minimum detectable r.m.s. voltage. This minimum voltage will depend on the noise level of the receiver. The minimum sensitivity will vary according to the communication standard but has a nominal range of -80dB_m to -105dB_m [8]. The GSM standard specifies a sensitivity of -102dB_m (1.8μV_{rms}). Obviously, this is a very stringent requirement imposed on the receiver.

The implication of the required high level of sensitivity is that the receiver must have a very low noise figure. Noise figure is defined as the ratio of the input signal-to-noise ratio (SNR) to the output SNR, or,

$$NF = \frac{(S/N)_i}{(S/N)_o} \tag{2.2}$$

From (2.2), the noise figure could be kept low by maintaining a low noise level (N_o) and a high output signal (S_o). The latter necessitates adequate gain. The noise figure of the receiver is determined mainly by its front-end, namely the LNA and

the mixer, according to (2.3):

$$NF = NF_{LNA} + \frac{NF_{mixer} - 1}{Gain_{LNA}} \tag{2.3}$$

These two components must thus have adequate gain and a sufficiently high current to maintain a low noise figure. A typical noise figure of a wireless front end will be around 2-3dB [7, 11, 20].

To complicate the designer's task, the receiver must not only be able to distinguish very weak signals, but must be able to simultaneously cope with very strong in-band signals. In the GSM standard, the front-end must cope with in-band signals as large as -23dB_m (80dB larger than the weakest signal) [21]. In addition, the channel directly adjacent to the desired channel can be 40-60dB larger [22].

The existence of a very weak signal in the presence of very strong signals has two implications. First, it necessitates a very large dynamic range. The dynamic range of the receiver is a measure of the ratio of the strongest signal that can be processed without saturation to the weakest signal that can be detected, which is usually equal to the noise level. This is represented as,

$$DR = 20log(\frac{V_{max}}{V_{min}}) \tag{2.4}$$

The dynamic range of the receiver is typically around 100dB [23]. This large dynamic range is needed at least for the front-end LNA and mixer that process the whole composite RF signal. As the signal is progressively filtered, the dynamic range requirement on the subsequent stages could be relaxed. In some wireless ar-

chitectures though, there is no filtering until baseband (direct-conversion) and the full dynamic range must be met by baseband components as well. This is one of the very challenging design considerations and has a large effect on the design of the wireless architecture.

The second implication of the presence of strong and weak signals simultaneously is the stringent linearity requirement imposed on the receiver. This results from the fact that two sinusoids that undergo any nonlinear operation will mix together producing tones at the sum and difference frequencies. This is depicted in Fig. 2.2. The value of the resulting spurii in Fig. 2.2 are called third-order intermodulation products (IM_3). Thus, if the nonlinearity of the system is not low enough, the desired channel might be severely contaminated by the intermodulation product of the two adjacent large channels.

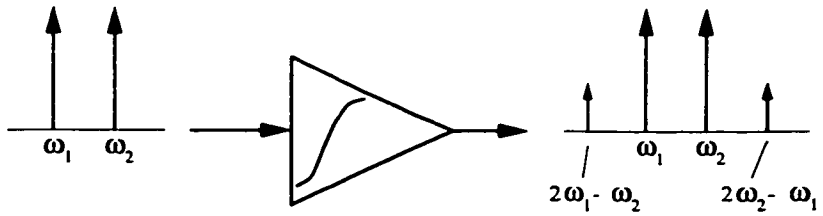

Figure 2.2: Third-order distortion effect.

Linearity of the system is quantified by the third-order intercept point (IP_3). The IP_3 of an element is depicted in Fig. 2.3. Here, the value of the input power is changed and each time the value of the fundamental and the third-order distortion component are measured and plotted. The two lines are extrapolated to yield the

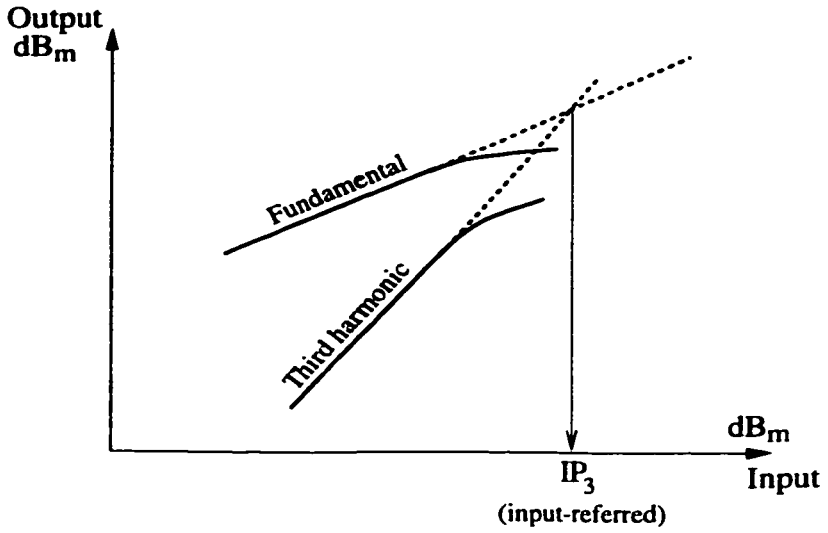

Figure 2.3: Third-order intercept point (IP_3).

IP_3 point. Referring to Fig. 2.3, IP_3 would be defined as

$$IP_3 = S_i|_{a_3=a_1} \tag{2.5}$$

This means that an input power of IP_3 would result in a third-order distortion that is as powerful as the fundamental. Of course, a large IP_3 indicates a more linear system. Each standard imposes a different linearity requirement in the form of a minimum IP_3. For DECT this value is -27dB_m [14], and is more stringent for GSM and other standards.

2.2.2 Architectural Variations

As mentioned above, the superheterodyne receiver of Fig. 2.1 is the architecture most widely used in market products due to its robustness. However, this architecture imposes some limitations on integration and power reduction of the wireless

radio.

The IF frequency in the superheterodyne receiver is around 10% of the RF (about 100MHz). It would be difficult to make it less than that since this will tighten the requirements on the image-rejection RF BPF that precedes the LNA. The filter that selects the narrow channel at such a high frequency would have a quality factor much higher than that attainable in an integrated realization. Hence, the necessity of the SAW filter. It is important to note that the presence of the SAW filter not only results in an increased form factor, but also in an increase in power dissipation due to the need to drive the off-chip 50Ωs.

In order to solve the integration problem of the IF stage, architectural innovations are necessary. Below is an outline of some receiver architectures that attempt to alleviate this problem.

Direct-Conversion (Zero-IF)

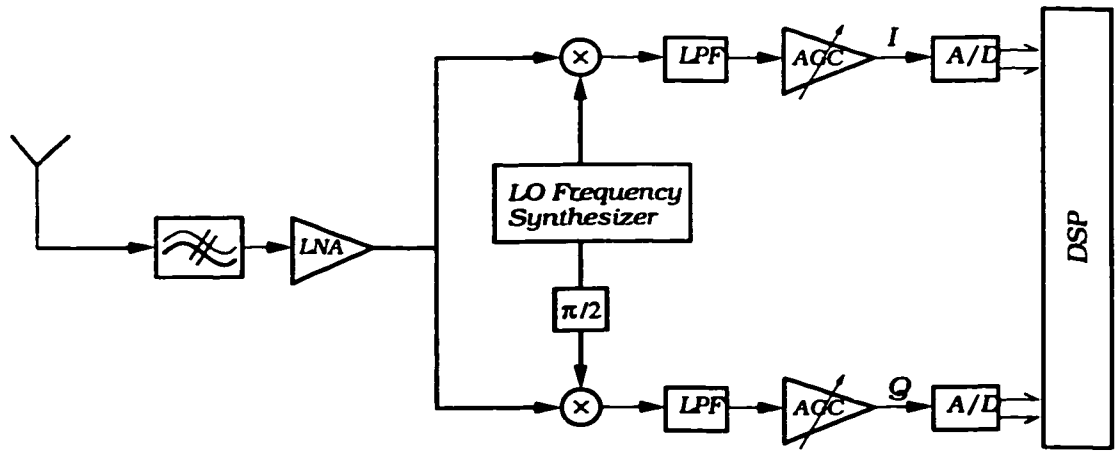

Figure 2.4: The direct-conversion receiver.

The direct conversion architecture is shown in Fig. 2.4 [24]. In this architecture, the LO frequency is tuned to exactly the RF frequency such that the desired channel is mixed down to d.c. (hence the name zero-IF). This has two important consequences. First, there is no image to the desired channel. This relaxes the requirements on the RF BPF, hence enhancing the possibility of integration. Second, and more significantly, the high frequency IF stage is completely eliminated. Any further signal processing, including channel selection, will be performed at the low frequency baseband part where integrated filters and components are readily available, hence much higher levels of integration are possible. This also leads to reduced power consumption because most of the signal processing is performed at low frequencies.

The direct conversion receiver has several drawbacks though. First, since the LO is tuned to exactly the RF, any leakage from the LO to the antenna will result in spurious radiation that is in the band of other receivers. The FCC imposes limits on the LO leakage typically between -60 to -80dB_c [25].

The most severe disadvantage of the direct conversion architecture, that is prohibiting its use in high performance radios, is the d.c. offset and the baseband distortion. In any receiver architecture there exists a relatively large d.c. offset. This offset results from mismatch between the I and Q paths and from LO leakage to the RF port and then self-mixing to d.c. This d.c. offset could easily be around 10mV [26]. In addition to that, even order distortion from LNA and mixer nonlinearity result in baseband spurii that can contaminate the desired signal. Also, the flicker noise of the active elements will lie in the band of the desired channel.

Given the limited gain of the LNA and mixer ($\approx 25dB$), we realize that the signal might not be strong enough to overpower the d.c. distortion and hence a severe SNR degradation can occur.

Solutions for the problem of the d.c. offset have been proposed. Among them is the insertion of a large off-chip decoupling capacitor between the mixer and the subsequent stage. This capacitor will form a HPF with a notch at d.c. hence removing the d.c. offset. Unfortunately, unless the system employs FSK modulation (such as pagers), this HPF tends to remove part of the channel spectrum as well, hence removing part of the information and degrading the SNR. The cutoff frequency of this HPF cannot be made too close to d.c. else the channel characteristics will be adversely affected by the group delay [27]. Another solution for this problem is to use the DSP for d.c. correction. Here the DSP periodically estimates the d.c. level of the signal and feeds it back to the analog domain where it is subtracted from the signal [27].

As mentioned above, pagers, which are a low performance systems that use FSK, do not suffer greatly from the d.c. problem. This is why, the direct conversion receiver architecture has been widely used in pagers by manufacturers such as ITT, NEC, Philips, Hitachi and Matsushita [24, 28]. In digital cellular phones, the use of direct conversion architectures has not been as widespread. Nevertheless, it has been employed by Alcatel in DECT [8] and Rockwell in the ISM band [6] and others [28].

Low-IF

In order to alleviate the problems in the direct conversion receiver, the low-IF architecture has been proposed. The idea behind the low-IF architecture is that if the signal is translated to a low frequency (100kHz), but not d.c., then the problems of the direct conversion architecture will be alleviated while still maintaining its advantage of eliminating the high frequency IF stage and processing the signal at low frequencies where integrated solutions are possible. Thus, in this architecture, the LO is tuned close to the RF and the signal is converted to 100kHz or so. The desired channel is then selected by a BPF or the whole/part of the spectrum is digitized and channel selected digitally in the DSP [26, 29].

Whereas in the zero-IF case, the channel has no image, the problem with this architecture is that the desired channel will have an image that is very close (200kHz away) that cannot be rejected by the RF BPF. Due to the quadrature nature of down conversion, the signal and its image are converted down to the same frequency but are in quadrature and can still be separated. The amount of image rejection though depends on the matching between the I and Q paths and on the accuracy of the 90° phase shifter that produces the quadrature component of the LO. The problem in this architecture is that the image channel is not attenuated by the RF BPF and hence could be more than 60dB larger than the desired channel. This means that a higher level of image rejection is needed and thus less mismatch and phase shift inaccuracy tolerance. While a phase error of 3° might be tolerated in the zero IF, an accuracy of less than 1° is necessary in this case.

Wideband-IF

Another problem that has been mentioned earlier is the need for an off-chip oscillator. The reason for that is that on-chip inductors have low quality factors and thus unacceptable phase-noise. The receiver architecture has been shown to greatly influence this problem.

As discussed in section 3.3.5, the bandwidth of the PLL must be made as large as possible to inhibit the phase-noise of the VCO. In a classical PLL, the output should be tunable to select the channel and thus the reference frequency is equal to one channel spacing. This results in a large division ratio and a small bandwidth. The wideband-IF (WBIF) architecture attempts to solve this problem through architectural innovation.

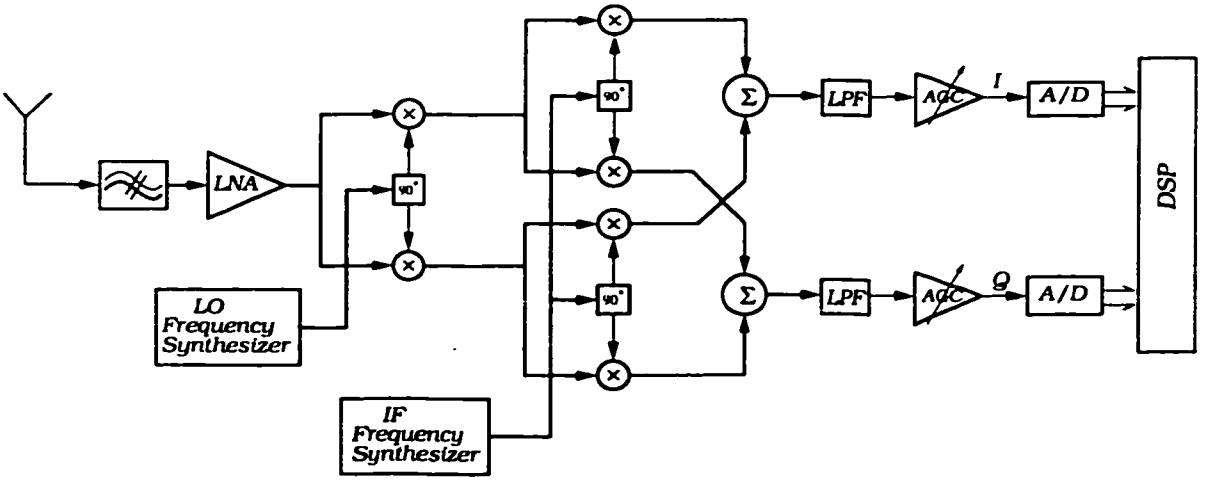

Figure 2.5: The wideband-IF receiver.

In the WBIF architecture [14], the signal is mixed down to baseband (or to a low frequency as in the case of low-IF) in two stages. First, the signal is mixed down to IF, then the *whole spectrum* is mixed down again to baseband with *no*

filtering. Obviously, since the image channel is not filtered out, the downconversion must be carried out by an image reject mixer as shown in Fig. 2.5.

In this architecture, there is two LO oscillators. The first LO is made fixed, and thus could have a much larger reference frequency and bandwidth. Since the increased bandwidth inhibits the close-in phase-noise, the phase noise requirements on the VCO is relaxed and a integrated VCO can be used. The second LO must be tunable and will have a reference frequency of one channel spacing and thus reduced bandwidth. However, this oscillator operates at a much lower center frequency and thus its phase-noise level is correspondingly lower.

The WBIF architecture attempts to achieve VCO integration, however, it still suffers from the same problems as the direct conversion (or low-IF) architectures.

2.3 The Transmitter

2.3.1 Design Considerations

Referring to the transmitter of Fig. 2.1, the signal is synthesized in the DSP, and then converted to the analog domain via a DAC. The signal is then up-converted to the RF frequency by the up-mixer, after which the signal is amplified by the preamplifier and fed to the, usually off-chip, power amplifier.

We note that many of the aspects that complicate the receiver design are not present in the transmitter. First, the signal is synthesized in the DSP and DAC. Hence, the signal generated is much larger than the noise level of the mixer and thus the up-mixer noise figure is not critical. Also, due to the absence of interfering

signal and due to the constant level of the I and Q signals, the dynamic range of the up-mixer is also not critical. The inherent quality of the signal that the transmitter is dealing with makes the design of the transmitter a much simpler task than that of a receiver.

An important consideration of the transmitter, though, is the amount of spurii transmitted. Spurs are discrete frequency components that appear at the spectrum of the transmitted signals. The main sources of these spurs are the LO leaking through the antenna thus creating a tone at the LO frequency, mismatches between the I and Q paths that cause one of the sidebands to leak and the power amplifier nonlinearities. The communication standard imposes a requirement on the maximum allowable spurious emission of the transmitter, and this requirement is one of the principle design parameters of the transmitter.

The main sources of spurii are the up-conversion mixers and the PA. Apart from the wanted sideband, the outputs of the mixers contain: 1) an unwanted sideband component due to the inaccuracy of the phase shifter and the mismatch between the I and Q paths, 2) an LO component due to LO leakage and d.c. offsets, and 3) a third-order intermodulation product due to nonlinearities in the mixer. This means that the linearity of the mixer is an important design consideration for achieving sufficient spurious rejection.

The other major design challenge in the transmitter is the power amplifier which is the main power drain in the system and a major contributor to the transmitter spurii.

2.3.2 Architectural Variations

Compared to the receiver, there is not many variations of the transmitter architecture. Many transmitters employ the direct conversion architecture which is shown in Fig. 2.1 [8,30]. In this architecture, the baseband signal is directly up-converted to RF. Due to the quadrature up-conversion, the image (unwanted sideband) is automatically rejected. The main problem with this architecture is the severe interference problem. Since the LO is tuned to the RF frequency, the powerful PA output is very close in frequency to the local oscillator and this causes the VCO frequency to drift. In such an architecture the required isolation between the PA and the VCO could be impractically high (more than 90dB [31]). In order to solve this problem, the VCO and PA frequencies should be sufficiently distant. This is accomplished by indirect up-modulation [2, 5], shown in Fig. 2.6 or by using an offset mixer [1] as shown in Fig. 2.7. In both cases, two VCOs are used such that neither is tuned to the output frequency. The drawback here however, is the need for more VCOs and more inter-stage filtering.

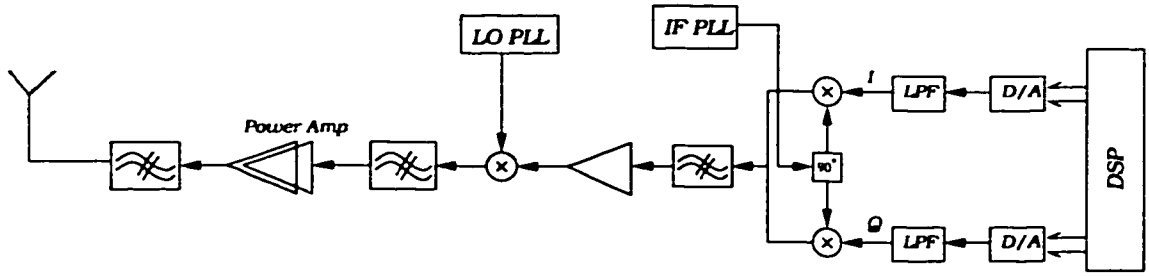

Figure 2.6: I-Q transmitter: Double conversion.

Another different transmitter architecture is the open-loop modulation architecture. In this architecture, after tuning the VCO to the channel center frequency,

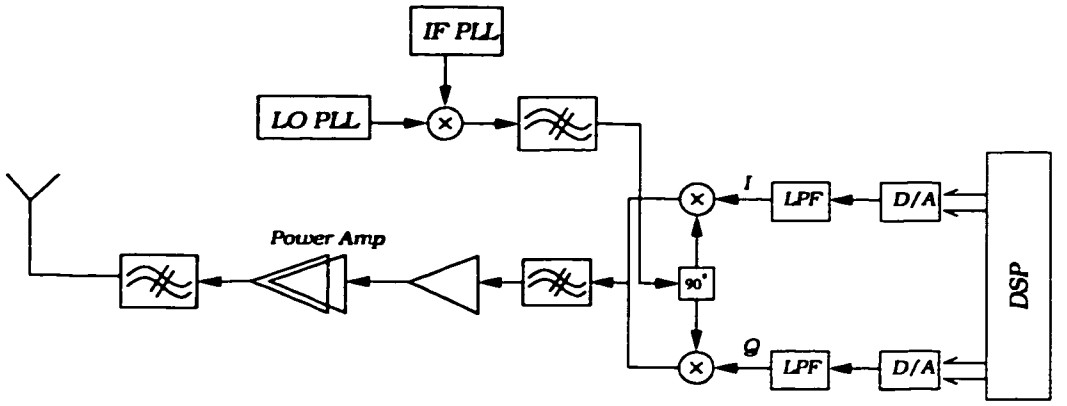

Figure 2.7: I-Q transmitter: Offset mixer.

the PLL loop is broken and the VCO is directly modulated by the baseband signal. The output is directly fed to the PA. This architecture is reported to have near ideal spurious rejection [4]. It also does not suffer from the frequency pulling problem since the instantaneous frequency of the PA is identical to the VCO. The drawback of this architecture is that the linearity of the VCO and its temperature dependence might be more critical. Also, in this scheme, the VCO will be either used in the transmit or receive mode (unless two VCO are used). Obviously this is not suitable for frequency division duplex (FDD). The open loop architecture has been employed in the DECT standard, which is TDD, in [4] and [18].

2.4 The Frequency Synthesizer

The third major part of the wireless transceiver is the frequency synthesizer. The frequency synthesizer, as shown in Fig. 2.1, is responsible for synthesizing the sinusoidal signals that are used to translate the frequency bands of the modulated/demodulated signal. The frequency synthesizer in a typical wireless radio

will synthesize one high frequency signal and one or two low frequency signals.

In order to synthesize an accurate frequency, the radio will have one crystal oscillator that produces a waveform with an exact pre-determined frequency that has no practical dependence on temperature. This crystal signal usually exhibits a very pure spectrum with very little jitter. The most common way of frequency synthesis is to use a phase-locked loop to lock the frequency of a voltage-controlled oscillator to a reference derived from the system clock. The PLL has to achieve the required accuracy and speed, while producing a signal with acceptable spectral purity.

Since the bulk of this work deals with frequency synthesizers, a comprehensive overview of frequency synthesizers, their requirements and different architectures is presented in the following chapter.

Chapter 3

The Frequency Synthesizer: An Overview

The frequency synthesizer is a critical component of the wireless radio. It is responsible for producing sinusoidal signals that translate the frequency bands in both the receive and transmit path. Hence, it is active for a large proportion of time, making it one of the power hungry parts of the radio. In addition, interference with other components can distort the synthesizer signal. The frequency synthesizer affects the sensitivity of the receiver and the quality of the transmitter.

In this chapter we review the general considerations related to the frequency synthesizer in a wireless transceiver. We then briefly discuss the operation and noise-performance of the VCO, the most critical part of the synthesizer. The operation of the widely used phase-locked loop is then explained and several architectural variations presented.

3.1 Requirements

3.1.1 Accuracy

The RF bands of transmission and reception are usually in the 1-2GHz range or more. The channel spacings can be as low as 12.5kHz for Mobitex [32] or 30kHz for IS-54. Under these circumstances, the frequency synthesizer must achieve a very high degree of accuracy. In the GSM system, the accuracy requirement is 0.1ppm (90Hz at 900MHz) [33]. Such stringent accuracy levels must be maintained over full temperature variations.

Since the system crystal acts as the clock to the entire radio, its accuracy will determine the accuracy of the frequency synthesizer. The crystal is usually temperature-compensated. Yet, automatic frequency control (AFC) [33] is often necessary to compensate for crystal aging.

3.1.2 Resolution

In the vast majority of wireless transceivers, channel selection occurs at the LO frequency, which is close to the RF band. Thus, while operating at high RF frequencies, the LO synthesizer must be tunable at steps equal to the channel spacings. Systems such as DECT and IS-95 have fairly large channel BWs (1.73MHz and 1.25MHz respectively). In other cellular systems, however, the channel spacings are much less (e.g. 200kHz in GSM).

As will be seen later, there is a tradeoff between the oscillator resolution and the remaining requirements such as phase-noise and settling speed. Therefore, it

becomes more and more difficult to satisfy all the requirements on the frequency synthesizer as the channel spacing gets smaller.

3.1.3 Spurious Signals

An important concern regarding the frequency synthesizer is the spurious content. The output of the frequency synthesizer is ideally a single sinusoidal tone located at the proper frequency. Practically, however, other frequency components appear.

First, the output of the VCO usually exhibits large harmonics of the fundamental (2^{nd} harmonic, 3^{rd} harmonics, and so on). While these harmonics can cause undesirable effects in some cases, they generally do not constitute a serious concern. This is because, being at a multiple of the fundamental, they do not cause any jitter in the zero-crossings of the signal.

In addition to the harmonics, however, other spurs can appear in the output spectrum. These spurs are objectionable in both the transmit and receive case. In transmission, the spurs can lie in the band of a nearby receiver and block the signal. Regulatory bodies usually impose a stringent limit on the spurious transmission.

During reception, any spur in the synthesizer output can mix with adjacent channels, thus producing 'noise' that lies in the downconverted band of the desired signal. This is illustrated in Fig. 3.1. The spur that is $\Delta\omega$ away from the LO signal, mixes with the blocking signal that is also $\Delta\omega$ away from the desired channel. The spur must be low enough so that the noise that aliases on top of the channel is still below the level determined by the eventual BER. Typical requirements on the spur level would be in the range of -50 to -60dBc.

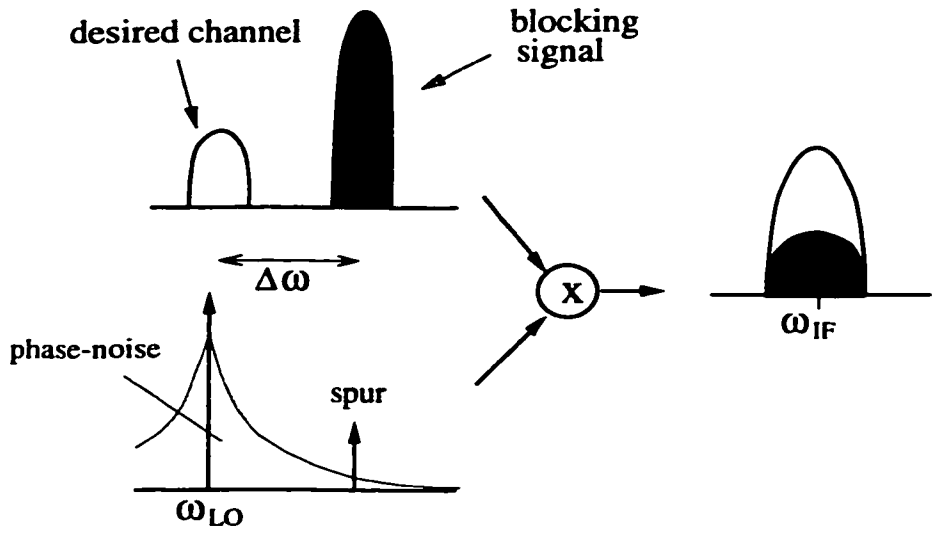

Figure 3.1: Effect of spurs on reception.

3.1.4 Phase-noise

Not only does the frequency synthesizer output exhibit spurious tones, but also, the fundamental LO frequency is not a discrete tone in the frequency domain, but rather a tone with phase-noise 'skirts' on both sides as shown in Fig. 3.1. In the time domain, this phase-noise leads to random jitters in the zero-crossing of the signal.

The effect of phase-noise is identical to the effect of spurious tones. During transmission, the transmitted signal will exhibit phase-noise. The *total* phase-noise in the band of a nearby receiver should be within the regulatory requirements. During reception, the phase-noise of the LO will again mix with adjacent channels and cause the same aliasing effect as in Fig. 3.1. This is called reciprocal mixing. With adjacent channels as large as 60-80dB above the desired signal, the phase-noise requirement is one of the very stringent and difficult requirements on the

frequency synthesizer. Phase-noise is quantified as the total noise-power in a BW of 1Hz at a specified offset from the carrier.

3.1.5 Speed

The final consideration in the design of frequency synthesizers is the speed. Communication standards will impose a limit on the maximum time needed to switch from one channel to the other. The frequency synthesizer needs to settle to within the required accuracy in that time. In GSM, the settling time should be less than 0.5ms [33]. This imposes a constraint on the BW of the frequency synthesizer that is often conflicting with other requirements.

3.2 The Voltage-Controlled Oscillator

Before we explain the operation of the PLL, we briefly discuss the structure of the VCO and its phase-noise. The VCO is the most critical component in the frequency synthesizer, since it is the device that actually produces the oscillation. It consists of an oscillator whose oscillating frequency can be controlled by a voltage (or indeed a current). The remainder of the PLL is only to lock this oscillator to a clean reference signal.

There are two distinct types of oscillators: the ring oscillator and the LC-based oscillator. The former is basically a chain of inverters connected in a ring with a net phase-shift of 180° around the loop. The frequency of oscillation depends on the number of inverters used and the delay per stage. Since the latter depends on the bias current, ring oscillators are usually tuned by varying the bias conditions.

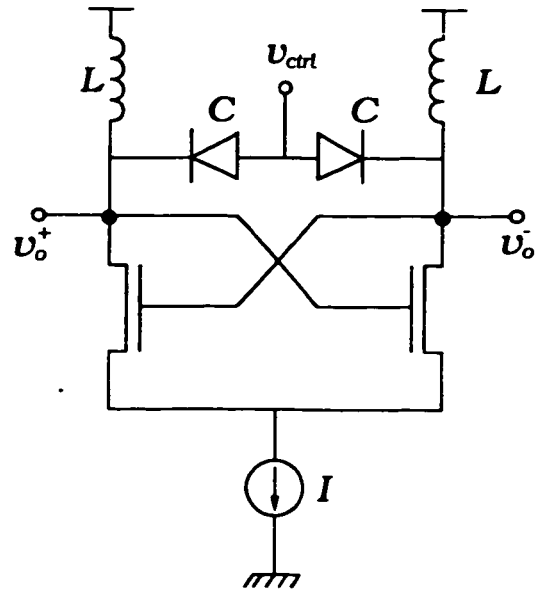

Figure 3.2: A shunt-LC tank VCO.

Ring oscillators exhibit very wide tuning ranges which is beneficial in covering the temperature and process variations. However, they exhibit a high level of phase-noise. It has been shown that ring oscillators have an equivalent open-loop Q of approximately one [34], much less than their LC counter parts. Thus ring oscillators are not popular in RF frequency synthesizer design and are rarely used to implement the LO VCO.

3.2.1 Operation of LC-VCO

LC VCOs consist of a inductor-capacitor resonance circuit with an active element providing the negative resistance. If the capacitance used is a varactor, then the control voltage will determine the tank resonance frequency. Fig. 3.2 is a popular example of a differential LC-VCO. The shunt inductor, L, and varactor, C form a

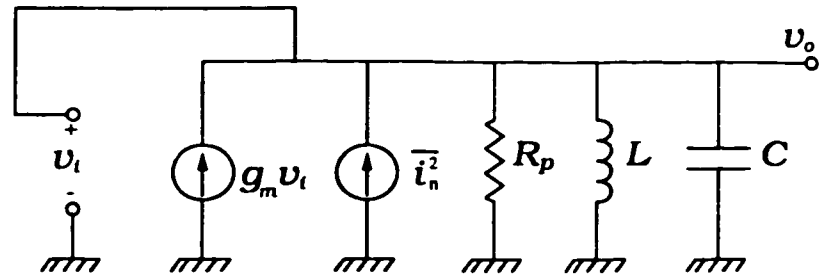

Figure 3.3: Equivalent circuit of LC VCO.

parallel resonance tank circuit. The cross-coupled transistors provide the negative resistance necessary to sustain the oscillations in the presence of the losses in the tank circuit.

3.2.2 Phase-noise

Fig. 3.3 shows the equivalent circuit of the LC-VCO. The shunt resistance, R_P, represents the equivalent shunt resistance of the inductor and is given by

$$R_P = \frac{(\omega_o L)^2}{r} = Q\omega_o L \tag{3.1}$$

where ω_o is the tank resonance frequency given by $1/\sqrt{LC}$ and r is the series resistance of the inductor. The tank quality factor, Q, is equal to $\omega_o L/r$. The current source $g_m v_i$ represents the gain of the active element and has an upward direction to signify the negative resistance. The noise current source $\overline{i_n^2}$ represents the noise of the transistor and tank resistance.

By expressing the tank impedance in terms of ω_o and Q, the output voltage is

given by

$$v_o = (g_m v_i + i_n) Z_{tank} = (g_m v_i + i_n) R_P \frac{\frac{1}{Q}\frac{s}{\omega_o}}{(\frac{s}{\omega_o})^2 + \frac{1}{Q}\frac{s}{\omega_o} + 1} \tag{3.2}$$

Since the output is directly fed-back to the input, then $v_o = v_i$ and

$$v_o = i_n R_P \frac{\frac{1}{Q}\frac{s}{\omega_o}}{(\frac{s}{\omega_o})^2 + (1 - g_m R_P)\frac{1}{Q}\frac{s}{\omega_o} + 1} \tag{3.3}$$

It is clear from (3.3) that when $g_m R_P = 1$, the middle term in the denominator vanishes and the gain of the VCO will be infinite at ω_o. Thus oscillations will build-up from any noise in the circuit. Under these conditions, the noise magnitude at the output will be given by

$$\overline{v_{on}^2} = \overline{i_n^2} R_P^2 \frac{\frac{1}{Q^2}(\frac{\omega}{\omega_o})^2}{(1 - (\frac{\omega}{\omega_o})^2)^2} \tag{3.4}$$

At an offset of $\Delta\omega$ from the carrier, the frequency will be given by $\omega_o + \Delta\omega$. Close to the carrier, the ratio $(\omega/\omega_o)^2$ is approximately 1. However, the denominator can be approximated by taking the first two terms of the Taylor expansion of that ratio. Thus the output noise will be given by

$$\overline{v_{on}^2} = \overline{i_n^2} R_P^2 (\frac{\omega_o}{2Q\Delta\omega})^2 \tag{3.5}$$

The noise $\overline{i_n^2}$ represents the noise of both the transistor and the shunt resistance

and is equal to

$$\overline{i_n^2} = 4kT(\gamma g_m + \frac{1}{R_P}) \ \Delta f \tag{3.6}$$

where γ is a bias-dependent constant [9]. Since this thermal noise causes fluctuations in both the amplitude and phase, the phase perturbation will be given by half the quantity in (3.6) [35]. Thus the output noise will be

$$\overline{v_{on}^2} = 2kT(1 + \gamma g_m R_P)R_P(\frac{\omega_o}{2Q\Delta\omega})^2 \ \Delta f \tag{3.7}$$

The phase-noise will be given by the output noise referred to the tank signal r.m.s. value, or

$$\Phi_n(\Delta\omega) = 10log(\frac{2kT(1 + \gamma g_m R_P)}{(v_{orms}^2/R_P)}(\frac{\omega_o}{2Q\Delta\omega})^2) \quad dBc/Hz \tag{3.8}$$

The first term in the logarithm argument is the ratio of the noise power to the signal power. The second term is the ratio of the VCO BW to the offset frequency.

Several key observations follow from (3.8). First, the phase-noise depends on the square of the operating frequency. Thus if a VCO is operated at half the frequency, the phase-noise decreases by 6dB. Second, the phase-noise decreases with the square of the offset frequency. The oscillator phase-noise will typically have a slope of -20dB/decade around the carrier. This constitutes the 'skirts' of the spectrum. A consequence of that is that meeting phase-noise requirements in standards with narrow channel BW is more difficult.

It is also clear from (3.8) that in order to reduce the phase-noise, the signal

power must be maximized. Since the tank peak-to-peak voltage is roughly $I_{bias}R_P$, this involves a trade-off with power consumption. Thus reducing the phase-noise for a particular technology would be at the expense of power consumption. In addition, using a smaller inductor is beneficial since R_P will be reduced. However, this again will be at the expense of power dissipation.

The phase-noise is also dependent on the square of the tank quality factor. It is thus very important to maximize Q. The value of Q in typical RF technologies range from 5-8 in the 1GHz range, and decreases beyond the 2GHz range due to substrate losses.

The term g_mR_P represents the open-loop gain of the VCO. This value is usually 2 to 3 in order to ensure that oscillations will start-up and sustain over all process and temperature variations. The value of the numerator of the first term in (3.8) will be approximately $14kT$.

3.2.3 Tuning

One of the parameters of the VCO is its tuning range. The VCO must be able to cover the full RF band over all temperature and process variations. The total capacitance of the tank circuit is made up of the varactor in addition to the inductor and transistor capacitances. Thus in order to increase the tuning range, the value of the varactor must be increased so that the variable part of the capacitance would dominate. This however involves a trade-off with power consumption, since as the value of the capacitance increases, the inductor must be decreased. For the same phase-noise, this will lead to a higher bias current as mentioned above.

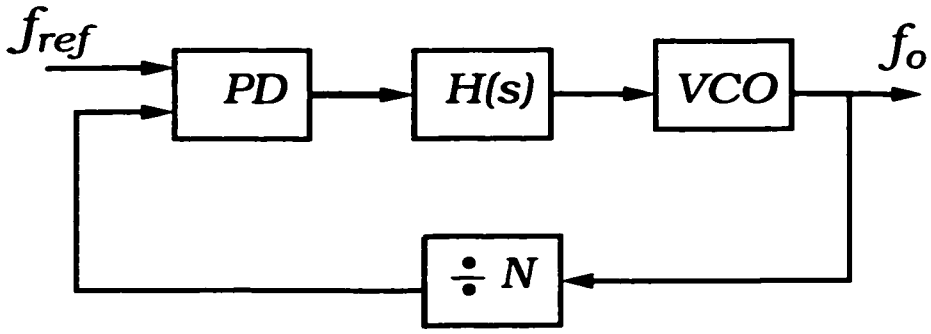

Figure 3.4: A classical PLL structure.

The tuning range of the VCO must be large enough to cover all variations. However, any further increase is not desirable as it only leads to a greater amplification of the noise at the VCO input.

3.3 The Classical PLL

The phase-locked loop is predominantly the most widely used means of frequency synthesis. A block diagram of a classical PLL is shown in Fig. 3.4. The output of the VCO is fed back to a frequency divider N, where N is an integer. The output of the divider is compared to a reference signal, f_{ref}, that is derived from a crystal oscillator. The comparison is done by a phase-detector, PD, that produces a signal proportional to the phase difference between the two inputs. The error signal is filtered by a filter, $H(s)$, and controls the VCO. If $H(s)$ includes an integrator, then the error signal will be zero at steady-state and the two inputs of the phase-detector will have equal phases (and hence frequencies). Thus at steady-state the output frequency will be equal to Nf_{ref}.

Since the division ratio, N, is an integer then the reference frequency, f_{ref} must

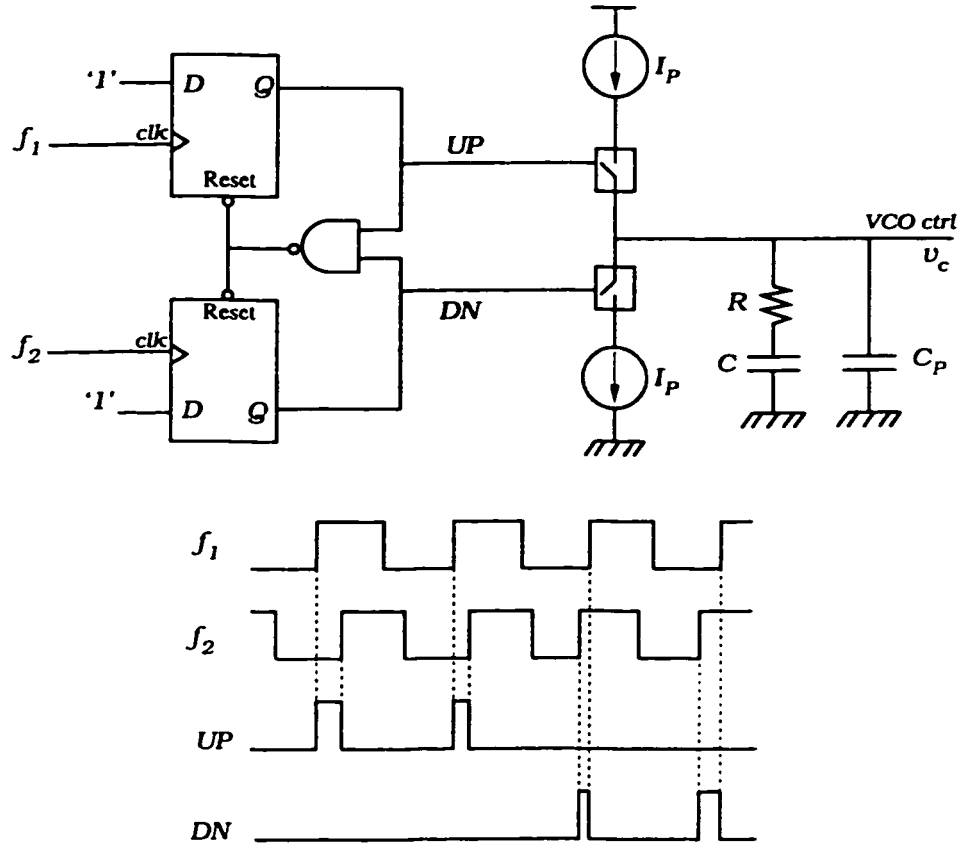

Figure 3.5: The phase-frequency detector and charge pump.

be equal to the channel width such that the output resolution will equal the channel spacing. As explained below, this is a severe limitation of the PLL.

3.3.1 The Phase-Detector and Charge-Pump

The phase-detector and charge-pump are responsible for producing an output current that is proportional to the phase difference between the two inputs. Fig. 3.5 shows a widely used phase-frequency detector (PFD) consisting of two D-type flip-flops [36]. The current sources and the passive components constitute the charge-

pump and loop filter. The timing diagram of the PFD is also shown. The PFD activates each cycle either the UP or DN signal, depending on which input is leading. The duration for which either signal is activated is equal to the phase difference. The amount of charge pumped into (or depleted from) the dominant capacitor C is thus proportional to the phase-difference.

It is clear that the average current flowing to the loop filter is equal to $I_P\theta_e/2\pi$, where θ_e is the phase difference. Assuming C_P is negligible with respect to C, the transfer function of the PFD/CP and loop filter will be given by

$$\frac{v_c}{\theta_e} = \frac{I_P}{2\pi} \cdot \frac{1 + sRC}{sC} \tag{3.9}$$

It is clear that the charge-pump, along with capacitor C, form an integrator. Since the VCO transfer function also involves integration (K_o/s), the resistor R has to be added to create a zero in the loop filter that provides phase-margin and thus stability.

The capacitor C_P is inserted to mitigate the ripples that would otherwise be present on the VCO control due to the square-wave nature of the charge-pump output. In order to ensure that the stability of the loop is maintained, this additional pole has to be far from the dominant one. Thus C_P is typically equal to $C/10$.

During lock, the two inputs will be in-phase and the output current would ideally be zero. However due to mismatches in the current sources, residual charges are pumped to the capacitor creating the reference spur. A critical problem with the PFD is also the fact that due to the delays in the PFD, there is a minimum phase difference below which the PFD produces no response (deadzone). The deadzone

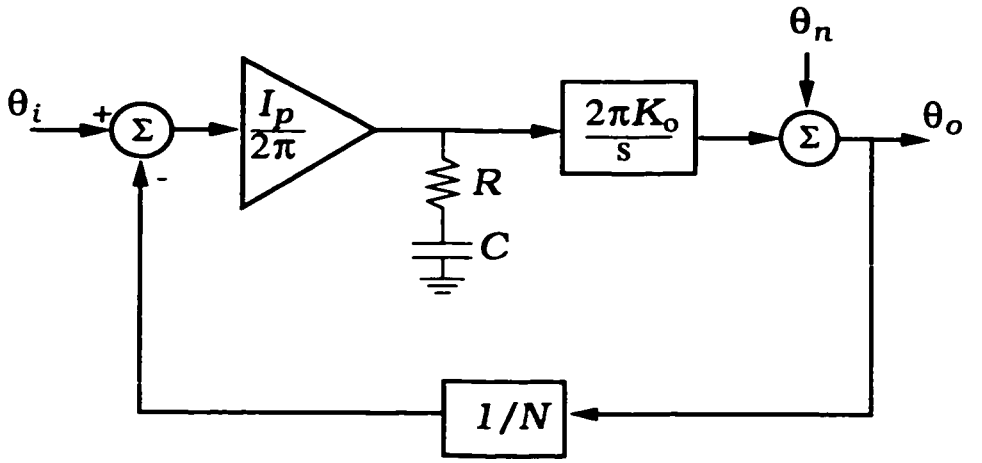

Figure 3.6: S-domain block diagram of a classical PLL.

is a range of uncertainty where the gain of the phase-detector is essentially zero. The deadzone is a critical source of spurs and phase-noise.

A deadzone-free phase-detector is a simple XOR gate. When the two inputs of the XOR gate are 90° out of phase, the XOR output is a square wave with a 50% duty cycle and hence zero average. The drawback of the XOR gate, beside higher spurs, is the limited pull-in range (section 4.5.3) [36].

3.3.2 Transfer Functions

An s-domain block diagram of the classical PLL is shown in Fig. 3.6. The amplifier $I_P/2\pi$ represents the current of the charge pump, and the RC combination constitute the loop filter. The VCO gain is K_o Hz/V. The input phase is θ_i, while θ_n is the phase-noise of the VCO. The transfer function between the output and the two

inputs will thus be

$$\theta_o = N\theta_i \cdot \frac{2\zeta(\frac{s}{\omega_n}) + 1}{(\frac{s}{\omega_n})^2 + 2\zeta(\frac{s}{\omega_n}) + 1} + \theta_n \cdot \frac{(\frac{s}{\omega_n})^2}{(\frac{s}{\omega_n})^2 + 2\zeta(\frac{s}{\omega_n}) + 1} \tag{3.10}$$

where

$$\omega_n = \sqrt{\frac{K_o}{N} \cdot \frac{I_P}{C}}, \quad \zeta = \frac{1}{2}RC\omega_n \tag{3.11}$$

Clearly, the system is a second-order system, with a natural frequency of ω_n and a damping factor ζ. The value of ζ is chosen so as to have enough phase-margin to guarantee stability, and to limit the overshoots during the transient response. The value of ω_n is then determined by the required loop bandwidth.

Several key observations are notable here. First, the input phase, as expected, is multiplied by N. This is very significant since it means that the power of any spur or input noise is multiplied by N^2. Second, the transfer function between the output and input is that of a low-pass response. Qualitatively, the loop cannot track those variations at the input that are faster than the loop BW. The slope of the θ_o/θ_i transfer function is -20dB/decade beyond the loop bandwidth.

The transfer function from the VCO phase-noise to the output is that of a high-pass response. For offsets beyond the loop BW, the transfer function is equal to unity. Below that, the function slopes downwards with a slope of 40dB/decade. Qualitatively, as long as the jitters (phase-noise) in the VCO are slower than the loop BW, then the loop will be able to track and correct them, thus cancelling their effect. For offsets greater than the BW, the loop is not fast enough to respond to

them and the phase-noise is passed to the output uncorrected.

3.3.3 Spurious Content

In the analytical models, the output of the phase-detector is equal to zero at steady state. However, the practical case is different. For example, if an XOR gate is used as a phase-detector, then the output of the charge-pump is a square wave whose *average* is zero but has a fundamental frequency and harmonics. This waveform will cause discrete sidebands (spurs) at the output. In the more popular case of using a phase-frequency detector, mismatches in the UP and DOWN current, leakage in the transistors and deadzones in the detector all cause spikes on the VCO control creating a spur at the reference frequency [37].

The spur problem is aggravated by the fact that input/output transfer function contains the factor N as shown in (3.10). Hence, if the input-referred spurs are at a certain level with respect to the carrier, then this level is increased by N^2 or $20log(N)$ at the output. Typical values of N in a classical PLL range from 1000 to above 30,000. The value $20log(N)$ can thus be very large. It is of great benefit to reduce the division ratio if possible [37].

3.3.4 Output Phase-noise

There are two main contributors to the output phase-noise: the phase-noise in the VCO itself, and the phase-noise at the input. As explained in section 3.2.2, the phase-noise of the VCO has a slope of -20dB/decade near the carrier. When configured in a PLL, the phase-noise of the VCO is attenuated within the loop BW.

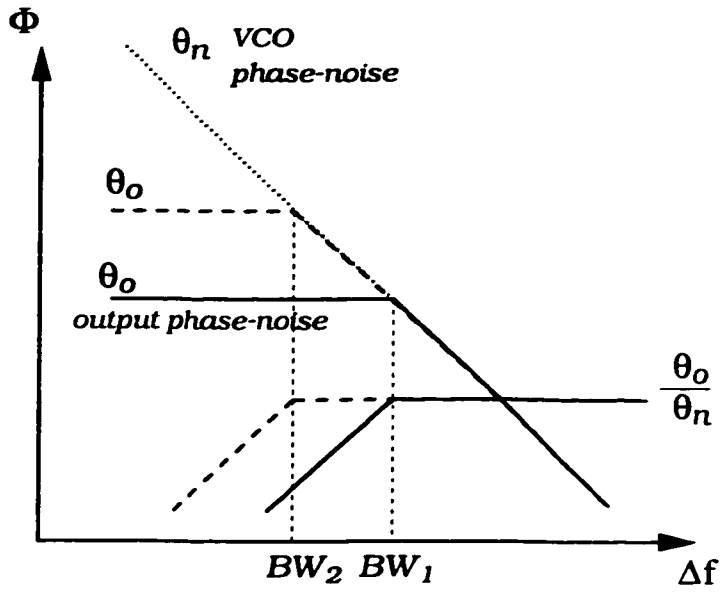

Figure 3.7: Effect of PLL on VCO phase-noise.

Hence, the phase-noise due to the VCO will typically take a LPF shape as shown in Fig. 3.7. At offsets larger than the BW, the PLL phase-noise tracks the VCO noise and slopes down 20dB/decade. Within the BW however, the positive slope of the high-pass transfer function (θ_o/θ_n) will cancel the negative slope of the VCO yielding a characteristic low-pass shape. Clearly, as the BW is increased, the VCO phase-noise is brought down. Thus using a wide-BW architecture is very efficient in reducing the VCO noise.

The other source of noise is the input noise which includes the phase-noise of the input reference signal, plus any input-referred noise from the charge-pump and frequency divider. This noise is again multiplied by N^2 when referred to the output. With very large values of N, this noise usually dominates the noise of the PLL within the BW. Unlike, the VCO, the phase-noise at the input favors a narrower BW in order to reject the noise at closer offsets.

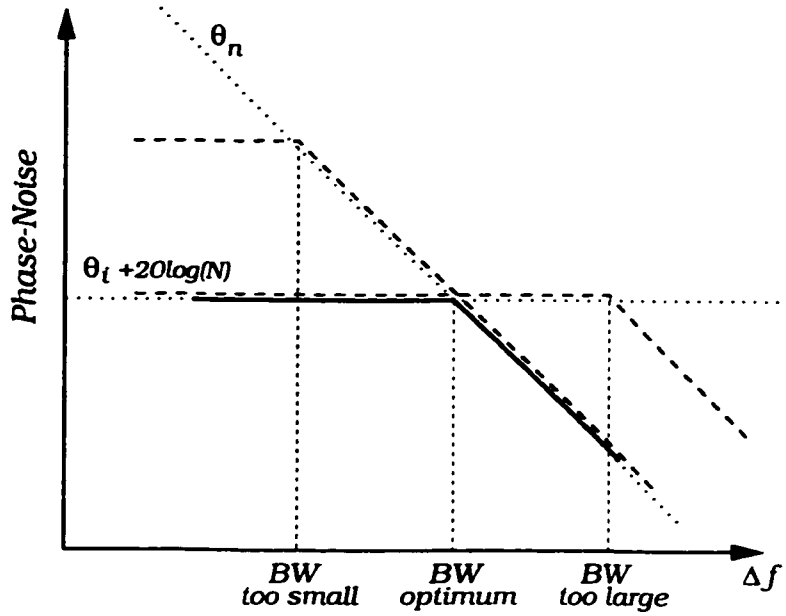

Figure 3.8: Total phase-noise at the PLL output.

3.3.5 The choice of BW

The choice of BW affects many parameters. From a phase-noise point-of-view, the optimum BW is the intersection of the VCO phase-noise and the input phase-noise (after being multiplied by N^2). This is depicted in Fig. 3.8. A BW that is less that BW_{OPT} will lead to a higher in-band phase-noise level, while a higher BW will only allow more input noise to pass through. In an optimum design, the noise level within the BW is dominated by the input noise, while beyond the BW, the VCO noise dominates.

Unfortunately, this value of BW is often not possible because of the other two fundamental concerns: spurs and speed. A BW that is larger than BW_{OPT} might have to be used in order to meet the settling time requirements, whereas a smaller BW in other cases must be used to keep the spur level within the specifications.

In the case where both the spur and the speed requirements cannot be met simultaneously, a small BW can be chosen to attenuate the spur and then speed-up methods to accelerate the loop are employed. Alternatively, the BW can be increased, while adding more poles in the loop filter to increase the filtering order. Both of these methods suffer from several problems.

From the above discussion we see that the classical PLL suffers from several problems. The low frequency f_{ref} signal leads to a very large division ratio that greatly increases the spur levels and the input phase-noise contribution. Due to this, the BW is often very narrow, leading to a slow loop that sometimes needs speed-up methods. The small BW means also that the VCO noise will not be attenuated in the range of interest, making an off-chip high-quality VCO often necessary.

3.4 Alternative Architectures

3.4.1 Fractional-N PLL

The fundamental problem with the classical PLL is that the division ratio is an integer. The fractional-N PLL tries to periodically alternate the division ratio between N and $N+1$ such that the *average* division ratio is a fraction in between. This is accomplished as shown in Fig. 3.9.

Suppose that the accumulator is a n-bit accumulator and increments its content by W every clock. Hence for every 2^n cycles of the reference frequency, the carry-out signal is high for W cycles. The average division ratio will thus be $N + \frac{W}{2^n}$. For the same resolution, the reference frequency can be increased by a factor of 2^n and

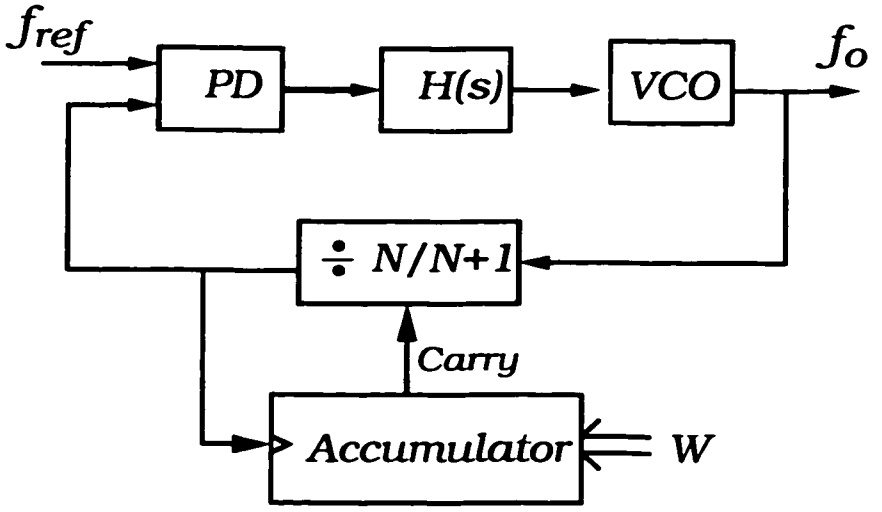

Figure 3.9: The fractional-N PLL.

the division ratio is correspondingly reduced.

The performance of fractional-N PLLs tends to deteriorate with increased fractionality (the value of 2^n). Commercial fractional-N PLLs usually have a fractionality of 8 (n=3) [38], 16 [39], and are limited to 32 [40]. This puts a limit on the achievable reduction in division ratio and in many cases is not enough. The main problem with fractional-N PLLs however is the fact that a spurious signal still persists at an offset equal to the channel spacing (rather than the increased reference frequency). This is because, the division ratio is changed from N to $N+1$ at a rate equal to $Wf_{ref}/2^n$, which is equal to the channel width when $W = 1$. Even with spur compensation methods, the loop BW remains an order of magnitude below the channel spacing [38, 40] and no improvement on the VCO noise is achieved at the range of interest.

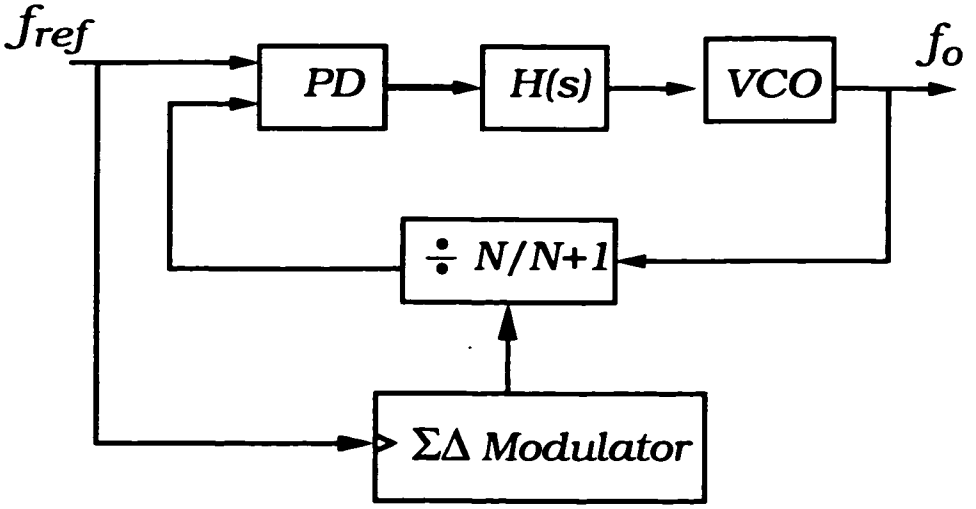

Figure 3.10: The ΣΔ PLL.

3.4.2 ΣΔ PLL

An interesting modification of the fractional-N PLL is the ΣΔ PLL [41]. The idea is that we want to represent a fractional number in terms of one of two values: N and $N + 1$. This is analogous to a 1-bit ΣΔ modulator that converts an analog signal to a 1-bit digital output. This is depicted in Fig. 3.10. Of course, in order to achieve the required resolution, the clock rate of the modulator must be very high.

The result of the ΣΔ PLL is that the instantaneous division ratio is varying very quickly producing a 'random' sequence whose average is the required fraction. Consequently, the discrete side-bands in the fractional-N PLL are converted to high-frequency random noise.

The ΣΔ architecture achieves a large reduction in the division ratio and uses reference frequencies to the PLL in the range of 10's of MHz. This results in a significant reduction in the phase-detector noise and spurs as well as a potential for

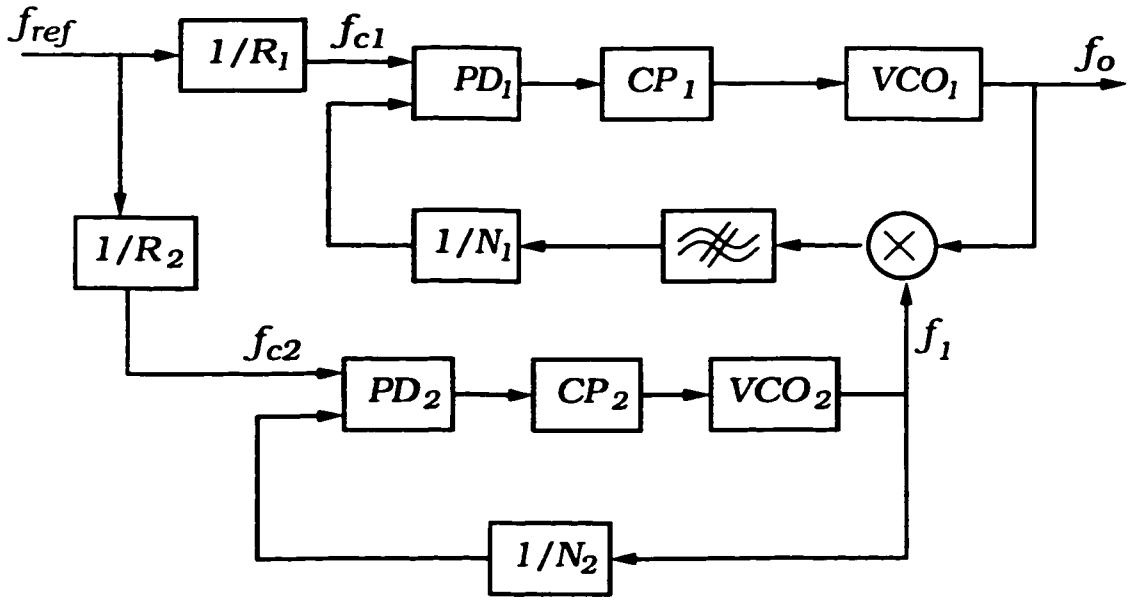

Figure 3.11: A double-loop PLL architecture.

a wide bandwidth. The fine resolution in division ratio (fractionality) that can be achieved can be utilized in directly modulating the carrier with data.

The problem with the ΣΔ PLL is the very high noise at larger offsets. To properly attenuate this noise, the clock to the ΣΔ modulator must be very high, especially for wider BW, which results in large power consumption.

3.4.3 Dual-loop PLLs

Another way to relax the requirements on the PLL is through the use of two loops [42]. Essentially the degrees of freedom are being increased hence easing the design. Of course the fundamental problem is that we need extra hardware and power consumption.

One such method is shown in Fig. 3.11. To increase the output frequency

by one step, N_1 is incremented by one while N_2 is decremented by one. Thus the frequency resolution of the output works out to be the *difference* between the two comparison frequencies f_{c1} and f_{c2}. The result is two coarse-tuned PLLs that provide fine tuning.

The benefit here is that only the difference $f_{c1} - f_{c2}$ needs to be equal to the channel spacing, while each one of the two frequencies can be much higher. This results in reduced division ratios and larger BW.

This approach does suffer from several problems, however. First, the hardware is essentially doubled and so will be the power consumption. Two high-frequency VCOs will be needed plus a high-frequency mixer. Also, the phase-noise of VCO_2 within the BW will add to the output. Second, since they are derived from the same system crystal, the values of f_{c1} and f_{c2} are limited to less than about 5MHz for typical values of crystal oscillators. Thus the reduction in division ratio and increase in BW is limited. Third, the mixing operation creates several spurs that, due to the non-linear nature of division, can be aliased within the loop BW [42]. In addition, the tuning range required for VCO_2 is very large.

3.4.4 Direct Digital Frequency Synthesis

The final method of frequency synthesis that we explain is a non-PLL-based, digital method. The idea is that a sine wave is a defined waveform that can be synthesized digitally. The easiest way is that shown in Fig. 3.12. An n-bit accumulator is incremented by a value W every clock cycle, thus producing a staircase waveform that represents the phase of the sinusoidal signal. A look-up table stored in the

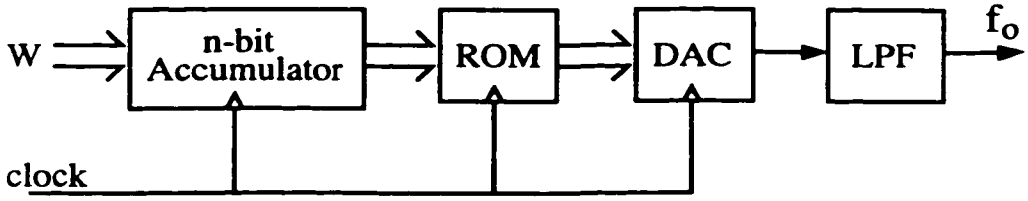

Figure 3.12: The DDFS structure.

ROM produces the proper amplitude that corresponds to the input phase, and a DAC converts that to the analog domain. The output frequency will be equal to $Wf_{clk}/2^n$, and is tunable by changing W.

To satisfy Nyquist criterion, the maximum frequency that can be synthesized is equal to $f_{CLK}/2$. However, this would require a 'brick-wall' low-pass filtering at the output. Thus the clock rate is typically 3 times the maximum desired frequency.

The DDFS is a very efficient way of frequency synthesis. It is very accurate and is an open-loop architecture with no feedback. Hence frequency transitions occur almost instantaneously. By making the accumulator width large, a very fine resolution can be achieved (e.g. 0.1Hz). This allows the DDFS to be used in modulation.

The clock of the DDFS is derived from the system crystal. Since the DDFS output is a fraction of the clock, the already-low phase-noise of the crystal is further reduced. The DDFS phase-noise performance is limited only by the jitter caused by the digital circuitry and the DAC and generally has excellent phase-noise performance [43].

The DDFS is an attractive mean of frequency synthesis at low frequencies. Operation at high frequencies however lead to excessive power dissipation and great difficulty in designing the DAC with adequate accuracy. The DDFS in [44] consumes

1.4W while operating at the maximum clock of 350MHz. The one in [45] consumes 0.6W while operating at the maximum clock of 150MHz.

The second limitation of the DDFS is the high spurious content. The DDFS exhibits large spurs at the output, primarily due to the DAC, that can be very close to the fundamental. Operation at higher frequencies aggravates the problem. Meeting the spur requirement imposes difficult constraints on the design of the DAC and/or the operation frequency.

One application of the DDFS is to use it to synthesize a low frequency, and then upconvert that frequency to RF using a PLL. While this may have good potential, the limiting problem here is that the spurs will be increased by the PLL division ratio, making the design of the DDFS even more difficult. We attempt to solve this problem in chapter 6.

Chapter 4

Nested-Loop PLL Architecture

A major goal of the wireless industry today is to reduce the number of components in the radio and integrate as many functions as possible inside the chip. Towards achieving this goal, the RF VCO is a candidate with a very high priority. Its integration not only reduces the cost and size, but also power consumption and interference problems. In this chapter we argue that a wide-bandwidth PLL (WBPLL) architecture is one of the most promising ways to integrate the VCO in today's technology. A simple wide-BW PLL architecture is proposed. A simple system-level modification to the transceiver architecture employs this WBPLL with no extra H/W or power dissipation. The advantages and considerations of this architecture are discussed. The implementation of the WBPLL along with measurement results are presented.

4.1 The Wireless Transceiver Frequency Synthesizer Problem

As discussed in section 3.1.4, a major parameter of the frequency synthesizer is its phase-noise. Phase-noise of the frequency synthesizer contributes to the receiver noise figure [46] and leads to reciprocal mixing with blocking signals that could desensitize the receiver [33,46]. It is thus essential to ensure a low phase-noise level at the output of the frequency synthesizer. Unfortunately, the phase-noise performance of integrated VCOs is relatively poor and does not meet many communications standards specifications over the full temperature and process variations.

In order to meet the specs, there are two alternatives for the designer: 1) resorting to external VCOs, generally having superior phase-noise performance, or 2) using an integrated VCO while attenuating its phase-noise through a wide PLL bandwidth [17, 33, 47]. The first alternative is much more popular [2, 3, 5]. This is because most PLL architectures cannot increase the PLL BW significantly while meeting the spurious requirements. Using an external VCO however does have many problems associated with it making it an undesirable alternative albeit necessary in most cases. In the following paragraphs, we highlight some of those problems.

Among the obvious and very significant problems of the external VCO is the extra ***cost*** and ***size***. In addition to the necessary decoupling capacitors, the external VCO significantly increases the component count and board area. It also increases the radio cost significantly. The savings in cost and size are two major reasons for

finding a way to integrate the VCO. They are not the only two reasons however.

The RF VCO is one of the sensitive components of the transceiver that is easily both affected by and affects other components of the system. It is thus of significant benefit to shield the VCO from other critical components. Once the VCO is taken off-chip the problem of ***interference*** and ***shielding*** is considerably exacerbated. An off-chip VCO is much more susceptible to interference from the off-chip power amplifier. This problem is even more critical in systems that transmit a large power level (e.g. Mobitex, 2W). The powerful PA signal couples to the VCO producing undesirable spurs and contaminating the spectrum [48]. The problem is amplified if the VCO frequency is close to that of the PA, as the PA will start 'pulling' the VCO frequency [49], [50]. This problem makes direct-conversion transmitter architectures very difficult to implement [23], since this interference produces many large spurs very close to the carrier causing 'spectral regrowth' [48, 51]. It is thus necessary to implement heterodyne transmitter architectures which need extra hardware and power consumption.

Another problem that follows from the use of external VCOs and inadequate shielding is the ***LO leakage*** effect. The VCO signal present on the board can easily leak to the antenna causing carrier feedthrough. Since the VCO signal is relatively strong, the amount of leakage is often above the regulatory specifications. This problem is even more profound in direct-conversion receivers where the LO frequency lies inside the receive band and is subject to in-band EMI requirements [31,49]. Also in this receiver architecture, the VCO signal can be reflected back from the antenna and self-mix in the downconversion mixer, thus creating a time-varying

d.c. offset that contaminates the desired signal and is difficult to remove [24, 28]. Contributing to the problems of interference and leakage is the ***lack of differential*** external VCOs, as opposed to integrated VCOs that are easily made differential [48].

Another consequence of the use of an external VCO is increased ***power consumption***. The off-chip VCO needs to drive a 50Ω load and the large board capacitances. This leads to a significant power dissipation in the buffer. It is typical for an external VCO to dissipate a whole 30mW [52]. Considering that, among the radio components, the VCO stays ON for the greatest percentage of time [3], this level of power consumption can have an adverse effect on the total battery life. Any savings in VCO power consumption is very significant.

Related to the issue of power consumption is the need for a ***high-gain preamplifier*** at the input of the prescaler to detect the usually weak signal from external VCOs. Since this amplifier operates at high frequencies, it is usually a significant power drain in the PLL. In addition, its very high gain could cause instabilities in the circuit.

It should be mentioned that using an on-chip VCO with an ***external tank*** circuit can alleviate some of the above problems. However, this approach suffers from significant drawbacks. The large board parasitics add directly to the tank circuit and become very significant especially at higher frequencies (e.g 1.8GHz). This limits the available tuning range and could even limit the maximum operating frequency. The parasitics complicate the design and make it dependent on the board parameters and the specific board layout. Also, the larger parasitic capacitors being driven in the GHz range mean inherently larger power dissipation. In addition

bondwire parasitics affect the tank circuit and can cause undesirable modes of oscillation [53]. Also, the differentiality of the signal can be degraded because the additional parasitics are not well controlled. Finally, from a reliability point of view, using external tanks means needing extra pins that are not ESD protected. In addition to these problems, the aforementioned problem of interference and leakage persists [28] as well as the extra cost of the tank circuits.

All of the above problems make the integration of the VCO very attractive. On-chip VCOs, by themselves, cannot meet the specifications in many cases. The theoretical value of the current needed is either very large or not possible [54]. Thus an efficient way to integrate the VCO is through the help of a wide-bandwidth PLL.

4.2 Wide-Bandwidth PLL Architectures

Communication systems have very stringent phase-noise requirements. These requirements gets even more stringent with reduced channel spacing. In Mobitex, a 2-way wireless data network [32], the channel spacing is 12.5kHz and the phase-noise of the frequency synthesizer needs to be more than 95dB down at a 10kHz offset (Mobitex operates in the 900MHz range). These very low levels of phase-noise cannot be achieved with today's standard-technology LC oscillators [55–57], or even oscillators employing bondwire inductances [58, 59]. This is also true for many other communication standards.

Integrated VCOs that meet such stringent phase-noise requirements will be possible only through innovative frequency synthesizer architectures that employ wide-BW PLLs [17], [33]. As seen from equation (3.10), the VCO phase-noise is

attenuated by 40dB/decade within the loop BW. Thus the use of a WBPLL is an efficient and promising way of integrating the VCO. This has to be done however, in a way that does not sacrifice the spurious performance of the frequency synthesizer.

In addition to the benefits achieved by integrating the VCO, a WBPLL has a shorter locking and settling time. This is a significant benefit as the settling time requirement is usually quite stringent (e.g. for GSM, the LO must settle to within 0.1ppm in 0.5ms [31]). This benefit is even more significant in frequency-hopping CDMA systems where channel frequency periodically and rapidly changes.

Another important advantage of a WBPLL is the low division ratio it usually utilizes. This leads to reduced output phase-noise due to the phase-detector and charge-pump as well as reduced spurii from the charge-pump (section 3.3.2).

Care must be exercised in the design of WBPLLs, however. The wide BW of the PLL will allow any noise at the input to pass to the output unattenuated. It is thus important to make sure that this noise level will be sufficiently low. In addition the ratio of the BW to the reference frequency must remain low enough to ensure adequate spurious rejection. (section 3.3.5). In this section we present some WBPLL architectures and discuss their merits and problems.

As discussed in section 3.4.1, fractional-N PLL still exhibit spurs at the channel spacing, necessitating a narrow BW. They are thus not considered wide-BW architectures unless $\Sigma\Delta$ modulation is employed (section 3.4.2).

The $\Sigma\Delta$ architecture suffers from some drawbacks, however. Not only is the power consumption relatively high (since the digital $\Sigma\Delta$ modulator operates at a high frequency), but also the high-frequency noise constitutes a major problem.

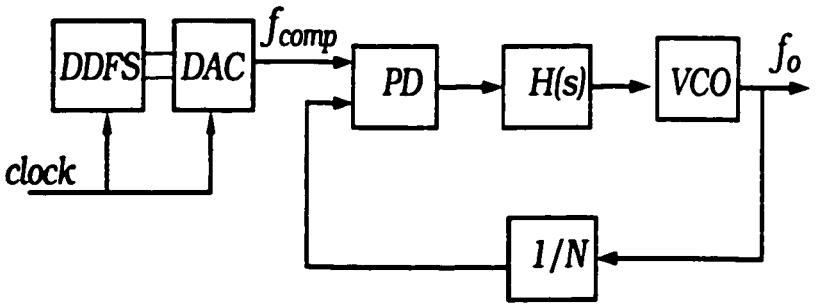

Figure 4.1: DDFS-driven PLL

This high-frequency noise is to be filtered out by the PLL transfer function. The PLL response rolls off at about 20dB/decade. For this to adequately reject the $\Sigma\Delta$ noise, the ratio of the clock frequency to the loop bandwidth (essentially the oversampling ratio), has to be very large. Thus an increase in BW means a need to generate a higher clock frequency (may well be limited by the available crystal reference) and involves a trade-off with power consumption [52]. Due to this fact, the BW of $\Sigma\Delta$ PLLs is limited to values that are often below those required for the VCO integration. The $\Sigma\Delta$ PLLs in [52] and [60] have a PLL reference (also the $\Sigma\Delta$ clock) of 20MHz and BWs of less than 100kHz.

Another wide-BW architecture is the use of a DDFS to synthesize a low-frequency signal and upconvert it using a PLL as shown in Fig. 4.1. If the synthesizer output frequency is f_o, then the DDFS output is to be f_o/N. A clock of approximately $3f_o/N$ will be needed. This allows the synthesis of a high frequency while operating the DDFS at a lower frequency and hence lower power. In addition the input to the PLL is of a relatively high frequency allowing for a larger BW and hence the potential for VCO integration. The reduced division ratio again means less

PD noise and CP spurs. The architecture is also suitable for constant-amplitude modulation (e.g. GMSK, QPSK, etc) if phase modulation is employed within the DDFS [61]. This architecture is a promising architecture that has the potential to achieve a monolithic low-power frequency synthesizer if the problems associated with it, listed below, are rectified.

Since the clock frequency is equal to $3f_o/N$, it is clear that there is a tradeoff between power consumption and the division ratio. While a small division ratio is desirable, this comes at the expense of increased power consumption in the DDFS and DAC due to the higher clock rate. The design of the DAC at a high frequency is also very problematic.

A major concern in the DDFS-driven PLL is the spur performance. The spurs produced at the DAC output are increased by $20log(N)$ when mapped to the output. In order to meet the regulatory requirements on the spur level the DAC output must be extremely 'clean'. For example even with a division ratio as low as 50, the DAC spurs must be more than 94dBc down to have a spur level of less than 60dBc at the output. This is very difficult to achieve. Even with the relatively superior performance of the DDFS in [62], the spur level is -70dBc when operating at 55MHz. The spur problem of the DDFS is a significant hurdle that has to be overcome for this architecture to achieve its goal.

4.3 A Novel Nested-Loop PLL Architecture

In this section we propose a simple novel PLL architecture that uses a low-frequency VCO to drive a WBPLL, thus alleviating many of the problems particular to the RF

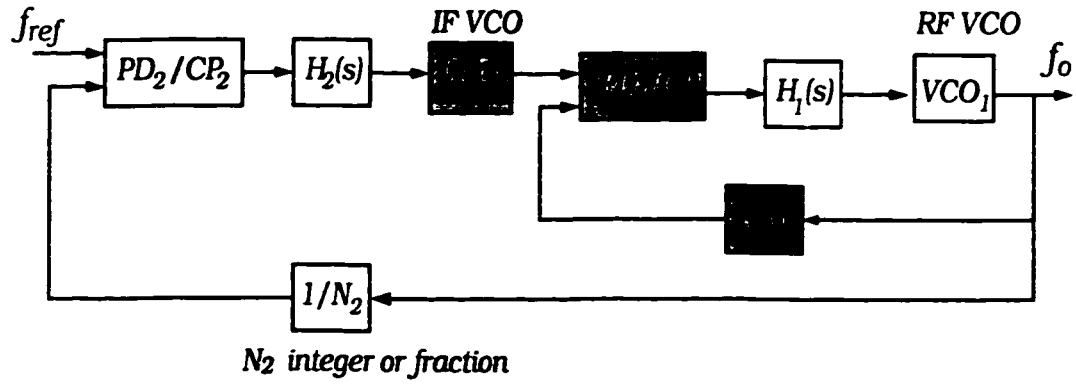

Figure 4.2: Nested-Loop PLL Architecture

VCO. The benefits of the architecture are presented and the stability considerations discussed.

The proposed PLL architecture is shown in Fig. 4.2. A PLL with a fixed division ratio, N_1, is configured within an outer phase-locked loop, thus creating a 'nested-loop' architecture. At steady-state, the output frequency, f_o, will be equal to $N_2 f_{ref}$ (like an ordinary PLL), and the output of VCO_2 will settle at f_o/N_1. With N_1 in the range of 10-20, VCO_2 will operate in the 100MHz range. With such a large input reference frequency, the inner PLL can have a very wide BW while maintaining adequate suppression of the spur at f_o/N_1. The BW of the overall loop is maintained one order of magnitude below f_{ref} to reject the spur at f_{ref} and the phase noise at the input of PD_2

As explained in the following section, the way the feedback is accomplished here ensures that, unlike cascaded PLL architectures, the settling time and input phase-noise rejection are not affected.

By adding an extra loop inside the PLL, one more degree of freedom is achieved

that allows us to increase the bandwidth of the inner loop to attenuate the RF VCO phase-noise, while keeping the spur level and input phase-noise low. This multiple-loop architecture achieves this benefit in a simple way with one additional low-frequency VCO, without the use of extra high-frequency VCOs, mixers or filters such as other double-loop architectures [3,33].

Note that the outer loop can still be a classical loop, fractional-N loop, $\Sigma\Delta$ loop or any other loop architecture desired.

4.3.1 Advantages

The main purpose of this architecture is to embed a wide-BW loop (with a sufficiently high reference) within a narrow-BW classical, fractional or $\Sigma\Delta$ loop that achieves fine frequency resolution while adequately rejecting the spur at the channel spacing. For example, if the LO frequency is in the 900MHz range and the frequency divider N_1 is fixed at 10, then the input to the inner loop is around 90MHz. With such a reference (and such a low division ratio, thus low impact of PD_1 phase-noise on output), the inner loop can easily have a bandwidth in the range of 1-5MHz. Such a wide bandwidth will greatly attenuate the phase-noise of the RF VCO (VCO_1), thus enabling its integration and achieving the associated benefits. With a differential integrated RF VCO, the cost and size of the radio will be significantly reduced as well as the interference with the power amplifier and the leakage to the antenna. The power needed to drive the off-chip 50Ω's and the board parasitics will be eliminated and so would the need for a high-gain power-hungry pre-amplifier. The IF VCO (VCO_2) operates at a low-frequency and can meet the

phase-noise specifications much more easily (section4.3.2).

The architecture of Fig. 4.2 is similar to a cascaded PLL architecture where a low-frequency PLL drives a subsequent one. In that case, the input of the divider N_2 is taken from the output of VCO_2. The problem with such an architecture is the much longer settling time and larger overall division ratio. This is because the frequency synthesizer output's resolution must be equal to the channel spacing (f_{CH}). If a cascaded PLL architecture is used, VCO_2 must be tunable at spacings of f_{CH}/N_1, leading to a reference of $f_{ref} = f_{CH}/N_1$ as opposed to f_{CH} in a classical PLL. This means the total division ratio is multiplied by N_1. For an N_1 of 10, this means a 20dB increase in the phase-noise and spur contribution at the output by the input phase-detector (PD_2) and charge-pump (CP_2). Also, since the reference frequency is reduced by a factor of N_1, so will the BW of the low-frequency PLL for adequate suppression of that spur. This means that the settling time would increase ten times. This cannot be tolerated in most cases.

In our architecture however, the novel method in which the feedback is tapped, means that the output of VCO_2 achieves the right resolution while the reference, f_{ref} remains equal to the channel spacing like a classical PLL. Thus our architecture achieves the required output resolution without affecting the overall division ratio or the settling time of the loop.

Another advantage of our architecture is that, even with a fixed f_{ref}, the division ratio N_1 can be made variable without affecting the resolution of the output. This may be beneficial in reducing the required tuning range of VCO_2 and/or reducing its gain variation (by utilizing only part of the tuning headroom available). Since

the gain of the VCO affects the loop BW, having a largely varying VCO gain is undesirable [63].

The above architecture achieves the integration of the RF VCO at the expense of an external IF VCO. However, the IF VCO operating at the lower IF frequency does not suffer from interference problems with the PA or leakage problems. Also, driving the large parasitic capacitances at this frequency will be much less power consuming. A pre-amplifier, if needed, will also consume much less power. In addition, the IF VCO can be easily integrated on-chip with an external tank, since at these low frequencies the board parasitics and bondwire inductances would be negligible with respect to the tank components' values. This way, a completely monolithic PLL can be achieved with only an external inductor and varactor. Furthermore, if the transceiver architecture described in section 4.4 is adopted, the already-available IF VCO can be utilized and the nested-loop PLL architecture can be built with no extra components.

Compared to an ordinary PLL, the architecture in Fig. 4.2 needs the extra components shaded in gray. Hence, from a power consumption point-of-view, the integration of the VCO is achieved at the expense of the extra power needed for VCO_2, divider N_1 and the indicated phase-detector and charge-pump. The phase-detector and charge-pump usually consume little power in a PLL, especially when the division ratio, N_1, is small and their phase-noise will not affect the output. Hence the power spent in PD_1 and CP_1 should not be significant. The power consumption of the frequency divider should also be small, since it is of fixed value and does not need the front-end dual-modulus counter needed in prescalers which

consumes most of the power. Also, since VCO_2 operates at a low frequency with an external tank, its current consumption would not be as high as the RF VCO. Thus the extra power consumption needed should only be moderate. This extra power is offset by the power savings achieved by integrating the VCO and by the fact that the integrated VCO does not need to achieve a very low level of phase-noise since it is being helped by the wide BW of the inner loop. Again, if the transceiver architecture of section 4.4 is employed, then no extra power will be needed.

4.3.2 Analysis

Loop Stability

Since the nested-loop PLL architecture involves one loop inside the other, it is important to study the stability issues. The following analysis shows that if the inner loop is much faster than the overall bandwidth, the loop will remain stable with no degradation in the phase margin.

An s-domain block diagram of the overall PLL is shown in Fig. 4.3. Inputs θ_{n1} and θ_{n2} represent the phase-noise of both VCOs. Assuming that loop filters, $H_1(s)$ and $H_2(s)$ are simple first-order RC combinations as in Fig. 3.6 then

$$H_1(s) = \frac{I_{P1}}{2\pi C_1} \cdot \frac{1+s\tau_1}{s}, \qquad H_2(s) = \frac{I_{P2}}{2\pi C_2} \cdot \frac{1+s\tau_2}{s} \tag{4.1}$$

where I_P is the charge-pump current and C and τ the filter capacitor and time constant respectively (section 3.3.2). Using (3.11) and (4.1), the input/output

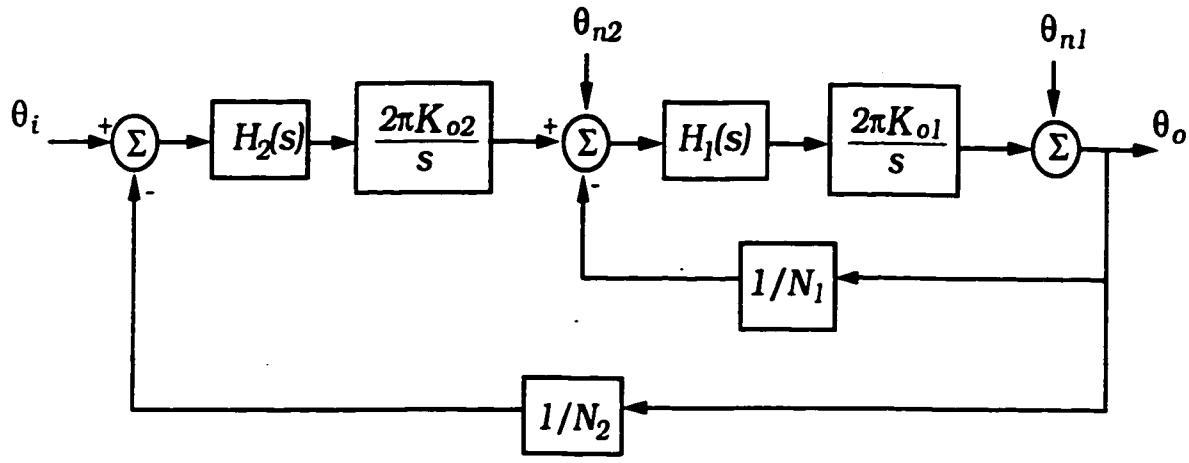

Figure 4.3: S-domain block diagram of the nested-loop PLL.

transfer function can be written as

$$\frac{\theta_o}{\theta_i} = N_2 \frac{4\zeta_2\zeta_1\frac{\omega_{n2}}{\omega_{n1}}(\frac{s}{\omega_{n2}})^2 + 2\zeta_2(1+\frac{\zeta_1\omega_{n2}}{\zeta_2\omega_{n1}})(\frac{s}{\omega_{n2}}) + 1}{(\frac{\omega_{n2}}{\omega_{n1}})^2(\frac{s}{\omega_{n2}})^4 + 2\zeta_1\frac{\omega_{n2}}{\omega_{n1}}(\frac{s}{\omega_{n2}})^3 + (1+4\zeta_2\zeta_1\frac{\omega_{n2}}{\omega_{n1}})(\frac{s}{\omega_{n2}})^2 + 2\zeta_2(1+\frac{\zeta_1\omega_{n2}}{\zeta_2\omega_{n1}})(\frac{s}{\omega_{n2}}) + 1} \tag{4.2}$$

where

$$\omega_{n2} = \sqrt{\frac{I_{p2}K_{o2}N_1}{CN_2}}, \quad \zeta_2 = \frac{\tau_2\omega_{n2}}{2} \tag{4.3}$$

Note that ω_{n2} and ζ_2 are the nominal natural frequency and damping factor of the overall loop respectively when the inner loop is considered to be an ideal frequency multiplier with gain N_1 and infinite BW.

The inner loop adds two extra poles to the transfer function. It is important to make sure that all poles lie in the left-hand S-plane so that the system remains

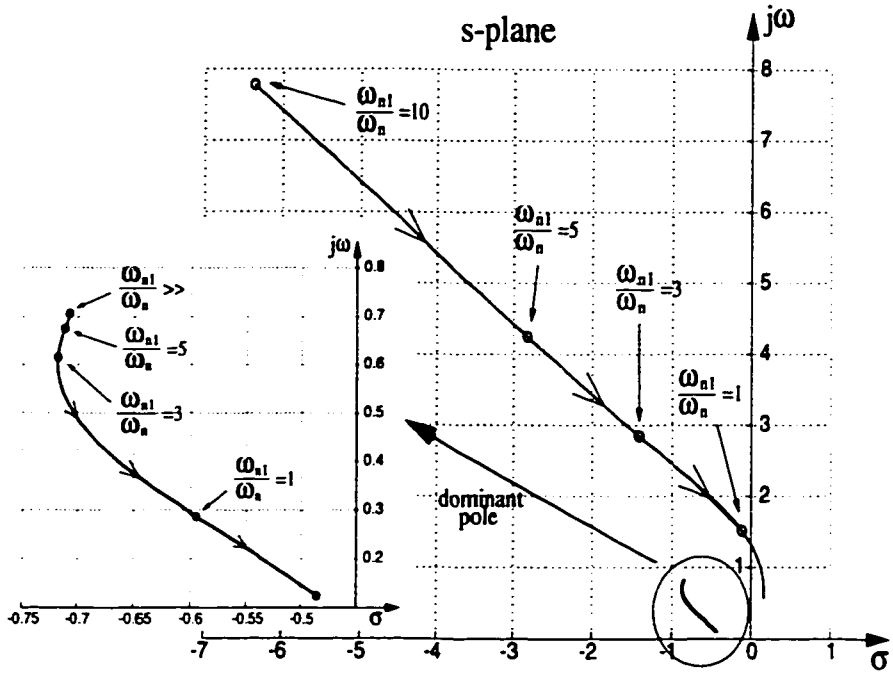

Figure 4.4: Normalized root locus of the poles of Fig. 4.3, $\zeta_2 = \zeta_1 = 0.707$.

stable, and that the extra poles do not appreciably move the nominal poles from their locations hence altering the nominal loop parameters and degrading the phase margin.

The locii of the poles of (4.2) are shown in Fig. 4.4. Nominal damping factors of $\zeta = \zeta_1 = 0.707$ are assumed and the natural frequency of the overall loop, ω_{n2}, is normalized to 1. Since the poles are conjugate pairs, only the poles in the upper half-plane are shown. It is clear that, when $\omega_{n1} \gg \omega_{n2}$, the dominant pole is at the nominal location of $-0.707 \pm 0.707j$ and the parasitic poles are at a distant negative frequency. As ω_{n1}/ω_{n2} starts decreasing, the parasitic poles move to the right, displacing the dominant poles from their nominal location slightly.

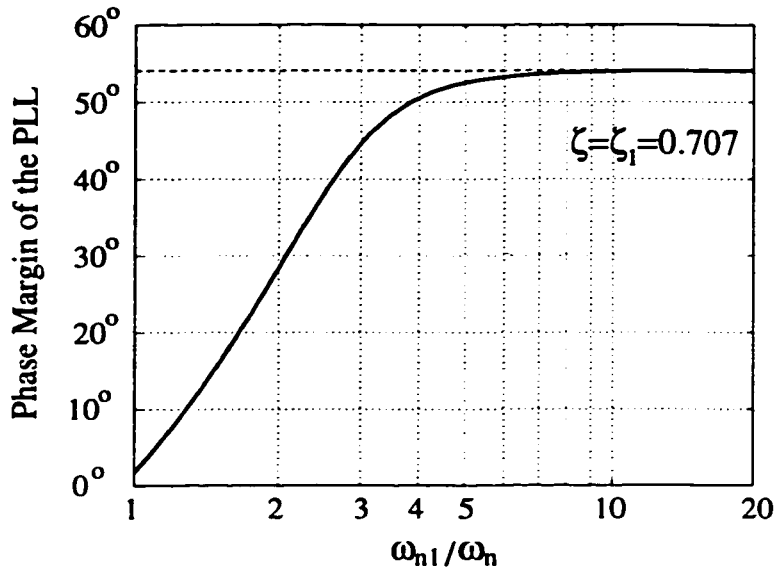

Figure 4.5: Effect of the inner-loop bandwidth on the phase-margin of the overall PLL.

As shown in the figure, for a ω_{n1}/ω_{n2} ratio as small as 3, the dominant pole moves only slightly causing a very small change in the loop natural frequency and damping factor. As the ω_{n1}/ω_{n2} ratio is decreased further (less than one in this case), the parasitic pole moves into the right half plane making the loop unstable. Hence, for this architecture, as long as the inner loop is sufficiently faster than the overall bandwidth (e.g. $\omega_{n1} > 5\omega_{n2}$), the loop stability is guaranteed and the natural frequency and damping factor of the overall loop will be very close to the nominal values given by (4.3).

The phase-margin of the overall PLL is shown in Fig. 4.5, for $\zeta_2 = \zeta_1 = 0.707$. A capacitance of value $C_2/10$ is assumed to be placed in parallel to the RC series combination of $H_2(s)$ in Fig. 4.3 in order to mitigate the ripples on the charge-

pump output. As shown, for $\omega_{n1} > 5\omega_{n2}$ the phase-margin degradation is less than 2°.

Transfer Functions

The input/output transfer function is given by (4.2). If $\omega_{n1} \gg \omega n2$, then this relation can be written as

$$\frac{\theta_o}{\theta_i} \approx N_2 \cdot \frac{2\zeta_1(\frac{s}{\omega_{n1}}) + 1}{(\frac{s}{\omega_{n1}})^2 + 2\zeta_1(\frac{s}{\omega_{n1}}) + 1} \cdot \frac{2\zeta_2(\frac{s}{\omega_{n2}}) + 1}{(\frac{s}{\omega_{n2}})^2 + 2\zeta_2(\frac{s}{\omega_{n2}}) + 1} \tag{4.4}$$

which is a cascade of two low-pass functions. Similarly, the transfer functions between the output and the phase-noise of the RF and IF VCO is given respectively by

$$\frac{\theta_o}{\theta_{n1}} \approx \frac{(\frac{s}{\omega_{n1}})^2}{(\frac{s}{\omega_{n1}})^2 + 2\zeta_1(\frac{s}{\omega_{n1}}) + 1} \cdot \frac{(\frac{s}{\omega_{n2}})^2}{(\frac{s}{\omega_{n2}})^2 + 2\zeta_2(\frac{s}{\omega_{n2}}) + 1} \tag{4.5}$$

and

$$\frac{\theta_o}{\theta_{n2}} \approx \frac{2\zeta_1(\frac{s}{\omega_{n1}}) + 1}{(\frac{s}{\omega_{n1}})^2 + 2\zeta_1(\frac{s}{\omega_{n1}}) + 1} \cdot \frac{(\frac{s}{\omega_{n2}})^2}{(\frac{s}{\omega_{n2}})^2 + 2\zeta_2(\frac{s}{\omega_{n2}}) + 1} \tag{4.6}$$

Equation (4.5) is a cascade of two high-pass functions, while (4.6) is a cascade of a low-pass and a high-pass function. All three transfer functions are illustrated in Fig. 4.6.

From that figure, it is clear that, like any ordinary PLL, any phase-noise at the input as well as any spurs from CP_2 will be amplified by $20log(N_2)$. It is thus necessary to keep ω_{n2} sufficiently small to reject those spurs and phase-noise. Input

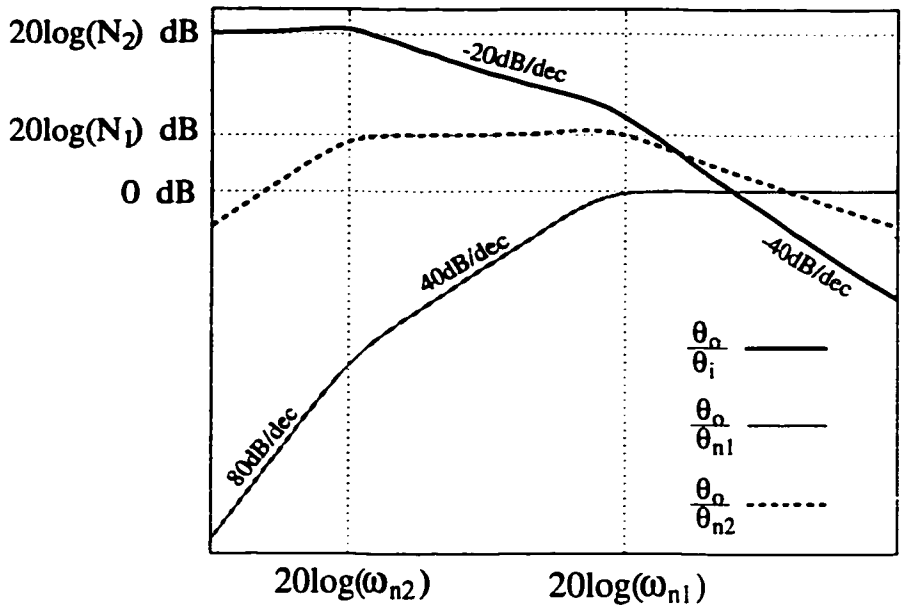

Figure 4.6: PLL transfer functions

noise and spurs at offsets larger than ω_{n1} will be attenuated by 40dB/decade. This extra order of attenuation could be useful in rejecting higher-frequency noise.

The RF VCO phase-noise that is within the inner loop BW (less than ω_{n1}) will be attenuated by 40dB/decade. This will greatly reduce the VCO phase-noise. We get an extra order of correction here as well due to the pole of the overall loop. Thus the VCO phase-noise that is very close to the carrier (less than ω_{n2}) will be attenuated by 80dB/decade. This might be useful in reducing the $1/f$ and $1/f^2$ noise of the oscillator.

The transfer function between the output and the phase-noise of the IF VCO is critical. As shown in the figure, any IF VCO phase-noise that lies at frequencies greater than ω_{n1} will be attenuated by the inner loop transfer function. Similarly,

any noise that is below ω_{n2} will be rejected by the overall loop. The phase-noise that lies in-between is critical though, as it will be increased by $20log(N_1)$ when mapped to the output. It is thus necessary to design the IF VCO with a phase-noise that is $20log(N_1)$ dB below the level required at the output. As discussed below, since the IF VCO operates at a low frequency and uses an external resonator, this should not be a problem.

IF VCO

The IF VCO operates in the 100MHz range. At that frequency the values of the resonance circuit components are too large to integrate on a chip. It is thus necessary to use an integrated VCO with external passive components. Even though the external tank does consume some board space, it does not suffer from most of the problems of an external RF VCO.

Due to the very large frequency offset, the problems of LO leakage and power-amplifier interference will not exist. Also, at this frequency, driving the external parasitic capacitances will not be power consuming, and no pre-amplifier will be needed. The IF VCO can still be designed in a differential manner and thus be more immune to noise and cross-talk. In addition, unlike an external RF VCO tank, the parasitic board capacitors and bondwire inductors are negligible with respect to the IF tank components. This makes the design much easier. Finally, the large values of the tank components may allow the inclusion of ESD protection on the extra pins.

As discussed above, the phase-noise of the IF VCO is a critical parameter and should be $20log(N_1)$ dB below the level required at the output. This should not

be a problem, however, since the IF VCO operates at $(1/N_1)^{th}$ the frequency of the output. Since the phase-noise of a VCO is proportional to the square of the operating frequency (equation (3.8)), the reduced operating frequency results in a reduction of $20log(N_1)$ dB in the phase-noise. Hence, with an external commercial inductor with a typical quality factor above 50 in that range of frequencies [64,65], meeting the phase-noise requirement should be easy. In addition, the high quality factor means that the phase-noise requirement can be met with a small bias current. The Colpitts oscillator in [66] operates at 180MHz and achieves a phase-noise of -118dBc/Hz at 25kHz offset using an external LC tank. This scales to -109dBc/Hz at 10kHz from 90MHz. The oscillator consumes a bias current of only 48μA. Since the phase-noise of a Colpitts oscillator decreases by 6dB when the bias current is quadrupled [54], a very low phase-noise level can be achieved while still consuming little current.

4.4 High-Integration Transceiver Architecture

Compared to a classical PLL, the nested-loop PLL architecture needs the additional components shaded in gray in Fig. 4.2. If the IF VCO is integrated on-chip with an external tank circuit, then the extra discrete components and board area will be minimal. Also, as argued above, the extra power consumption needed for those components should not be large. Hence, if the nested-loop PLL replaces the classical PLL in the frequency synthesizer of a wireless radio (with no other system modifications), a relatively inexpensive price is paid in order to realize the benefits listed earlier.

In this section, however, we present a transceiver architecture that attempts to extend the level of integration and power reduction by utilizing the IF VCO that is already available in any superheterodyne architecture in the nested-loop PLL. In this way, the transceiver can reach a very high level of integration while minimizing power consumption. The drawbacks of such an architecture are also discussed.

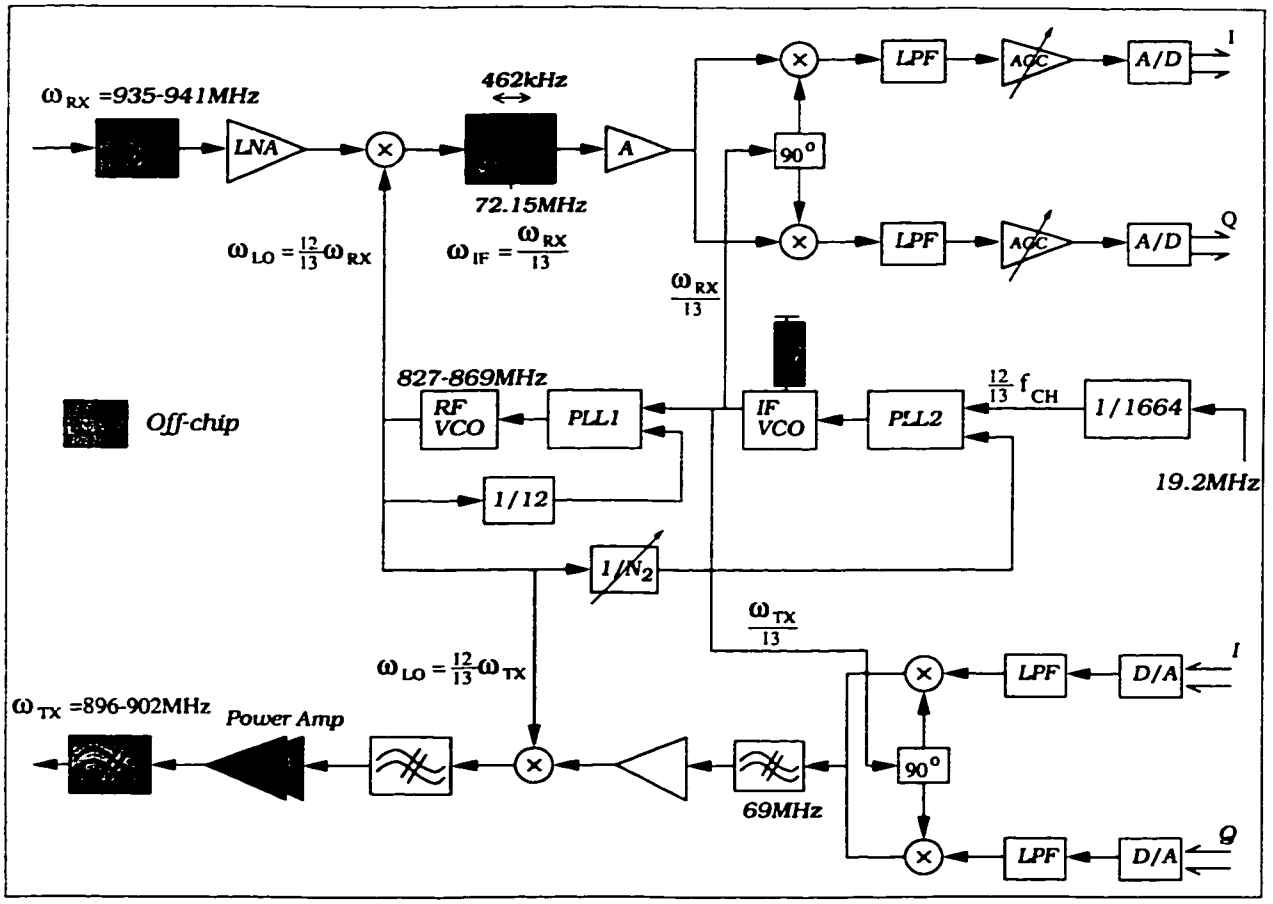

Figure 4.7: Highly-integrated transceiver architecture.

The transceiver architecture is shown in Fig. 4.7. The specific numbers and division ratios are geared towards the Mobitex data network [32]. This architecture is a double conversion transmitter as well as a double-conversion receiver architec-

ture. The low-frequency VCO that drives the WBPLL, is utilized to downconvert the received signal from IF to baseband and upconvert the transmitted one from baseband to IF.

In Fig. 4.7, the inner division ratio was chosen to be 12. For the shown bands, this gives a receive IF frequency of about $72MHz$ and a transmit IF of $69MHz$. Since the LO frequency is 12 times the IF frequency and $\omega_{RF} = \omega_{LO} + \omega_{IF}$, then $\omega_{IF} = \omega_{RF}/13$ and $\omega_{LO} = 12\omega_{RF}/13$. Hence the IF frequency is not fixed, but depends on the particular channel location. This has several implications discussed below. Since the LO frequency is equal to 12/13 of the RF frequency, then the reference frequency to PLL_2 must also be equal to 12/13 of the channel spacing. For a channel spacing of $12.5kHz$, this value could be easily obtained from a commercial crystal of $19.2MHz$ as shown.

This architecture achieves a very high level of integration. The only components in the above architecture that need to be off-chip are the ones indicated. As shown, the complete transceiver can potentially be integrated on a single-chip with only the front-end SAWs and power amplifier, the external IF filter and an external tank circuit. This leads to a great saving in cost and area. Since this architecture does not use any additional components as compared to the widely used superheterodyne architecture [1], there is thus no power penalty. In fact, power consumption should be reduced due to the integration of the UHF VCO, as discussed earlier.

When compared to the wide-band IF architecture of [14], we realize that this architecture attains the same benefits of having a wide-bandwidth main PLL for VCO integration. This is done while still maintaining the UHF VCO tunable, thus

eliminating the need for quadrature downconversion and the use of extra mixers, as well as having to deal with the full receive band at IF (section 2.2.2).

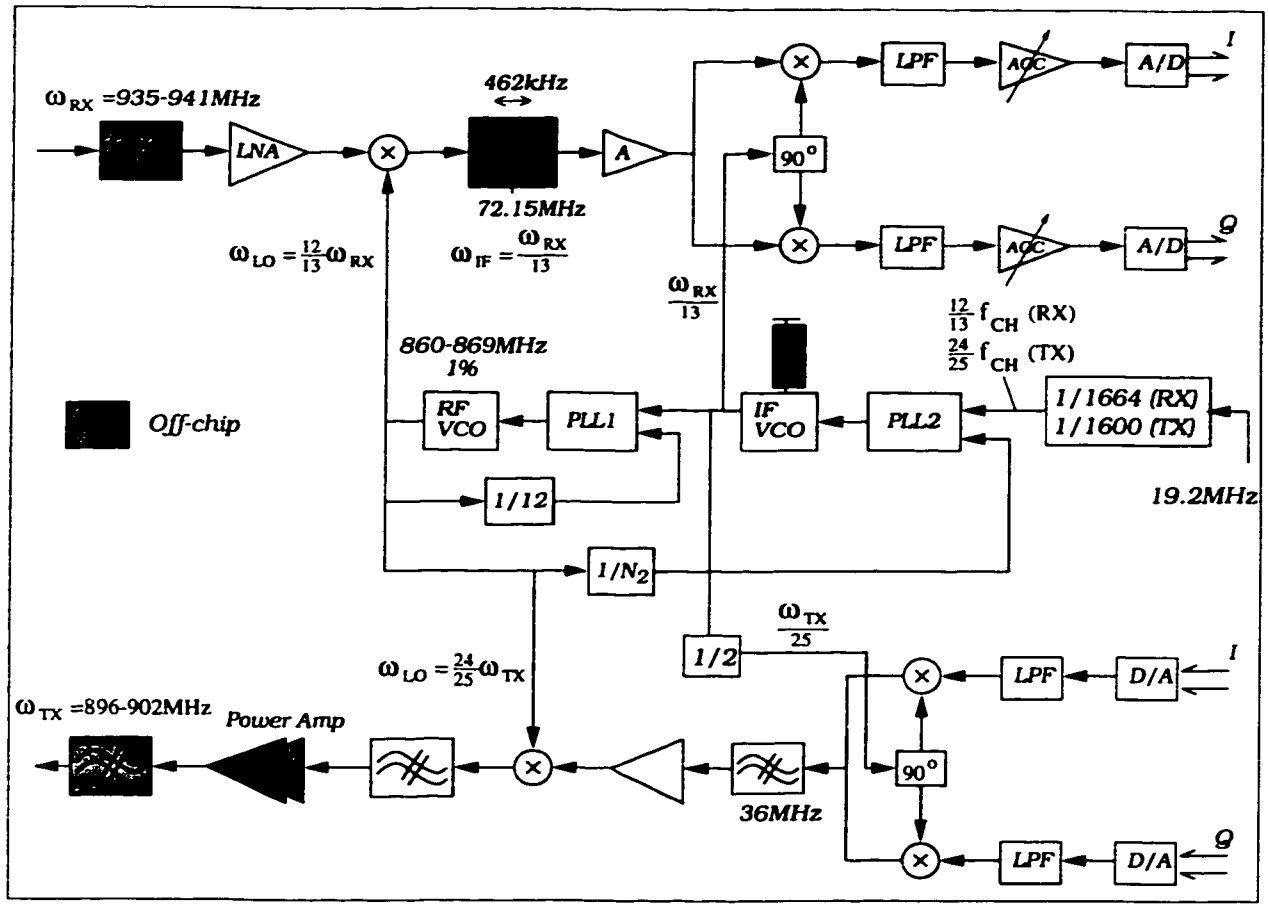

Figure 4.8: Highly-integrated transceiver architecture, with reduced tuning range of the UHF VCO.

Since the values of both the transmit and receive IF frequencies are close in the architecture of Fig. 4.7, the RF VCO must be able to cover the full range between the transmit and receive bands (800-868MHz). This means an 8.2% tuning range. Although this could be easily achieved by monolithic VCOs [56, 59], the required tuning range can be reduced by changing the transmit IF frequency, as shown

in Fig. 4.8. By adding an extra divide-by-2 flip-flop, the transmit IF frequency becomes 36MHz. In the receive mode, the LO frequency and the PLL_2 reference are the same. In the transmit mode, the reference to PLL_2 is changed to $24f_{CH}/25$ (by changing the division ratio of the preceding counter), and the LO frequency becomes $24\omega_{RF}/25$. This results in a tuning range of the UHF VCO of only 1%.

4.4.1 Design Considerations

The main drawbacks of the architecture of Fig. 4.7 stem from the fact that the IF frequency is not fixed. In the shown example, the IF frequency is equal to $\omega_{RF}/13$. Since the receive band has a 6-MHz span, then the pass band of the IF filter has to be 385kHz wide (6MHz divided by 13). This means that, unlike a superheterodyne architecture where the IF filter passes only the desired channel (and possibly a few adjacent ones), the IF filter will pass 31 channels unattenuated (the desired channel being in the middle). This has undesirable implications.

Since the adjacent channels are not completely filtered-out at IF, then this imposes a requirement on the IF VCO phase-noise that is more stringent than the superheterodyne architecture. However, in our case this is not a drawback, since such a requirement is already imposed, because the phase-noise of the IF VCO will be mapped to the frequency synthesizer output. As discussed earlier, meeting this phase-noise requirement should not be difficult.

The main problem with having a wide IF filter, however, is that the subsequent stages will have to deal with a larger dynamic range. The amplifier A in Fig. 4.7 and the subsequent quadrature mixers will have to have a sufficiently large

dynamic range to accommodate the power of 30 channels, as opposed to the only one channel in a superheterodyne architecture. We do note that the LNA and front mixer deal with the power of the full band (480 channels), and the IF SAW filters-out more than 92% of the power of the composite signal. Thus such a drawback may not be a problem in many cases [67]. In the WBIF architecture of [14], the IF amplifiers and mixers accommodate the full band with no filtering, thus requiring a much larger dynamic range than in our case. The wide IF pass-band can be an undesirable consideration nonetheless, especially in systems where the noise at lower frequencies is large. The wide IF pass-band also prevents placing the AGC before the baseband LPF.

4.5 Wide-bandwidth PLL Circuit Design

The wide-BW PLL at the core of the nested-loop architecture was designed and fabricated in a 25GHz triple-metal bipolar process. A block diagram of the fully integrated loop is shown in Fig. 4.9. The VCO used was a monolithic LC-oscillator, with on-chip single-layer spiral inductors having an unloaded Q of around 7.5 at 1GHz. The division ratio, N_1, was chosen to be 16 implemented by a simple 4-stage ripple counter. The target operating frequency was in the 1GHz range, giving a reference frequency at the input of the phase detector, PD_1, around 65MHz.

The loop bandwidth is approximately 1.5MHz. Since the bandwidth is relatively large, it was possible to integrate the loop filter on-chip, thus avoiding exposure to external noise sources. The grounded bottom plates of the loop capacitors help in shielding them from the substrate noise.

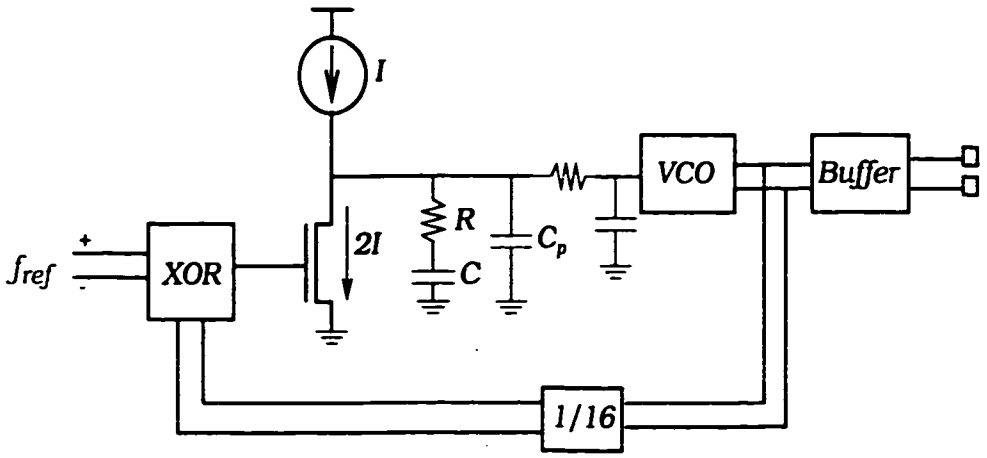

Figure 4.9: Wide-BW PLL schematic.

An XOR gate was used as the phase-detector. This results in a simple deadzone-free phase detector with low in-band noise as discussed in section 4.5.3. The charge-pump consists of a switched npn current source and a constant-current pull-up pnp transistor. Although it dissipates slightly more d.c. current, the constant-current pull-up was necessary due to the lack of a sufficiently fast pull-up device in the technology.

4.5.1 VCO

The schematic of the VCO is shown in Fig. 4.10 [68]. The inductor, L, and varactor, D, form a parallel-resonance tank. The emitter coupled pair, Q_1 and Q_2, provide the negative resistance necessary to sustain oscillations. The value of the negative resistance is determined by the g_m of the transistors and the emitter resistances.

The emitter-follower buffers Q_3 and Q_4 serve two purposes. First, they provide a d.c. shift between the bases of the cross-coupled pair and their collectors. This

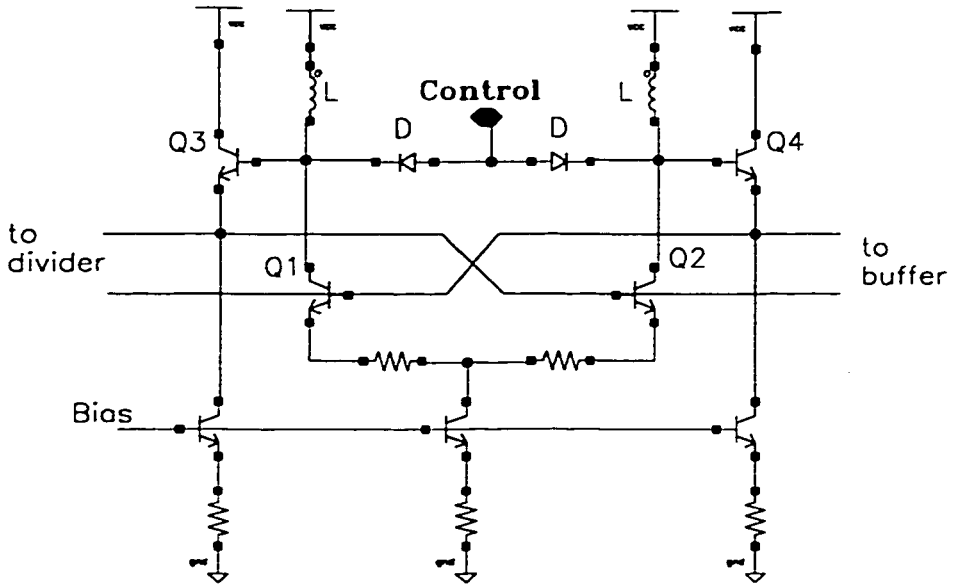

Figure 4.10: VCO schematic.

allows for a larger tank swing without driving Q_1 and Q_2 into saturation. Second, Q_3 and Q_4 buffer the loads driven by the VCO from the tank, thus increasing the loaded quality factor.

The inductor used was a monolithic metal-3 spiral inductor, with a value of 7.5nH and a loaded Q of approximately 7 at 1GHz. The peak-to-peak swing on the tank was 1.1V giving an r.m.s. value of 0.39V. With an open-loop gain of 6.5, the phase-noise at 1MHz offset from 1.1GHz is expected to be around -117dBc/Hz (equation (3.8)). The phase-noise level could be reduced by decreasing the excessive open-loop gain. The VCO core consumes 2.3mA from a 3.3V supply.

4.5.2 Frequency Divider

The frequency divider is a simple divide-by-16 counter composed of a cascade of four divide-by-two flip-flops. One stage is shown in Fig. 4.11. This one-bit counter consists of a master-slave D flip-flop. Emitter follower buffers are added to provide a d.c. shift so that the transistors are not driven in the saturation region.

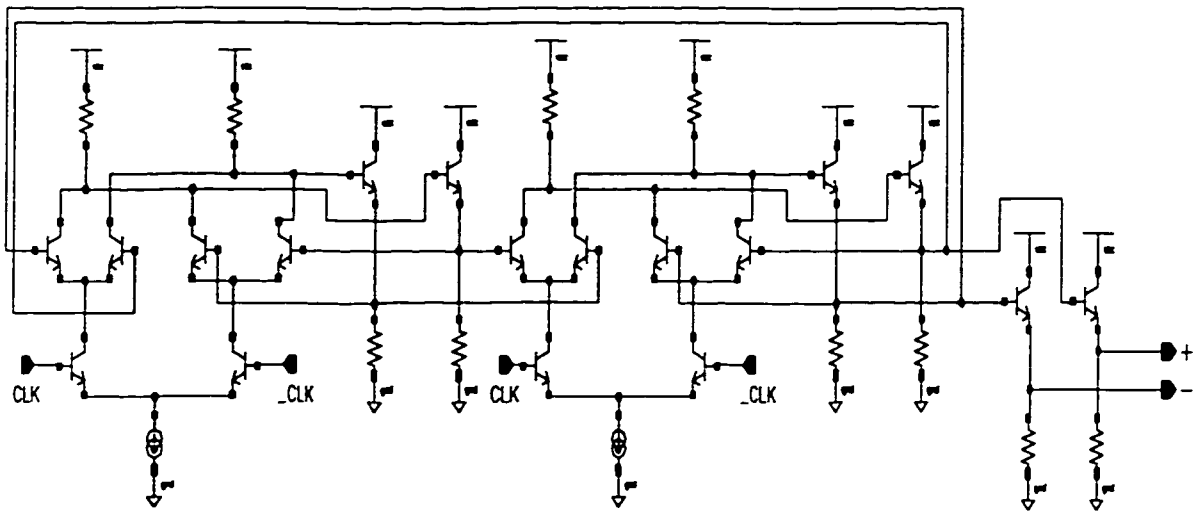

Figure 4.11: Divide-by-2 stage.

The output frequency of each stage is half the input frequency. The pull-up resistors can thus be doubled each stage while maintaining enough bandwidth. Hence, for a constant swing, the current in each stage is approximately half that of the preceding stage.

4.5.3 Phase-detector

The phase-detector was implemented with an XOR gate. In addition to its simplicity, the XOR gate has several benefits. First, the XOR gate has no dead-zone,

which can significantly increase the in-band noise, [52,69] particularly when operating at a high reference frequency [70]. This is particularly important here since the loop bandwidth is wide. Second, when an XOR gate is used, its two inputs are 90° out-of-phase during lock. This ensures time-domain isolation of the transitions of the reference signal and the divider output, thus reducing the 'digital feedthrough'. It was shown that in-band noise is lowest when the edges were separated as widely as possible (90°) [52]. A third benefit is that the XOR phase-detector is immune to mismatches in the UP and DOWN current of the charge-pump and to rise/fall time variations, since the loop will automatically respond to these effects by adjusting the duty-cycle of the XOR output such that the average current integrated on the capacitor remains zero. Finally, the XOR phase-detector has a high gain ($2I_p/\pi$ as opposed to $I_p/2\pi$ for a type-3 PFD [36]) for the same bias current. This reduces the input-referred noise due to any noise at the detector output.

The disadvantages associated with the XOR phase-detector would not be a problem in our case. The use of an XOR phase-detector leads to a higher reference feedthrough. However, the spur caused is located at twice the reference frequency. The spacing between the reference frequency and the loop natural frequency is thus very large allowing for adequate spur suppression (especially since there is more room to add extra poles in the loop filter). The other disadvantage of the XOR gate is the limited lock range [51]. In case of a frequency step that is larger than the lock range, the loop will take several additional cycles to regain lock (pull-in time) and hence a longer time to settle (as in Fig. 4.14). However, this is not a problem for two reasons. First, the very wide-BW of the loop ensures that the

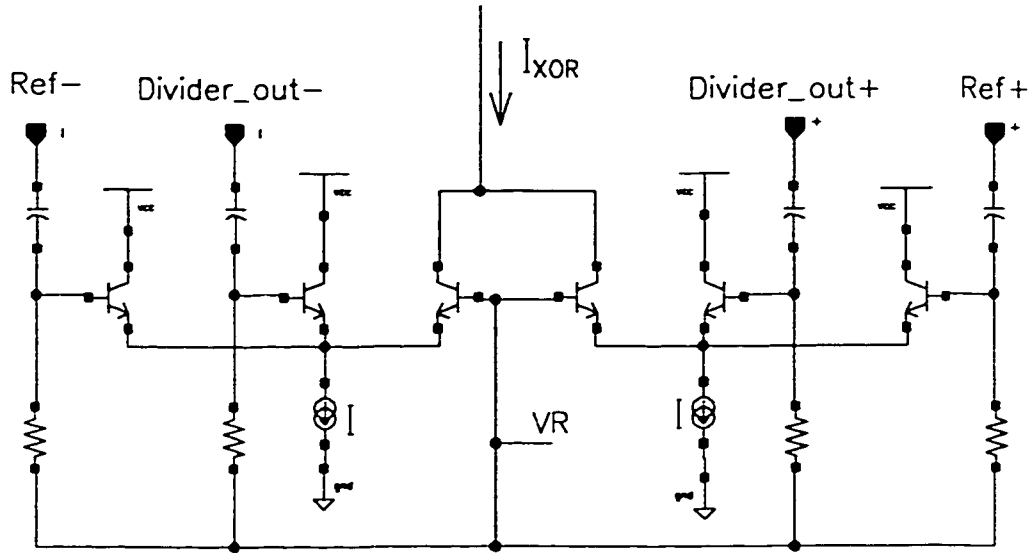

Figure 4.12: Phase-detector schematic.

settling time remains very short even if several additional cycles are needed for pull-in. Second, this loop is to be configured in a larger PLL as shown in Fig. 4.2. Since the bandwidth of the outer loop is very small, then the frequency of the IF VCO will be slowly varying, and no sudden frequency steps of magnitude greater than the lock range will occur at the input.

The differential XOR gate was implemented as shown in Fig. 4.12. The differential output of the frequency divider and the differential reference frequency are a.c. coupled to the inputs of two differential pairs. All the inputs are at the same d.c. level, V_R. If the differential value of both inputs is either positive or negative simultaneously, then the XOR output, I_{XOR} will be zero. If one input is positive while the other negative, I_{XOR} will be equal to the tail current I. Thus this differential gate implements a differential XOR function with a current output.

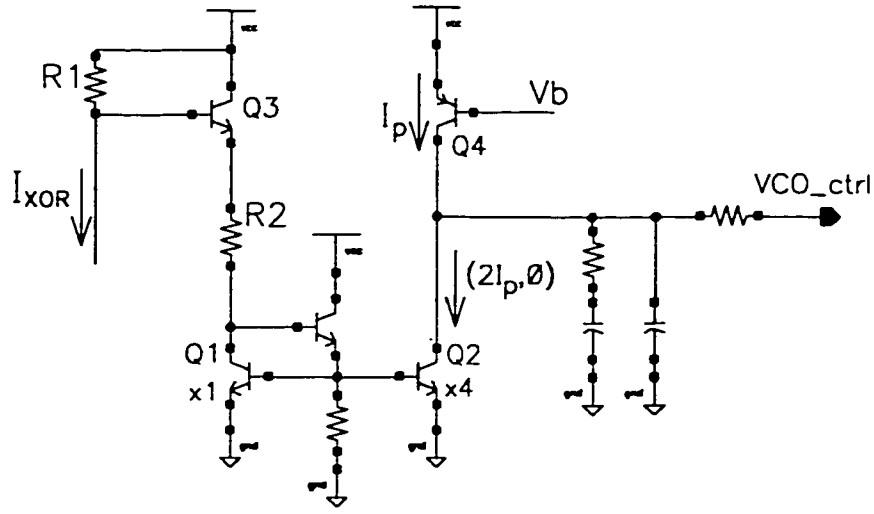

Figure 4.13: Charge-pump schematic.

4.5.4 Charge-Pump

The output of the XOR phase-detector is used to activate the charge-pump shown in Fig. 4.13. Due to the lack of a fast p-type device in the technology, it is not possible to use a current mirror to switch the charge-pump. Instead, an npn voltage-follower is used. When the XOR output current is zero, the drop on R_2 is $V_{cc} - 3V_{BE}$ and the current mirror is turned ON. When I_{XOR} is non-zero, the voltage headroom is reduced by $I_{XOR}R_1$ which is chosen to be large enough to make the current in R_2, and hence Q_2, zero.

The UP current of the charge-pump is a constant current, I_p, coming from current source Q_4. This constant d.c. current dissipated is necessary due to the low f_t of the pnp bipolar in the technology. During lock, the DOWN current of the charge-pump is a square-wave with a 50% duty cycle with levels of 0 and $2I_p$. The average current to the loop filter is thus zero. If the ratio of the UP and DOWN

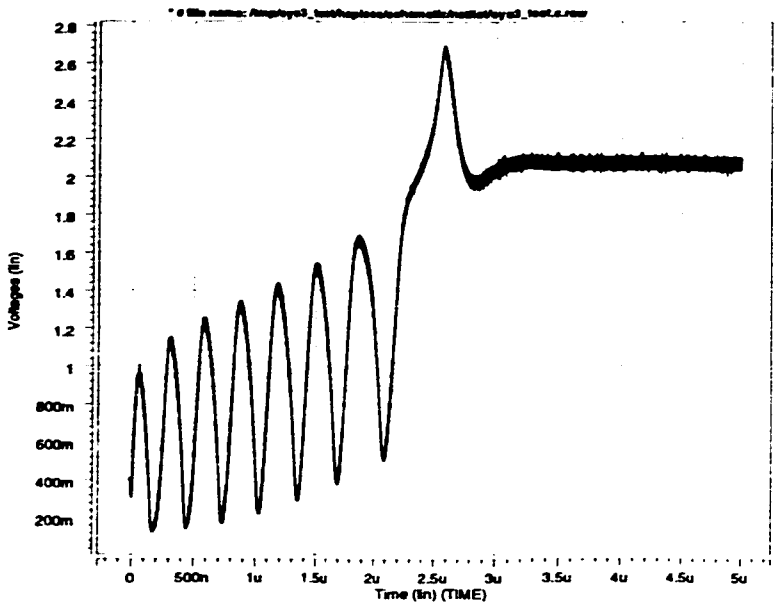

Figure 4.14: Hspice transient simulation of the WBPLL.

current is not properly matched, or the '0' current of Q_4 is finite, this will only lead to the loop adjusting the duty cycle so as to maintain a zero average current during lock.

The loop filter, also shown in Fig. 4.13, consists of the classical RC series combination creating the dominant pole plus two extra high-frequency poles that reduce the high-frequency noise while maintaining adequate phase-margin. The complete loop filter was integrated on the chip. The total capacitance in the filter was approximately $90pF$.

Post-layout simulations were conducted on the complete PLL using Hspice. Fig. 4.14 shows the transient response of the VCO control voltage for a 70MHz frequency step. The first part of the curve shows the pull-in process of the loop since the frequency step is larger than the lock range. After 2.25μs, the loop re-enters

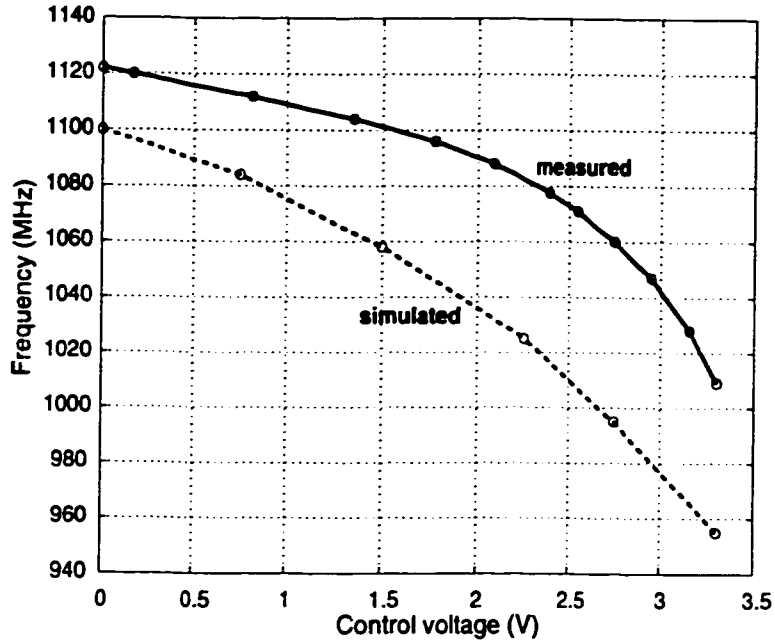

Figure 4.15: V-F characteristics of the VCO.

the linear region and will take approximately 2.6μs more for settling, resulting in a total settling time of only 5μs. The high frequency spurs on the VCO control are due to the use of an XOR phase-detector.

4.6 Measurements

The wide-BW PLL was fabricated in a 25GHz triple-metal bipolar technology. The circuit core occupied $1.7mm^2$. The power consumption was 9.9mA from a 3.3V supply plus 5mA for the output buffer.

The tuning characteristics of the LC-VCO is shown in Fig. 4.15. Due to process variations, the center frequency is 3.6% higher. The measured tuning range is

10%. The phase-noise of the free-running VCO, shown in Fig. 4.16, is equal to -114dBc/Hz at 1MHz offset, close to the calculated value of -117.

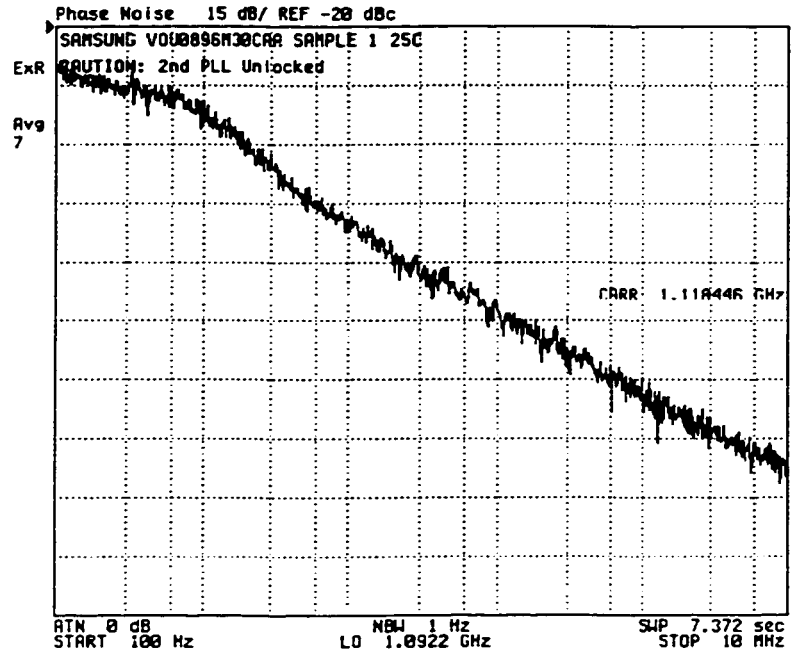

Figure 4.16: Phase-noise of the free-running VCO, reference at -20dBc and vertical scale=15dB/division.

The reference signal to the PLL was derived from a signal generator. The input reference's measured phase-noise was -126dBc/Hz at 10kHz and had a noise floor of -133dBc/Hz. The phase-noise performance of the closed-loop synthesizer is shown in Fig. 4.17. Within the loop BW, the noise is -108dBc/Hz, indicating that it is limited by the noise floor of the input reference. At offsets less than 20kHz, the phase-noise starts to rise, also limited by the input reference phase-noise. The phase-noise at 10kHz is equal to -100dBc/Hz. To our knowledge, this is the lowest phase-noise reported in the literature at this frequency. At this offset, this is equivalent to

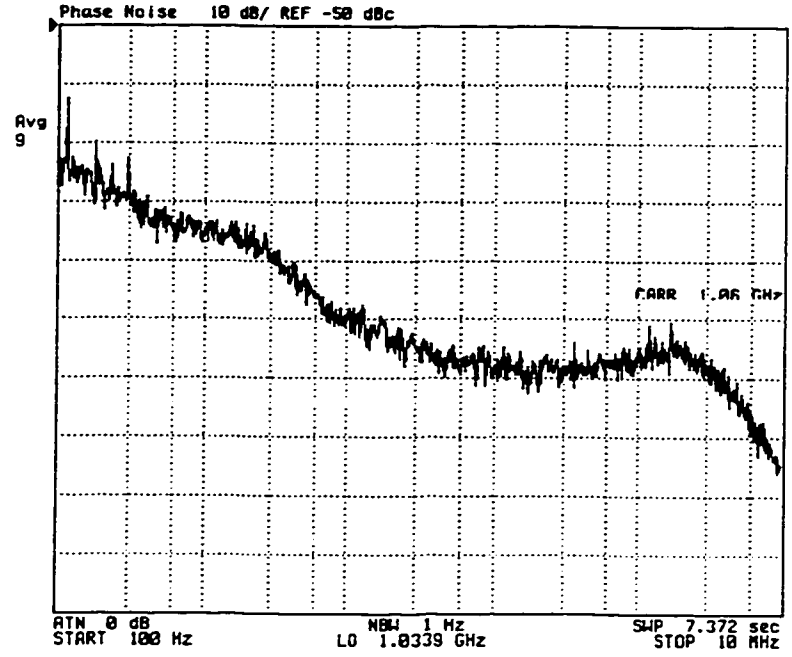

Figure 4.17: Phase-noise of the PLL output, from 1.06GHz.

an oscillator with an effective Q of more than 20. The use of a better reference source would even improve the output phase-noise performance. The frequency synthesizer exhibited spurs of -52dBc at twice the reference frequency.

Table 4.1 compares this work to other LC-VCO PLLs in the literature that have a wide bandwidth. The figure-of-merit proposed is defined as

$$FOM = 20log(\frac{f_o}{f_{ref}}) - \Phi - 10logP \tag{4.7}$$

where P is the power consumption of the PLL core (excluding any digital control or output buffers) and Φ is the phase-noise level within the BW. The ratio f_o/f_{ref} indicates the division ratio used which directly affects the phase-noise within the

Ref.	f_o (Hz)	f_{ref} (Hz)	BW (Hz)	Phase-noise (dBc/Hz)		P (mW)	Comments	FOM
				in BW	at 10kHz			
[71]	1.6G	53.4M	450k	-91	-90	185	BJT, Bondwire L	127
[63]	1.8G	26.6M	45k	-90	-88	51	CMOS	140
[72]	1.6G	61.5M	200k	$\approx$ -97	$\approx$ -95	62	CMOS, On-chip loop filter	137
[73]	2G	2G	200M	-110	?	1.6	BJT, Needs 2GHz reference	138
[52]	0.9G	20M	100k	-95	-95	48	BJT, External VCO	141
[60]	1.8G	20M	84k	-73	-73	25 + VCO	CMOS, External VCO	128
This work	1.06G	66M	1.5M	-108	-100	33	BJT, On-chip loop filter	147

Table 4.1: Comparison of this work to other wide-BW PLLs with LC-oscillators in literature.

BW. This is more meaningful than the ratio of f_o to the offset frequency. The figure-of-merit of our work is 6dB better than [52], which uses an external VCO, and is 7dB higher than the nearest monolithic VCO implementation. A micrograph of the chip is shown in Fig. 4.18. The total chip area was 2.5x2.4mm^2.

4.7 Conclusion

Today's RF IC technologies attempt to maximize the quality factors of monolithic inductors. While this results in incremental improvements, the realization of a completely integrated VCO that meets the required phase-noise levels proves very

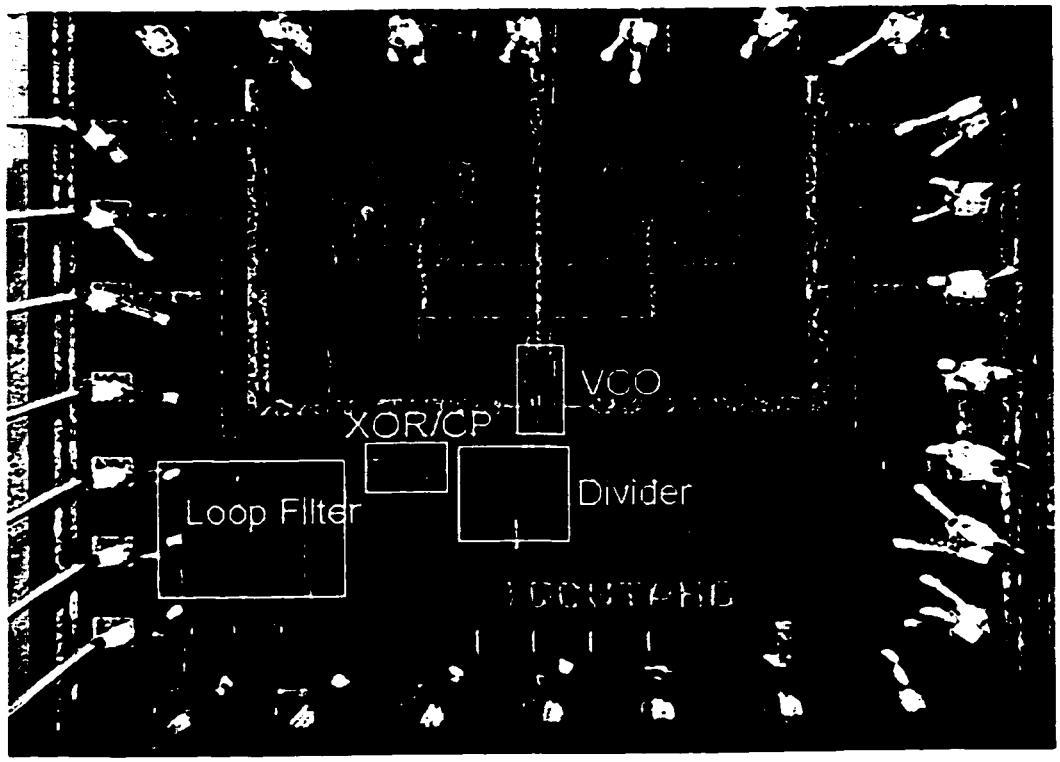

Figure 4.18: Chip Micrograph.

difficult and a distant goal for many communications standards, especially over all process and temperature variations. A very effective way of integrating the VCO, and hence gaining numerous benefits, is by configuring it in a wide-bandwidth PLL. We have presented a simple nested-loop PLL architecture that achieves wide-BW while maintaining fine resolution and spur reduction. The expense is little in terms of extra components and power dissipation. A possible utilization of this idea in a transceiver architecture that will eliminate the overhead has been discussed. The potential of the idea was demonstrated by designing a completely integrated PLL that achieves the lowest phase-noise reported to date.

Chapter 5

PLLs using Closed-Loop Voltage-to-frequency Converters

5.1 Introduction

In the previous chapter we have presented the nested-loop PLL architecture, reproduced in Fig. 5.1. With two loops, one inside the other, the designer can optimize the BW of both loops to reject the reference spurs and the VCO phase-noise simultaneously. The expense involved in the nested-loop architecture is the extra IF VCO needed. While this VCO operates at a low frequency and should not represent a design bottleneck, it would be beneficial to try to eliminate it along with the associated external components.

A significant improvement can be made to the architecture of Fig. 5.1 by observing the nature of the signal at each point. The output of the first loop filter, $H(s)$, is a voltage signal that is converted to a frequency via VCO_1. This sig-

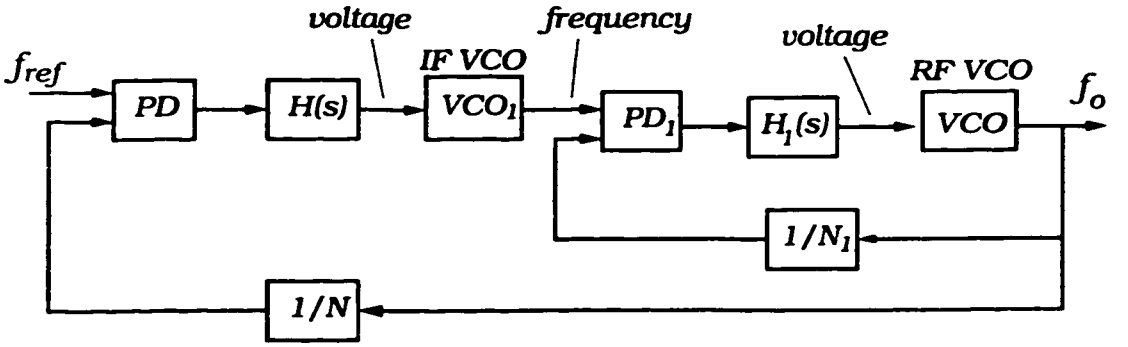

Figure 5.1: Dual-loop PLL.

nal is converted back to a voltage signal at the output of the second loop filter $H_1(s)$ via the phase-detector and filter. By making the inner loop a voltage-locked loop rather than a phase-locked loop, the IF VCO could be eliminated all-together. This is accomplished by inserting a frequency-to-voltage converter in the feedback path of the inner PLL and locking the loop to a voltage reference, creating a novel voltage-to-frequency converter.

5.2 Closed-Loop Voltage-to-Frequency Converter

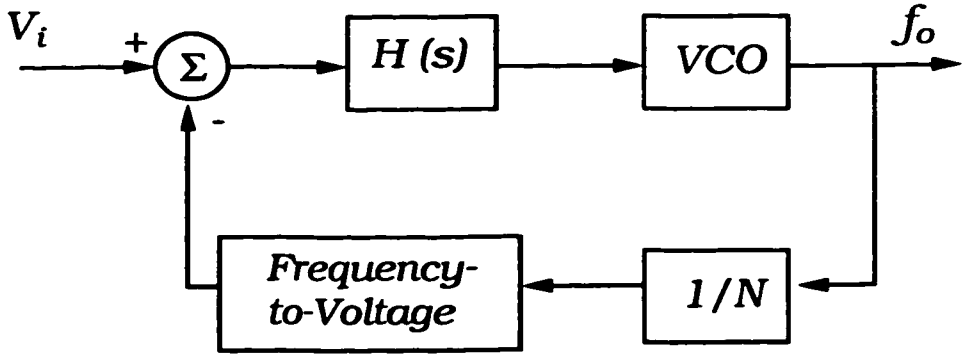

Figure 5.2: Closed-loop V/F architecture.

The proposed architecture is shown in Fig. 5.2 [74]. An ordinary VCO is configured in a feedback loop. The output of the VCO is first divided down to an intermediate frequency (e.g. 50MHz) where a frequency-to-voltage converter can be easily designed. The F/V produces a voltage output that is proportional to the input frequency. This output is subtracted from a reference voltage and the error signal is filtered by the loop filter, $H(s)$, which determines the bandwidth and spurious rejection. If the loop filter is chosen to have a pole at d.c. (integrator), then the steady-state error signal will be zero and the output of the F/V will be equal to the input reference. The output frequency will thus be equal to

$$f_o = K_{oeq} V_i = \frac{N}{K'} V_i \tag{5.1}$$

where K' is the gain of the F/V in V/Hz.

Hence, this architecture constitutes a voltage-to-frequency converter with an equivalent gain of K_{oeq} Hz/V. However, unlike an ordinary VCO which is an open loop architecture, this V/F is a closed-loop, voltage-mode feedback architecture, and will thus be referred to as a closed-loop V/F converter (CLV/F).

The closed-loop V/F architecture achieves the advantages of feedback systems, the most significant of which is the correction effect of the loop. S-domain analysis shows that this loop, similar to a PLL, will inhibit the VCO phase-noise within its bandwidth. Since the reference in our case is a voltage level, not a frequency signal, it will not result in any spurii at the output, and hence the bandwidth can be made very wide resulting in a much better VCO phase-noise suppression. This is achieved without the use of multiple VCOs as the case in multiple-loop PLLs.

The spurious in this architecture comes from the F/V output. According to the design of the F/V, there will be a certain amount of voltage spikes on its output. The frequency of these spikes will be equal to the input frequency of the F/V. Since this frequency could be relatively very large (e.g. 30-60 MHz), then even a high loop bandwidth (e.g. 2MHz) will still achieve adequate spurious suppression. In addition, since the spurious frequency is in the tens of MHz range away from the carrier (as opposed to the channel spacing in a classical PLL case), then the spurious rejection requirement might even be more relaxed. This is because channels at the spur frequency (which are feared to mix with the spur and alias on top of the desired channel), being at such large offsets, would have probably been attenuated by the receive-band filters preceding and/or following the LNA.

The frequency divider used in the closed-loop V/F is of fixed division ratio and, hence, does not need a prescaler. Thus this block is not expected to dissipate a large amount of power.

Another advantage of this architecture is the tunability of the gain of the V/F. By varying the gain of the F/V in the feedback path (through varying the bias current, for example), or the division ratio, the gain of the V/F gain be tuned. This might be useful in covering process and temperature tolerances or in other applications.

5.3 Low Phase-Noise PLL Architecture

While the closed-loop V/F is a feedback system, it would not be sufficient for frequency synthesis in most application. This is due to the fact that the linearity of

the relation between the output frequency and the reference voltage will be limited by the linearity of the frequency-to-voltage converter. In addition, temperature and process dependence of the F/V will cause the output frequency to drift with such variations.

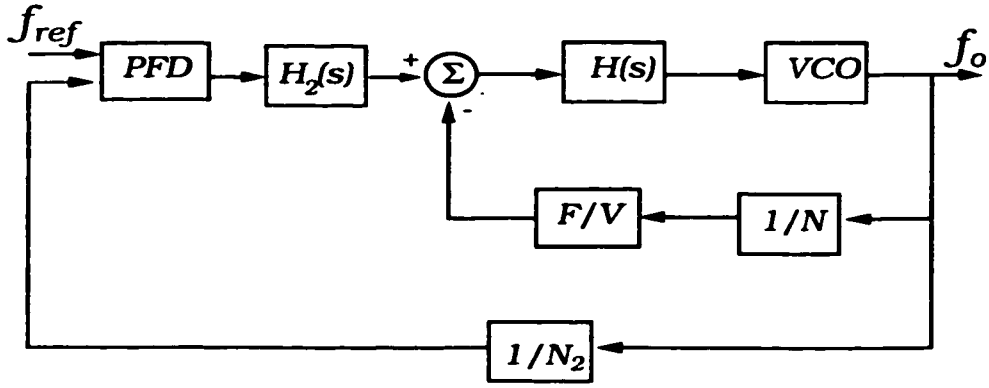

Figure 5.3: Low phase-noise PLL architecture.

In order to alleviate these drawbacks, the closed-loop V/F is further configured in a larger PLL as shown in Fig. 5.3. This architecture is essentially a classical PLL with the ordinary open-loop VCO replaced by the closed-loop V/F architecture. At steady state, the inputs of the phase/frequency detector (PFD) will be phase-locked and the output frequency will be equal to $N_2 f_{ref}$. The outer loop can be a classical PLL, a fractional-N PLL, a $\Sigma\Delta$ PLL or any other architecture.

In wireless applications, f_{ref} will be chosen to be equal to the channel spacing (or a multiple thereof in the case of a fractional-N PLL) and will be derived from an accurate, temperature-compensated crystal oscillator. Thus accurate frequency synthesis will be achieved, independent of temperature and process variation. For adequate spurious rejection and PFD phase-noise attenuation, the bandwidth of the overall loop will be kept one decade less than f_{ref}. The VCO phase-noise will

still be inhibited however within the wide bandwidth of the closed-loop V/F.

For the loop stability not to be affected, it is necessary that the inner CLV/F be much faster than the overall loop. This is in agreement with the purpose of the architecture, which is having a very large closed-loop V/F bandwidth, for good phase-noise rejection, and a small overall loop bandwidth for adequate reference spurious and phase-noise rejection. It can be shown that with the inner loop ten times faster than the overall one, the phase margin will degrade by less than 0.2°.

5.4 S-Domain Analysis

5.4.1 Closed-Loop V/F

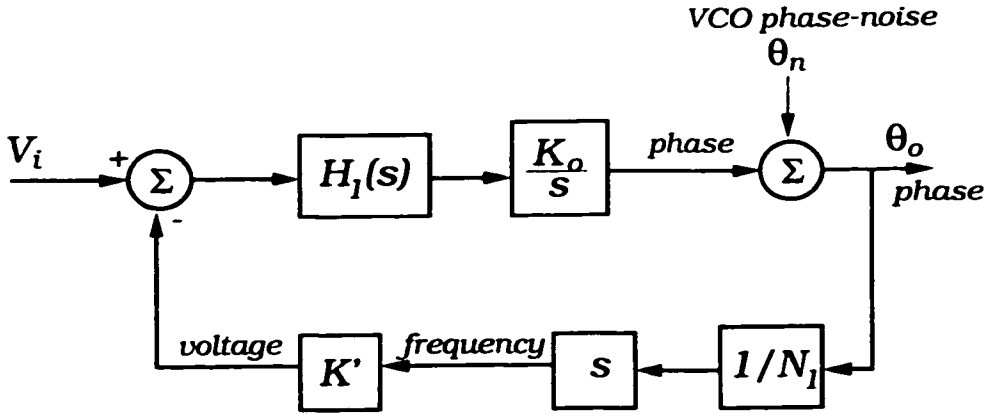

Figure 5.4: S-domain block diagram of the closed-loop V/F.

Fig. 5.4 shows the s-domain block diagram of the V/F with the nature of the signal indicated at each point. The F/V is modeled by its gain, K', assuming that its delay is negligible with respect to the loop bandwidth. The output of the loop is taken to be the VCO phase, hence the VCO is modeled by K_o/s. This allows for

proper comparison to an ordinary PLL, and because the VCO phase-noise is in the phase domain. The s-block in the feedback path translates the signal from the phase to the frequency domain (since it is the frequency of the signal that is sensed by the F/V).

The transfer function between the output phase and input reference voltage is given by

$$\theta_o = \frac{K_{oeq}}{s} V_i \cdot \frac{\frac{K_o}{K_{oeq}} H_1(s)}{1 + \frac{K_o}{K_{oeq}} H_1(s)} \tag{5.2}$$

where K_o is the VCO gain and K_{oeq} is the equivalent gain of the V/F as defined in (5.1). It is interesting to compare this result with that of an ordinary PLL. In a PLL, the input/output transfer function is similar to (5.2) but with each $H_1(s)$ replaced by $H_1(s)/s$. This leads to the fact that a PLL with a first-order loop filter will have a second-order transfer function. In our case, this extra order is missing. This is because of the s-block in the feedback path. Indeed, since we are sensing the frequency of the VCO output rather than its phase, one order of correction would be lost. However, we also note that the presence of the s-block in the feedback path *adds* 90° to the loop phase margin. This allows the order of $H_1(s)$ to be increased without affecting the loop stability. We thus can achieve the same level of correction as an ordinary PLL. For example, if the loop filter is chosen to be

$$H_1(s) = K_1 \frac{1 + s\tau_1}{s^2} \tag{5.3}$$

which is similar to a PLL loop filter but with an extra integrator, then the output

phase would be

$$\theta_o = \frac{K_{oeq}}{s} V_i \cdot \frac{2\zeta_1(s/\omega_{n1}) + 1}{(s/\omega_{n1})^2 + 2\zeta_1(s/\omega_{n1}) + 1} \tag{5.4}$$

where

$$\omega_{n1} = \sqrt{\frac{K_o K_1}{K_{oeq}}}, \quad \zeta_1 = \frac{\tau_1 \omega_{n1}}{2} \tag{5.5}$$

The relation between the output and the VCO phase-noise is given by

$$\theta_o = \theta_n \frac{(s/\omega_{n1})^2}{(s/\omega_{n1})^2 + 2\zeta_1(s/\omega_{n1}) + 1} \tag{5.6}$$

These relations are exactly similar to those of an ordinary PLL with a first-order loop filter [75]. Thus the closed-loop V/F converter attenuates the VCO phase-noise within the loop bandwidth by 40dB/decade, and attenuates the input noise beyond the loop bandwidth by 20dB/decade.

5.4.2 Overall PLL

Loop Stability

An s-domain block diagram of the overall PLL is shown in Fig. 5.5. The outer loop is assumed to have a classical phase-frequency detector and a first-order charge pump with a gain of $I_p/2\pi$ A/rad [36, 75]. Since the overall PLL consists of one loop inside the other, it is important to investigate the issue of stability. An identical problem has been considered in the case of the nested-loop PLL. It has been

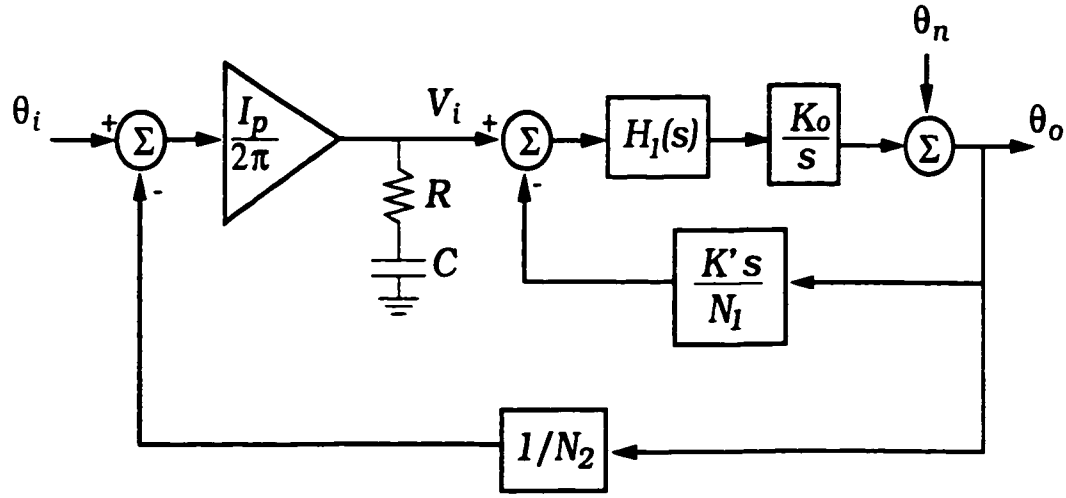

Figure 5.5: S-domain block diagram of the overall PLL.

demonstrated in section 4.3.2 that such an architecture remains stable as long as the inner loop is faster than the overall BW. If the ratio of the respective BWs is greater than 5-10, then the dominant poles remain very close to their nominal locations and the phase-margin of the overall loop is virtually unchanged. Hence, the designer can design the natural frequency and damping factor of the overall loop by assuming that the inner CLV/F is an ideal VCO with gain K_{oeq} Hz/V.

Transfer Characteristics

The input/output frequency response is shown in Fig. 5.6. ω_n and ω_{n1} are the natural frequencies of the overall loop and the CLV/F respectively. Input spurs or phase-noise at offsets greater than the PLL bandwidth are attenuated by 20dB/decade. At even larger offsets, the poles of the inner loop start to take effect. Thus noise at large offsets from the input (that are beyond the bandwidth of the inner loop) are attenuated by 40dB/decade.

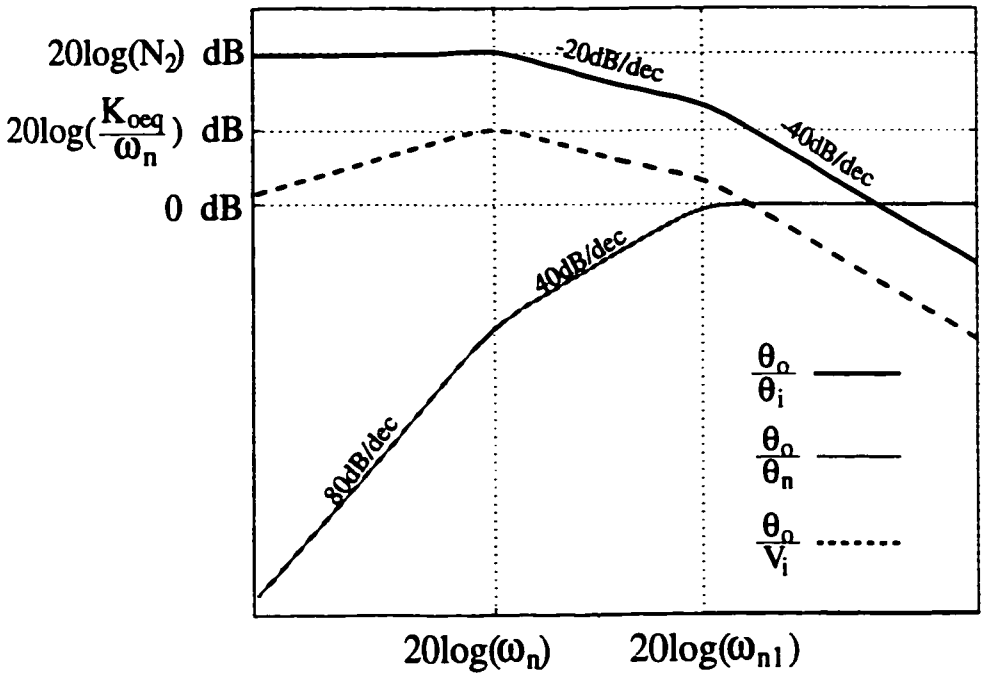

Figure 5.6: Bode plot of CLV/F transfer functions.

Similarly, the VCO phase-noise that is very close to the carrier (within the bandwidth of the overall PLL) gets an extra order of correction. As shown in Fig. 5.6, such noise is attenuated by 80dB/decade. This could be useful in reducing the $1/f$ and $1/f^2$ noise of the oscillator. VCO phase-noise that lies between ω_n and ω_{n1} is attenuated by 40dB/decade.

It is important to note the transfer function between the output phase and the input to the inner closed-loop V/F. If ω_{n1} is much larger than ω_n, this transfer function can be expressed as

$$\frac{\theta_o}{V_i} \approx \frac{K_{oeq}}{s} \cdot \frac{2\zeta_1(\frac{s}{\omega_{n1}}) + 1}{(\frac{s}{\omega_{n1}})^2 + 2\zeta_1(\frac{s}{\omega_{n1}}) + 1} \cdot \frac{(\frac{s}{\omega_n})^2}{(\frac{s}{\omega_n})^2 + 2\zeta(\frac{s}{\omega_n}) + 1} \tag{5.7}$$

Thus the overall PLL will inhibit any noise at the input of the closed-loop V/F that is within its bandwidth. Noise that is beyond the PLL bandwidth will appear at the output multiplied by $K_{ocq}/\Delta\omega$. However, noise that is at an offset greater than the inner BW will again be attenuated by an extra 40dB/decade. Hence, as discussed below, it is important to minimize the noise at the the input of the CLV/F that lies between the BWs of the two loops.

5.5 Design Considerations

As mentioned earlier, the wide BW of the inner closed-loop V/F attenuates the VCO phase-noise. Also, the narrow BW of the overall PLL attenuates the reference phase-noise and spur. However, the critical point in the proposed PLL is the input of the CLV/F. As discussed above, the noise at this point is amplified by the CLV/F gain and not attenuated in the range of interest. This dictates two important design considerations. First, the noise at that point should be minimized. Second, the equivalent gain of the CLV/F, K_{ocq}, should be as small as possible. Of course, K_{ocq} has to be large enough to cover the full tuning range over process and temperature variation. However, increasing it further will only amplify the noise at the input. These two considerations are similar to those of a classical narrow-BW PLL, where it is also important to minimize the gain of the VCO and the noise at its input.

The first consideration necessitates a low-noise design of the loop filter and F/V. In essence, the proposed frequency synthesizer architectures trades an off-chip, high-Q, bulky and power-consuming VCO for a low-noise opamp and frequency-to-voltage converter. As shown below, the design of such components should not

be an obstacle. The second consideration relates to the design of the F/V, as this element determines the equivalent gain.

5.5.1 Loop filter design

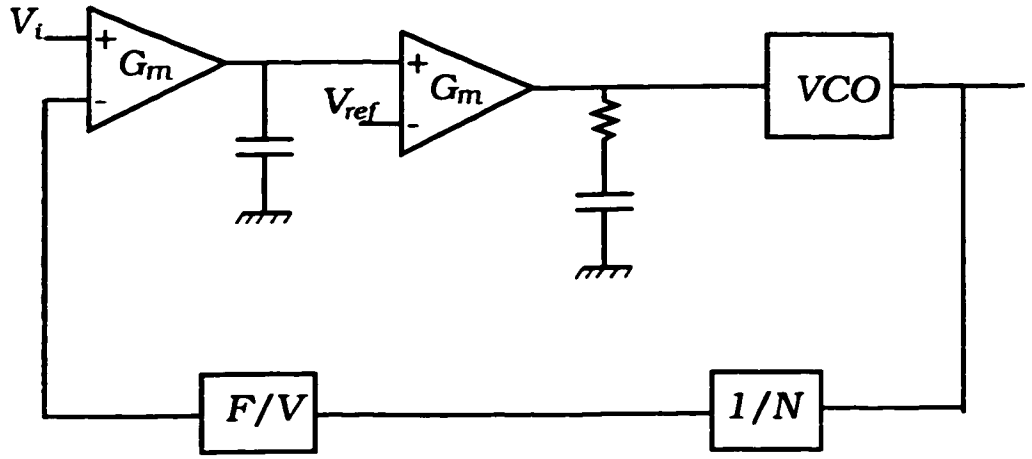

Figure 5.7: Closed-loop V/F with the second-order loop filter

A possible implementation of the loop filter given by (5.3) is shown in Fig. 5.7. The first opamp and capacitor form one integrator, and the remaining opamp and RC section form the other integrator and the zero. The proposed PLL architecture dictates some requirements on the opamps and offers some significant design simplifications.

As mentioned above, the input-referred noise level of the opamp must be minimized. This applies to the front opamp only, as any noise in the subsequent opamp will be attenuated by the gain of the first integrator. The total amount of noise that would be allowed at the opamp input is determined by the required phase-noise at the output. For example, the GSM standard has the relatively stringent

phase-noise specification of approximately -121dBc/Hz at an offset of 600kHz [56]. Assuming that the equivalent gain of the CLV/F is 50MHz/V, the noise at the opamp input must not exceed $10.5nV/\sqrt{Hz}$ at that frequency. Opamps implemented in ordinary CMOS processes can achieve better noise levels with very little power consumption [76–79]. The opamp in [76] is implemented in a CMOS process and achieves an input-referred noise level of $7nV/\sqrt{Hz}$ and gain of 120dB while dissipating 1.4mW. Opamps in bipolar technology can meet even more stringent requirements [80].

Another critical parameter of the front opamp is its input common-mode range, since it inversely affects the minimum equivalent gain of the CLV/F. In other words, if the input range of the front opamp is maximized, K_{oeq} can be made lower, hence reducing the noise at the output, while still covering the full tuning range over temperature and process variations. On the other hand, the output swing of the front opamp is not critical, because according to this configuration, the output will always settle at V_{ref}.

The offset of the opamp is not important since this only leads to a slight shift in the V-F curve. Since, the final value of the output frequency is determined by the overall loop, this is not a problem. The same argument holds for the linearity of the filter.

Careful observation reveals that the gain of the opamp and its slew rate are also not critical. The first opamp acts as an integrator, ideally with infinite gain at d.c. The limited gain of the opamp means that at very low frequencies, the opamp/capacitor pair will operate as an amplifier rather than an integrator. Specif-

ically, if the integrator time constant is τ_i, and the opamp d.c gain is A_o, then the response of the front opamp/capacitor pair will be that of Fig. 5.8a. This is referred to as integrator leakage [81].

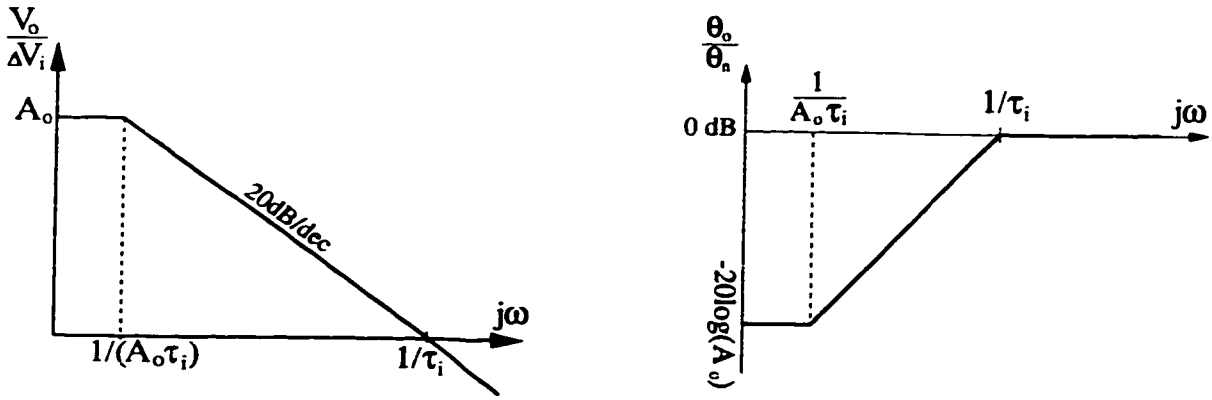

(a) Transfer function of integrator with limited d.c gain

(b) Effect on VCO phase-noise rejection

Figure 5.8: Integrator leakage effect.

The effect of integrator leakage on the rejection of the VCO phase-noise is shown in Fig. 5.8b. The maximum noise rejection is limited by A_o and occurs for frequencies less than $1/A_o\tau_i$. This means that the output phase-noise will start to slope upwards at this frequency. However, the integrator time constant, τ_i, will be approximately equal to the loop natural frequency ω_{n1}. Since the natural frequency is in the range of 2-3MHz, even a moderate gain of only 40dB will make the value of $1/A_o\tau_i$ well below the channel spacing and hence below the range of interest. Thus, a very high opamp d.c. gain, while beneficial, is usually not critical.

The slew rate limitation of the opamp is also not critical. When a sudden jump

occurs in the input voltage of the CLV/F, a slew-rate limited opamp will affect the dynamics of the loop and prolong the settling time. However, when the CLV/F is configured in the overall PLL, its input will not experience sudden changes and will be a slowly varying input due to the narrow BW of the overall loop. Thus the slew-rate of the opamps do not have to be high.

In conclusion, this specific architecture presents relatively relaxed constraints on the opamps with respect to the d.c gain, offset and slew-rate. This makes the design of a low-noise, large input-swing opamp with low power consumption easier.

5.5.2 Frequency-to-Voltage Converter Design

The F/V in the feedback path is the most critical component of the proposed closed-loop V/F since its determines the equivalent gain and the V-F characteristics. A major consideration of the design of the F/V is its gain. The gain of the F/V must be large in order to reduce the equivalent gain of the CLV/F (equation (5.1)) and thus the contribution of the input noise.

One possible F/V is shown in Fig. 5.9a. The input signal is XORed with a delay version of itself. The average d.c. value of the output will be proportional to the ratio of the fixed delay τ and the time period of the input. The output will thus be directly proportional to the input frequency. The problem with this F/V, apart from the large ripples that have to be filtered out, is that the gain is very limited: Assuming a full-swing XOR output, the d.c. value of the output is

$$\overline{V_o} = \frac{\tau}{T/2} V_{DD} = (2\tau V_{DD}) f_i \tag{5.8}$$

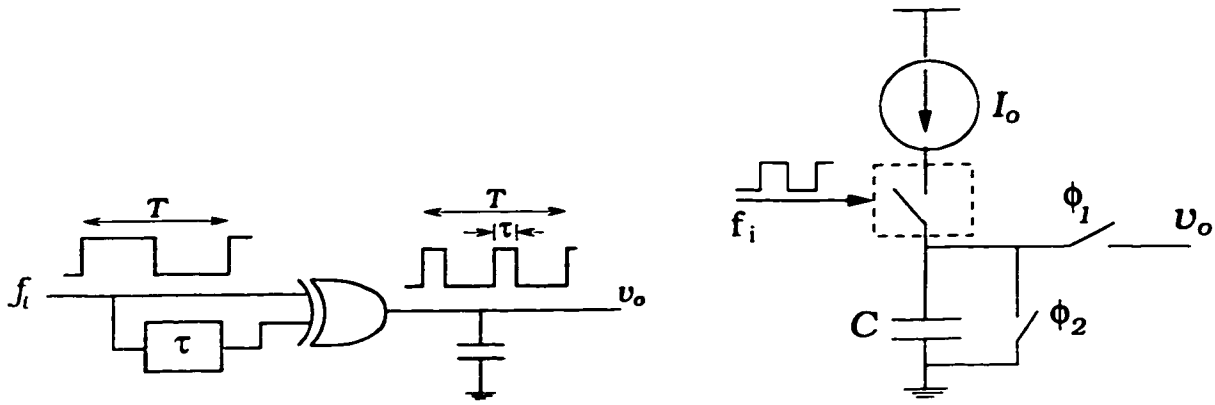

(a) (b)

Figure 5.9: Low-gain frequency-to-voltage converter architectures

However τ will always be less than half the input period, and the F/V gain is limited by

$$K' = 2\tau V_{DD} < \frac{V_{DD}}{f_i} \tag{5.9}$$

which is a very small gain. Since the input to the F/V is the output of the VCO divided by N (Fig. 5.2), then the equivalent gain of the CLV/F will be at least

$$K_{oeq} = \frac{N}{K'} > \frac{Nf_i}{V_{DD}} = \frac{f_o}{V_{DD}} \tag{5.10}$$

Hence, if the LO frequency is in the range of 1GHz, then for a 3-V supply, the equivalent gain will be more than 330MHz/V. This is a very large gain that will significantly amplify any noise at the input.

Another F/V architecture is shown in Fig. 5.9b. When the input signal is high,

a switch is closed and a capacitor, C, starts to integrate a constant current, I_o. The duration of integration is controlled by the period of the input signal, and thus the final voltage on the capacitor will be inversely proportional to the input frequency. This will result in a non-linear relation, but this is not important as discussed earlier. When the capacitor is charged, switch ϕ_1 is closed to transfer the voltage to the output, and then switch ϕ_2 discharges the capacitor before the next rising edge of the input.

This architecture also suffers from a limited maximum gain. Assuming the input signal has a 50% duty cycle, the output voltage will be equal to

$$V_o = \frac{I_o}{C} \cdot \frac{1}{2f_i} \tag{5.11}$$

and the F/V gain will be

$$K' = \frac{\partial V_o}{\partial f_i} = \frac{-I_o}{2Cf_i^2} \tag{5.12}$$

Since the capacitor voltage cannot exceed V_{DD}, the I_o/C ratio, which controls the gain, is limited by

$$\frac{I_o}{C} \cdot \frac{1}{2f_{i_{min}}} < V_{DD} \tag{5.13}$$

where $f_{i_{min}}$ is the minimum input frequency to the F/V. From (5.1), (5.12) and

(5.13), the CLV/F equivalent gain will thus be at least

$$|K_{oeq}| > \frac{2Nf_i^2}{2f_{i_{min}}V_{DD}} = \frac{f_o^2}{f_{o_{min}}V_{DD}} \tag{5.14}$$

which is an even higher value than the XOR F/V. This circuit, as it stands, is thus inadequate as K_{oeq} will be excessively high.

The high equivalent gain is caused by the fact that we start integrating the current on the rising edge of the input signal. However, in wireless systems, the center frequency of the synthesizer is much larger than the tuning range and the period of the input signal varies slightly around a nominal value. Hence, we only need to start integration towards the end of the HIGH half-cycle, allowing for a higher I_o/C ratio without saturating the output.

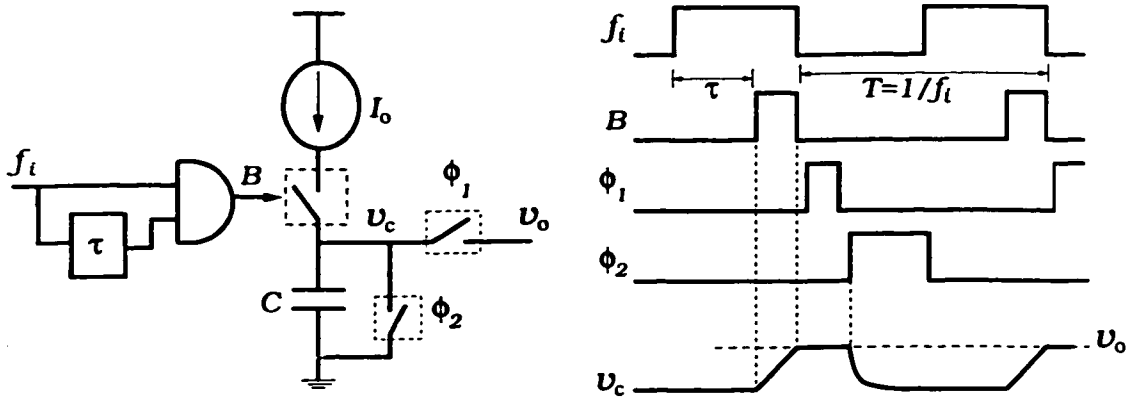

Figure 5.10: Frequency-to-voltage converter.

The circuit implementing this, along with the timing diagrams, are shown in Fig. 5.10. The AND/delay-block combination delays the integration by a *fixed*

duration τ. The output voltage will be

$$V_o = \frac{I_o}{C} \cdot (\frac{1}{2f_i} - \tau) \tag{5.15}$$

and the gain K' will again be given by (5.12), but with a larger I_o/C ratio. The delay τ will determine the frequency of operation. If the output voltage ranges from 0 to V_{DD}, then

$$\begin{aligned} \frac{I_o}{C}(\frac{1}{2f_{i_{max}}} - \tau) &= 0 \\ \frac{I_o}{C}(\frac{1}{2f_{i_{min}}} - \tau) &= V_{DD} \end{aligned} \tag{5.16}$$

and hence,

$$\tau = \frac{1}{2f_{i_{max}}} \tag{5.17}$$

and

$$\begin{aligned} \frac{I_o}{C} &= \frac{2V_{DD}}{\frac{1}{f_{i,min}} - \frac{1}{f_{i,max}}} \\ &\approx \frac{2f_i^2 V_{DD}}{\Delta f_i} \end{aligned} \tag{5.18}$$

where f_i and Δf_i in (5.18) are the center frequency and tuning range at the F/V input respectively. As shown, the value of the delay τ determines the maximum frequency, while the I_o/C ratio determines the tuning range. Obviously, the tuning range should be designed to have enough margin to account for process and temperature variations in values of the current, capacitor and delay.

From (5.12) and (5.18), the equivalent gain of the closed-loop F/V will be

$$K_{ocq} = \frac{N}{K'} = \frac{2NCf_i^2}{I_o} \approx \frac{\Delta f_o}{V_{DD}} \tag{5.19}$$

where Δf_o is the tuning range at the output of the VCO. Unlike (5.10) and (5.14), K_{ocq} in this case is equal to the ratio of the required tuning range to the available voltage headroom. This yields a significantly lower equivalent gain. Using this architecture, equivalent gains of as low as $40MHz/V$ were achieved.

F/V Noise Analysis

It is very important to analyze the noise contribution of the F/V. The constant current transistor in Fig. 5.10 will usually be implemented using a transistor, possibly with a degeneration resistor. The noise associated with this current source is going to be integrated on the capacitor and then sampled by the output at the end of the integration period. Between the capacitor terminals, the low-frequency components of the noise are enhanced due to the integration process. However, the capacitor is discharged every cycle. Also, the variations of the capacitor voltage are masked from the output voltage that only samples the final value at the end of the integration period. This final value remains a random value with a Gaussian distribution (if the noise source is white noise) and hence has a constant spectrum over frequency. As the following analysis shows, the transfer function between the current-source noise and the output is not of an integrator, but rather of an equivalent resistor.

Assume that the noise associated with the current source is designated by i_n,

where $\overline{i_n^2}$ represents the total noise power or the variance of the Gaussian process (if the noise is white). During the integration period, the capacitor voltage is equal to

$$v_c(t) = \frac{1}{C}\int_0^t i_n dt \tag{5.20}$$

and the value of the output voltage at the end of the first integration cycle would be

$$v_{n1} = \frac{1}{C}\int_0^{T_{int}} i_n dt \tag{5.21}$$

where T_{int} is the duration of integration. From Fig. 5.10

$$T_{int} = \frac{1}{2f_i} - \tau \tag{5.22}$$

By dividing T_{int} into infinitesimal simulation steps with a time step of t_s, the output voltage can be written as

$$v_{n1} = \frac{t_s}{C}\sum_{p=1}^{N_s} i_{n_p}\ , \qquad t_s \to 0 \tag{5.23}$$

where i_{n_p} is the noise current at the p^{th} sample and N_s is the total number of samples in T_{int}, i.e.

$$N_s = \frac{T_{int}}{t_s} \tag{5.24}$$

After the first integration cycle, the capacitor is discharged, and the next value of the output would be given by another value similar to (5.23). Over a large number of cycles, the mean square value of the output noise would be

$$\overline{v_n^2} = \overline{v_{n1}^2 + v_{n2}^2 + \ldots} = \frac{t_s^2}{C^2} \cdot \overline{\left(\sum_{p=1}^{N_s} i_{n_p} \right)^2}_{1,2,\ldots} \tag{5.25}$$

But since the noise i_n is uncorrelated, the power-of-two in the righthand side of (5.25) can be entered into the summation, yielding

$$\begin{aligned} \overline{v_n^2} &= \frac{t_s^2}{C^2} \cdot \overline{\left(\sum_{p=1}^{N_s} i_{n_p}^2 \right)}_{1,2,\ldots} \\ &= \frac{t_s^2}{C^2} \cdot N_s \overline{i_n^2} \end{aligned} \tag{5.26}$$

where the second step in (5.26) follows from the definition of the r.m.s. value.

Equation (5.26) gives the total noise power at the output. However, the output samples-and-holds these values at a rate of f_i (the input frequency to the F/V). Hence, this noise power is concentrated in frequencies from $0 \rightarrow f_i/2$ only, and the power spectral density of the noise at the output is

$$PSD_{v_n} = \frac{N_s t_s^2}{C^2} \cdot \frac{\overline{i_n^2}}{f_i/2} \quad V^2/Hz \tag{5.27}$$

From (5.24) and rearranging,

$$\begin{aligned} PSD_{v_n} &= \frac{T_{int}}{f_i C^2} \cdot \frac{\overline{i_n^2}}{1/(2t_s)} \\ &= \frac{T_{int}}{f_i C^2} \cdot PSD_{i_n} \end{aligned} \tag{5.28}$$

Hence, the PSD of the output noise voltage is equal to that of the input noise current multiplied by a noiseless equivalent resistance of value

$$R_{eq}^2 = \frac{T_{int}}{f_i C^2} \tag{5.29}$$

This analysis can be easily verified using Matlab. From (5.22) and (5.17) we can re-write (5.29) as

$$R_{eq} = \frac{4.44N}{\omega_o C}\sqrt{\frac{\omega_{o_{max}} - \omega_o}{\omega_{o_{max}}}} \tag{5.30}$$

where ω_o and $\omega_{o_{max}}$ are the RF frequencies at the VCO output. This shows that as the operating frequency approaches the maximum frequency (limited by $1/\tau$), the noise decreases, because the integration period gets shorter.

In order to achieve low-noise operation we must not only reduce the current-source noise (PSD_{i_n} in (5.28)), but also R_{eq}. This can be done by increasing the integrating capacitor, C. To maintain the same F/V gain, the current I_o will have to be proportionally increased. Hence, as in the case in many situations, noise reduction involves a trade-off with power consumption.

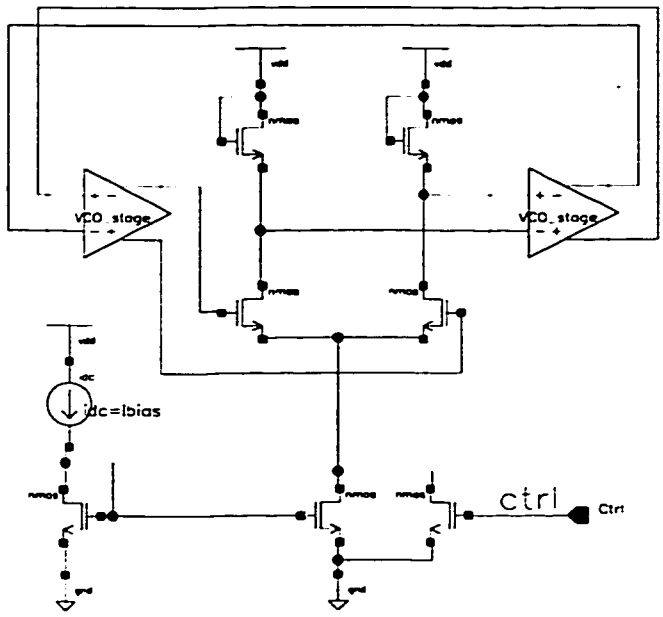

Figure 5.11: Three-stage ring oscillator VCO

5.6 Circuit design

The architecture of Fig. 5.7 was designed in a readily available 12GHz, 0.8μm BiCMOS technology. A simple three-stage ring oscillator shown in Fig. 5.11, was used for the VCO. Even though ring oscillators exhibit poor phase-noise, the implementation demonstrates the improvement that could be achieved by using the proposed architecture. The VCO center frequency was 1.05GHz, with a gain of approximately 150MHz/V. By varying the bias current a very wide tuning range can be achieved. The nominal power consumption of the VCO was 11mW.

Four stages of the D-flip flop shown in Fig. 5.12 were used to implement a divide-by-16 ripple-counter frequency divider. ECL-mode logic was used with BJT

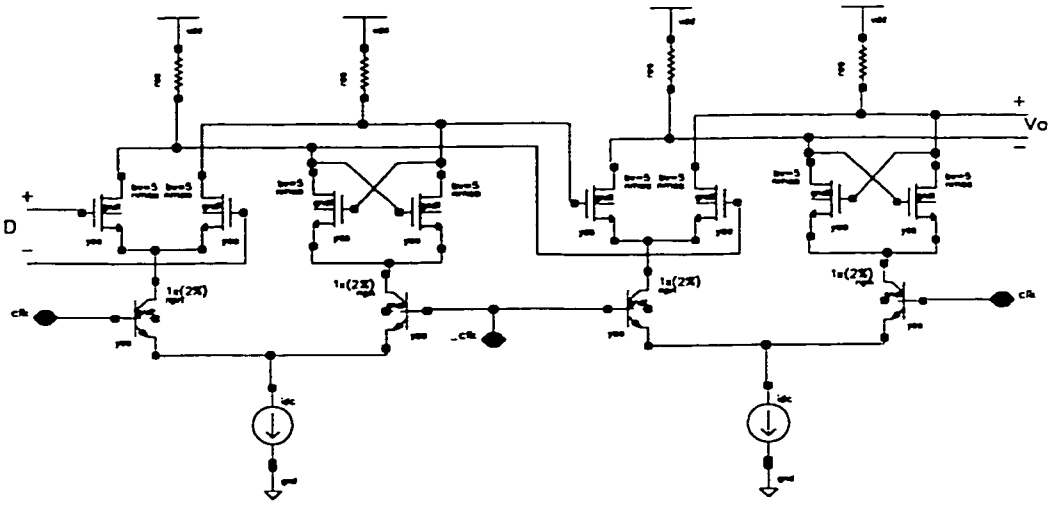

Figure 5.12: D-FF used to implement a divide-by-16 counter

as the driver elements. This allows for reduced swing and less noise injection. A differential-to-single-ended stage at the end of the divider produces a rail-to-rail signal for the subsequent clock generation block.

The clock generation circuit is shown in Fig. 5.13. The delay element is implemented using three inverters with an added capacitance load of 0.4pF on each inverter to increase the delay. The simulated delay τ was 7.5ns. The two lower branches shown in the figure generate clocks ϕ_2 and ϕ_3. ϕ_2 is wide enough to allow the transfer of charge from the main capacitor to the output capacitor, while ϕ_3 is wide enough to fully discharge the capacitor. The connections shown along with the values of the delay capacitors ensure that there is no overlap between the clocks. The simulated outputs of the clock generator are shown in Fig. 5.14.

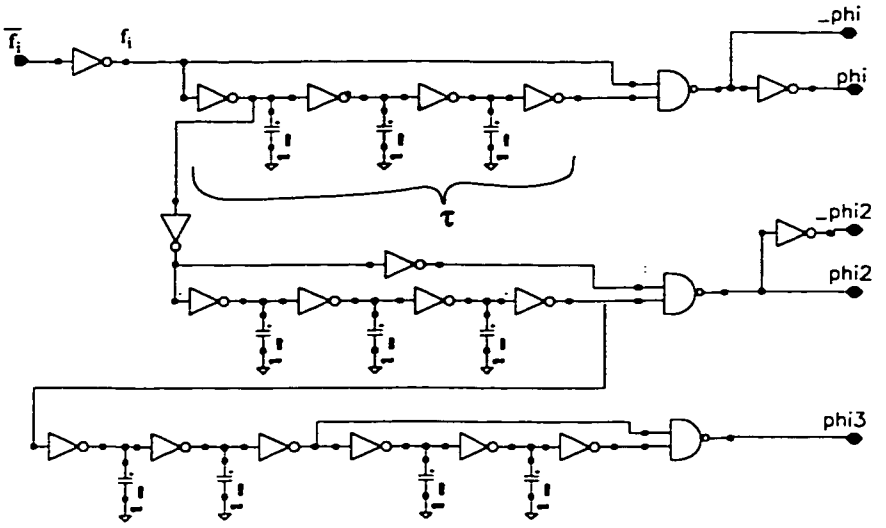

Figure 5.13: Schematic of the clock generation circuit.

The schematic of the frequency-to-voltage converter is shown in Fig. 5.15. A BJT is used as the constant current source and the capacitor is connected to V_{DD}. This results in a positive-slope closed-loop V/F. The start/stop of integration is controlled by switching the bias current. A resistance, R, is added in shunt with the base of the BJTs to ensure sharp fall time. Transistor M1 samples the capacitor voltage to the output at the end of the integration cycle, while transistor M2 is a dummy transistor that reduces charge injection and clock-feedthrough. PMOS transistor M3 discharges the capacitor to V_{DD}. The simulated F/V output is also shown in Fig. 5.14. The I_o/C ratio was set to 1.5V/ns yielding an equivalent gain of the CLV/F of approximately 40MHz/V.

The opamps used to implement the loop filter were folded-cascode opamps with gain of more than 50dB and input range of more than 3V. The input referred noise-

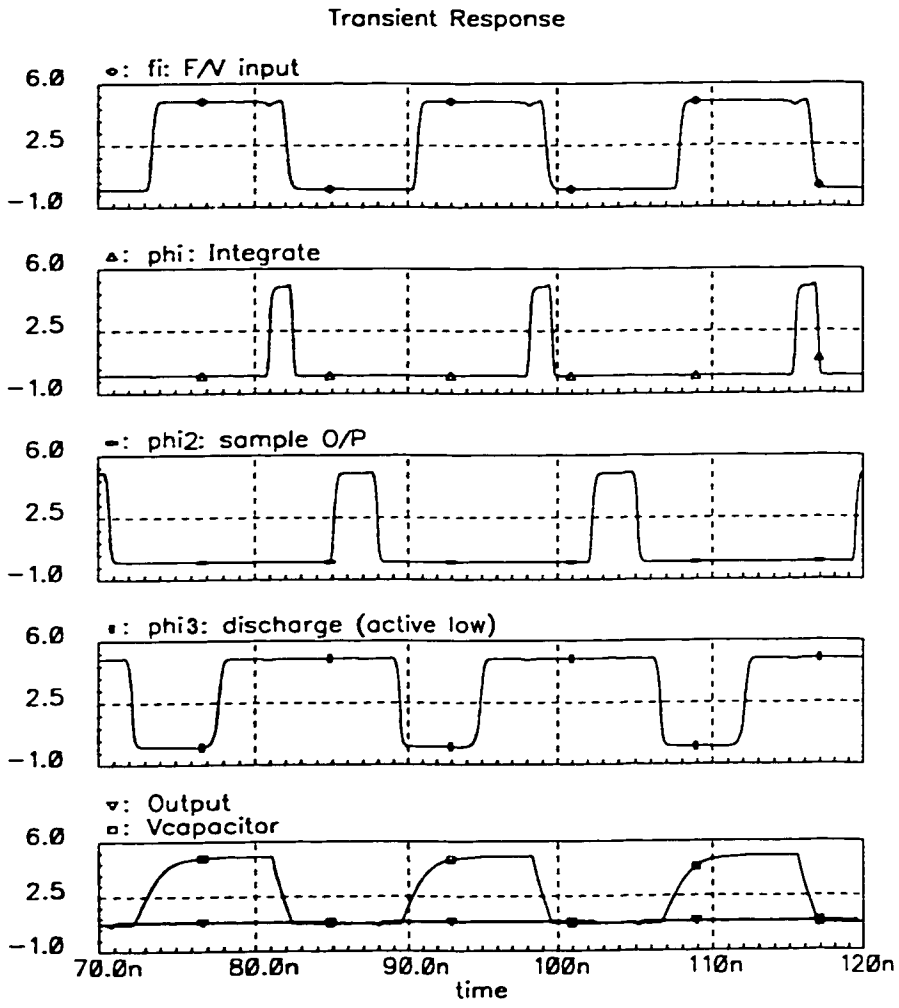

Figure 5.14: Hspice simulations of the clock generator and F/V outputs.

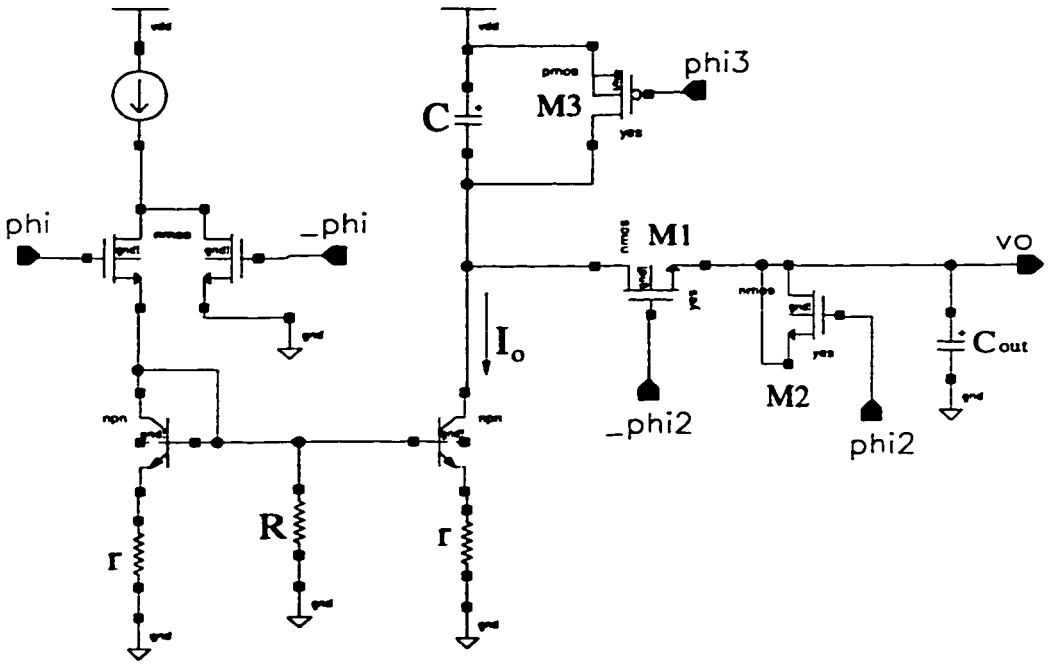

Figure 5.15: Schematic of the frequency-to-voltage converter.

floor of the opamp was approximately 13nV. The passive components of the loop filter were off-chip.

5.7 Simulations

Using Simulink from Matlab, system-level simulations were performed on the proposed architecture. In order to verify the s-domain analysis, the transient response of the system was compared to that of an ideal second-order system. Fig. 5.16 shows the transient response of the input of the VCO for a frequency step of 110MHz. The natural frequency and damping factor were set at 2.8MHz and 0.7 respectively. Also shown is the response of an ideal second-order system with similar parameters. As show, the two responses are in close agreement and exhibit the same settling time. The CLV/F has a higher overshoot due to delays in the loop that were not accounted for (like the F/V delay), which reduce the loop phase margin.

In order to demonstrate the loop rejection of the VCO phase-noise, three equal tones, 26dB below the carrier, were injected in the VCO phase at offsets of 6MHz, 11MHz and 30MHz. The loop bandwidth was increased to 12.5MHz. Fig. 5.17 shows the PSD of the VCO output when configured in the loop. As shown, the tone within the loop BW was attenuated by more than 10dB while the tone at 30 MHz was not attenuated. The tone at 11MHz (close to the cutoff frequency) was only slightly attenuated.

Circuit-level simulations were also performed on the CLV/F. Hspice was used to simulate the step response of the complete loop. As mentioned in section 5.6, the loop was designed with an equivalent gain of 40MHz/V. Fig. 5.18 shows the

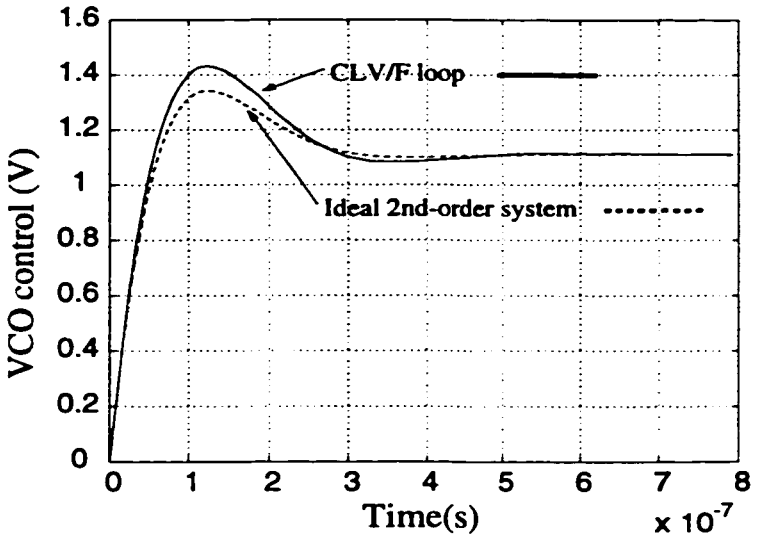

Figure 5.16: Matlab simulations of the transient response of the CLV/F loop.

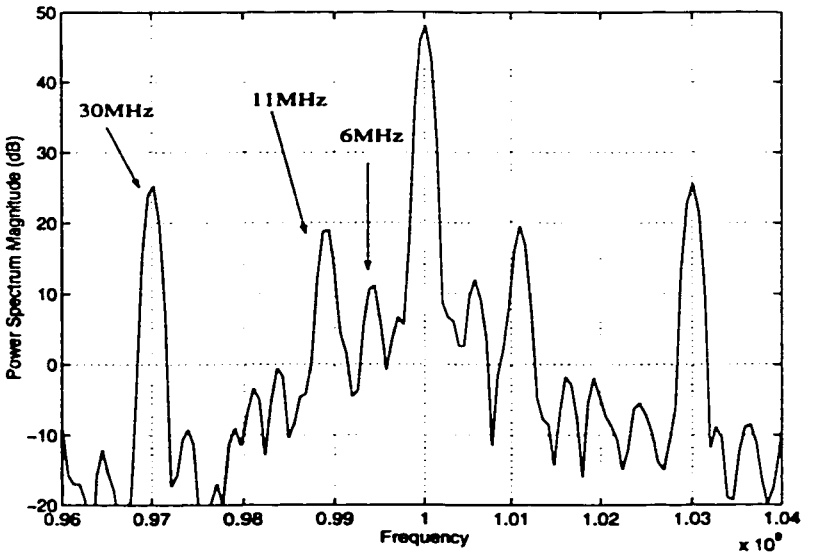

Figure 5.17: Matlab simulations of the loop rejection of the VCO phase-noise.

transient response of the loop, with the natural frequency and damping factor set at 2.6MHz and 0.7 respectively. The input voltage was stepped from 0.8V to 1.4V at 0.6μs, resulting in a frequency step of 24MHz. The upper portion of the figure shows the F/V output tracking the input voltage. The spikes on the F/V output are due to asymmetry in the clocks and their complements, creating some clock-feedthrough. The F/V output exhibits the predicted oscillating frequency of 2.6MHz. However, it requires a slightly longer time for settling as compared to Fig. 5.16 due to slew-rate limitations at the beginning of the step. The lower part of the Fig. 5.18 shows the VCO control.

5.8 Experimental results

The designed circuit was fabricated in the 0.8μm BiCMOS technology and tested. The measured voltage-frequency characteristics of the chip is shown in Fig. 5.19. The slope of the curve is in good agreement with the simulations. The center frequency of operation is approximately 8.5% higher than the simulated value due to process and temperature tolerance.

The phase-noise measurements are shown in Fig. 5.20. Since, the purpose of this prototype was to demonstrate the phase-noise improvement that could be achieved with the proposed architecture, no effort was made to optimize the phase-noise of the ring VCO used. Because of this and because of the proximity of noisy digital circuits, the phase-noise of the free-running VCO is relatively poor as shown in the figure (dotted line). However, when configured in the closed-loop V/F architecture, with a bandwidth of approximately 800kHz, the phase-noise of the output (solid-

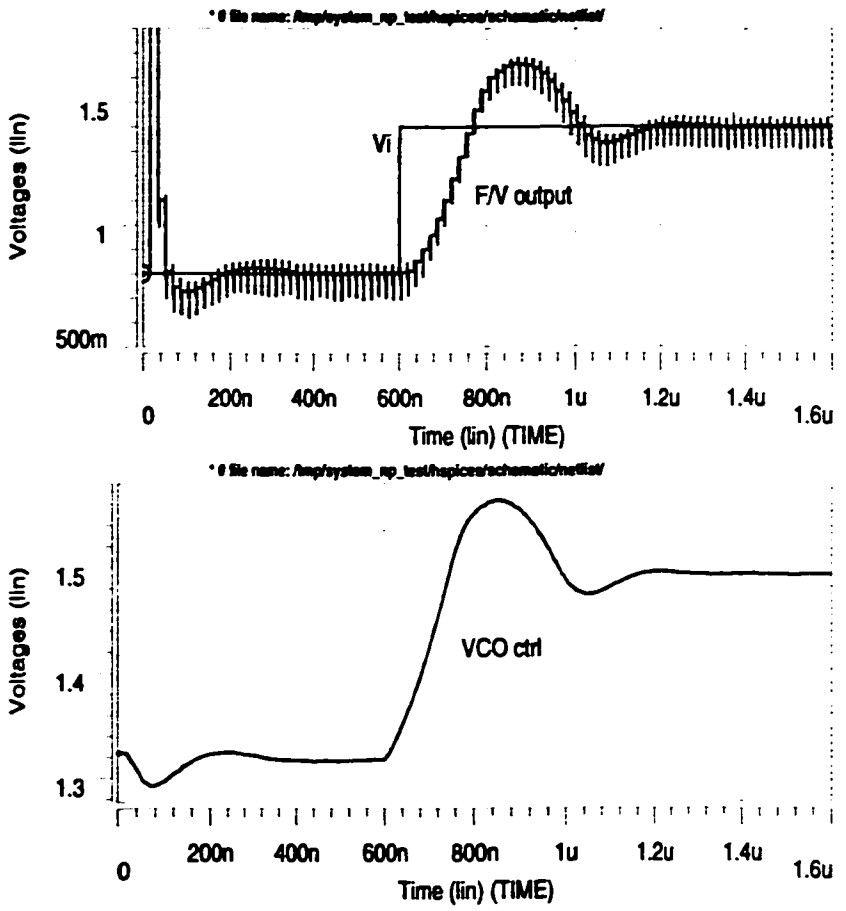

Figure 5.18: Hspice simulations of the CLV/F for a frequency step of 24MHz.

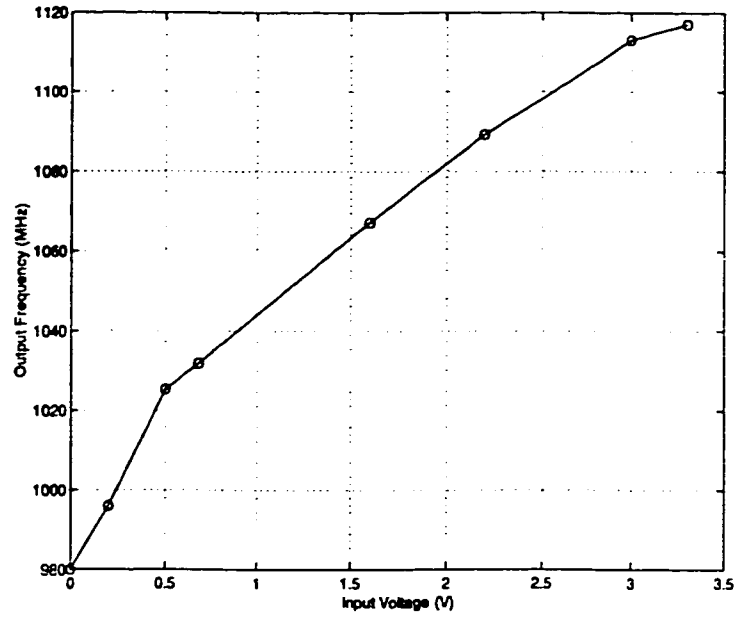

Figure 5.19: V-F characteristics

line) is improved by up to 20dB. As shown, the curve is approximately flat for frequencies less than the loop bandwidth. The phase-noise starts to slope upwards around 2kHz where the input noise starts to dominate and/or the integrator leakage effect starts to appear.

We do note that had an LC-VCO with a lower free-running phase-noise level been used, then the improvement would not necessarily have been as much, since the F/V and opamp noise might start to dominate. In that case, a lower noise level at the input might be necessary.

The spur levels were 35dBc below the carrier with little dependence on the loop filter. This indicates that the main source of the spurs was the noise from the digital circuits (which operate at that frequency) coupling through the substrate. Unless

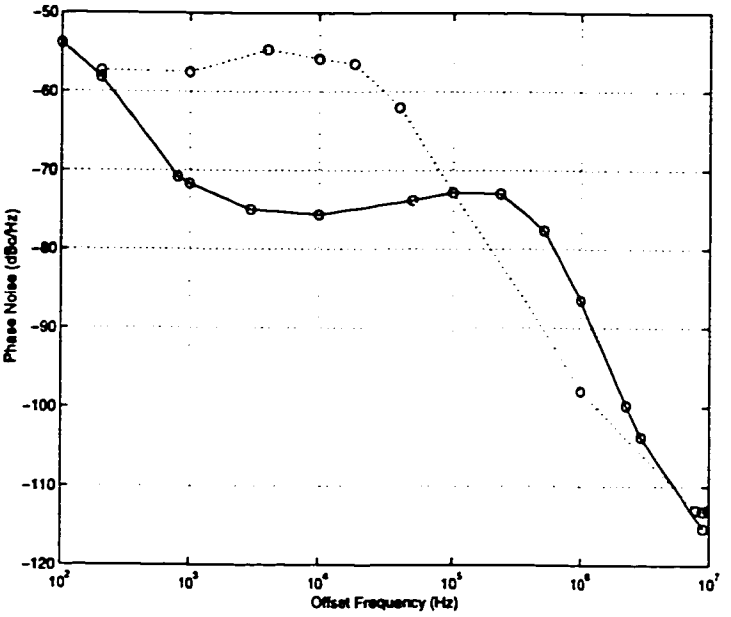

Figure 5.20: Phase-noise performance of the free-running VCO (dotted line) and the closed-loop V/F architecture (solid-line).

the digital and analog parts are integrated on separate chips, a lower level of spurs could only be attained through better layout and isolation. The chip core consumes approximately 14.2mA from a 5V supply in addition to 9mA for the output buffer. A micro-photograph of the chip is shown in Fig. 5.21.

5.9 Conclusion

In this chapter we have presented a novel voltage-locked loop that constitutes a means of closed-loop voltage-to-frequency conversion. By making the loop BW very large, 20dB reduction in the phase-noise of an integrated ring oscillator was

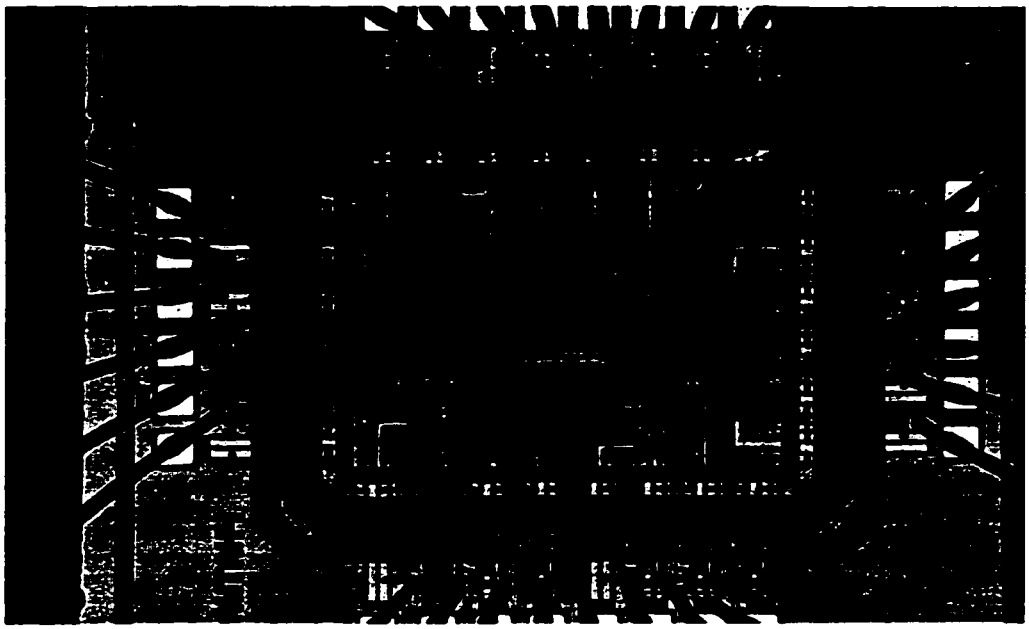

Figure 5.21: CLV/F Chip Micrograph.

achieved. When this architecture is configured in a phase-locked loop, completely integrated frequency synthesizers could be achieved.

The fundamental trade-off involved in the CLV/F architecture is between the VCO phase-noise and the noise at the input. Hence, for the architecture to achieve its goal, a low-noise F/V with a high gain should be used. The architecture will thus be useful where the noise contribution by the F/V and opamp is within the phase-noise specifications.

Chapter 6

Subsampling in Frequency Synthesis

Subsampling has been proposed before as a means of downconverting the RF band to d.c. by using a low-frequency LO signal [24, 82]. The drawback of such a subsampling mixer is the high noise-figure due to aliasing of noise. In this section we extend the idea of subsampling to frequency synthesis [83] and show that the problem of noise aliasing will not affect the performance in this case.

6.1 Introduction

A explained in chapter 3, the main problem with the classical PLL is that the reference frequency is equal to the channel spacing. This leads to two major consequences. First, the division ratio, N, will be very large. As explained in section 3.3.2, any phase-noise at the input (due to the input signal or the phase-detector)

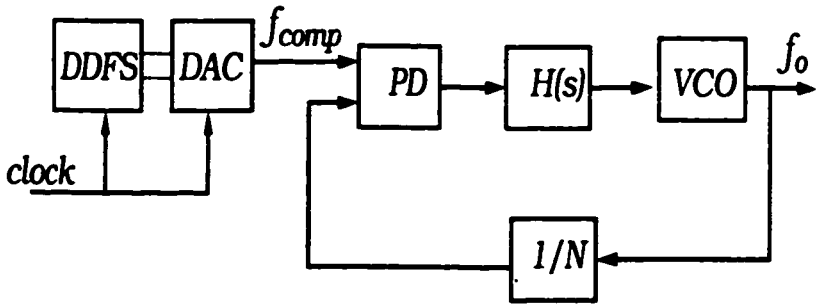

Figure 6.1: DDFS-driven PLL

as well as any spur at the charge-pump output, will be multiplied by N^2 when referred to the output. The very large division ratios in classical PLLs result in the input phase-noise limiting the frequency synthesizer performance.

Second, the loop bandwidth is generally an order of magnitude below the reference frequency [40]. With the reference equal to the channel spacing, the loop BW can be very small leading to an overly slow loop. In addition, the narrow BW means that the loop does not inhibit the phase-noise of the VCO in the range of interest. This means that a high-Q, external VCO must be used to meet the phase-noise requirements.

Fractional-N PLLs do not solve the problem. While moderately reducing the division ratio, a spur at the channel spacing persists, preventing any increase in the BW and hence accelerating the settling time or reducing the VCO phase-noise. $\Sigma\Delta$ PLLs involve a tradeoff between BW and power consumption and dissipate a large level of power.

An attractive method of frequency synthesis is using a direct digital frequency synthesizer (DDFS) to generate the signal and then upconverting the signal to RF via a PLL, as shown in Fig. 4.1 and reproduced in Fig. 6.1 [42, 84]. Since

the reference frequency is high, the division ratio, N, is reduced and hence the effect of the phase-detector noise and the charge pump spur. With a high reference frequency, the BW of the loop can be increased hence accelerating the settling time and attenuating the VCO phase-noise.

While this architecture is promising, the limiting problem here is the dilemma of spurious content and operating frequency. The DDS with the DAC can obtain a good spurious performance, (e.g -70dBc, -80dBc narrowband), provided that the output frequency is low (e.g. 5MHz-10MHz) [61, 62]. At this frequency however, the division ratio will be around 100-200. Since any DAC spurs will be multiplied by N^2 when referred to the output, this results in -40dBc to -30dBc spurs at the output. These levels are not acceptable [42]. If the division ratio is reduced on the other hand, the DDS must operate at a higher frequency and the DAC spurs start to increase significantly, resulting in even larger spurs at the output. In addition, a higher operating frequency, means larger power dissipation in the DDS and DAC.

Another problem with the DDS/PLL architecture in Fig. 6.1 is the large number of bits needed in the DDFS. Since the output frequency of the DDS is f_o/N, the DDS clock will be around $3f_o/N$. The required frequency resolution is f_{CH}/N. Thus the number of bits of the accumulator in the DDS will be $log_2(\frac{3f_o/N}{f_{CH}/N})$ which is independent of the division ratio. This translates to around 14-bits for GSM and 18-bits for Mobitex. The large number of bits leads to more power consumption.

A technique to reduce the division ratio while maintaining low operating frequency of the DDFS is to downconvert the VCO signal mostly through mixing, rather than division, as shown in Fig. 6.2. The obvious problem with this method

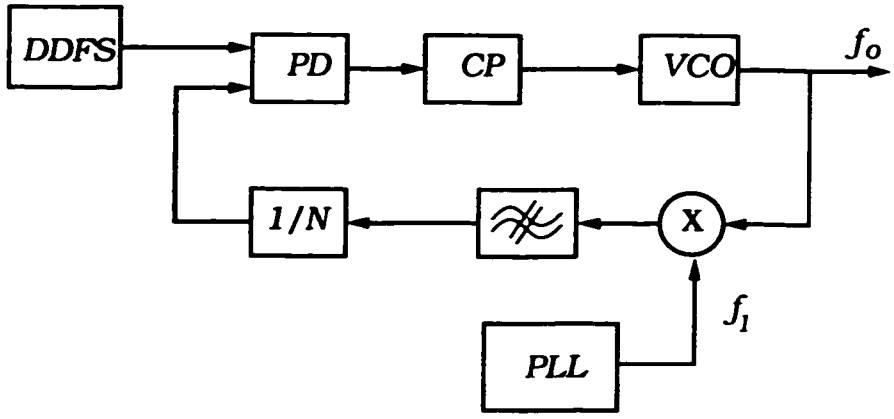

Figure 6.2: Hybrid DDS/double-loop PLL

however, is the need for an additional PLL, VCO and mixer which makes this solution expensive [42]. In addition, the noise in the secondary PLL and VCO will add to the output.

In this chapter we propose a novel architecture that alleviates these problems by employing sub-sampling in the feedback path. This architecture does not need extra PLLs/VCOs or mixers and does not degrade the output phase-noise. In the following sections we outline the architecture and list its benefits. The design considerations pertaining to subsampling and harmonics are discussed. System-level simulations of the architecture are then presented.

6.2 Subsampled PLL

The proposed architecture is shown in Fig. 6.3, where the frequency band and specific values were geared towards the Mobitex standard. The output of the VCO is sampled-and-held at a frequency, f_s, which is low, hence the name subsampling.

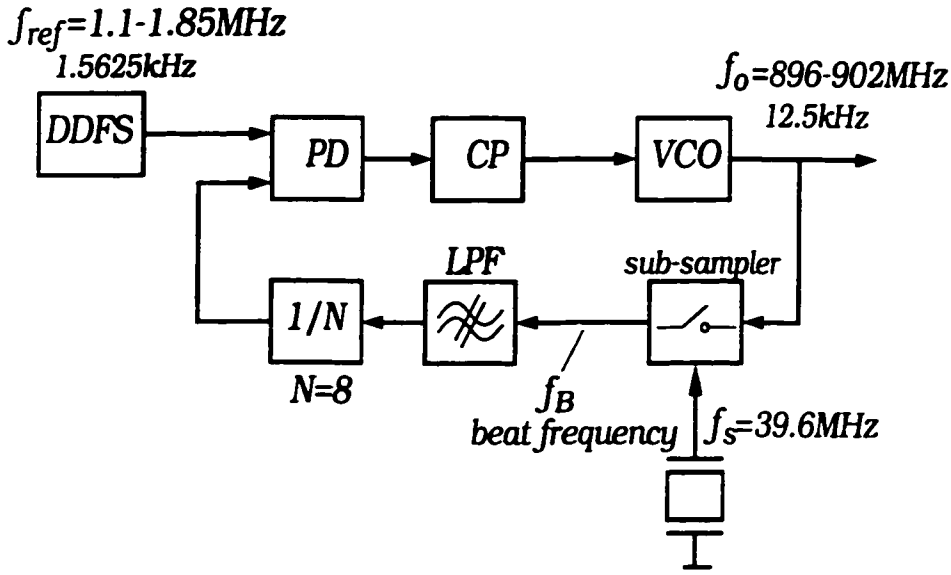

Figure 6.3: PLL using subsampling.

Since sampling creates replicas of the input spectrum spaced at the sampling frequency, the fundamental frequency of the output of the sampler would be the beat frequency between f_o and f_s. This is defined as the difference in frequency between the sampling frequency, f_s, and the nearest subharmonic of the input frequency, f_o. In a mathematical form, the beat frequency is given by

$$f_B = |f_o - N_s f_s| \tag{6.1}$$

where we define N_s as the subsampling ratio which is equal to the ratio f_o/f_s rounded to the nearest integer. For example if f_o is 900MHz and f_s is 39.6MHz as in the figure, then the subsampling ratio, N_s, will be 23 and the beat frequency 10.8MHz.

The output of the sampler is filtered in order to remove the harmonics created

by sampling. The LPF output is then optionally divided by a small, fixed division ratio, N. This allows the DDFS which provides the reference to the phase-detector to operate at a very low frequency thus greatly simplifying its design. In the above architecture, the sampling frequency f_s is fixed and tuning is achieved by tuning the DDFS frequency. While the DDFS is particularly suitable, any alternative frequency source can be used to provide the reference signal.

The main idea in this architecture is the use of subsampling in the feedback path. By using subsampling, the signal f_s will be of a fixed and low frequency, and thus can be easily obtained from the system crystal oscillator. Hence the downconversion of the VCO signal is accomplished with a small division ratio without the need for additional PLLs or VCOs, etc. In addition, since f_s is obtained from a crystal oscillator, its phase-noise is very low and will not affect the output.

6.3 Advantages

This architecture has several advantages. The most important advantage is the great reduction in the division ratio. The division ratio in the above example is 8, as opposed to 72,000 if a classical, integer-N PLL was used. The leads to a 79dB ($20log(72,000/8)$) reduction in the phase-noise contribution of the phase-detector, making it negligible.

The lower division ratio also means that the effect of the DAC spurs is greatly lowered. For example, if subsampling was not utilized, the division ratio would have to be around 500 for the DDFS/DAC to operate at the same frequency as Fig. 6.3. Our architecture uses a division ratio that is 62 times smaller, resulting in

36dB reduction in the effect of the DAC spurs. This allows for a much lower DAC resolution and/or power consumption and makes the hybrid DDS/PLL feasible.

Another advantage of the subsampled-PLL is the reduction in the required DDFS phase-resolution. Since N in Fig. 6.3 is 8, the resolution of the DDFS must be 1.5625kHz to yield a channel spacing of 12.5kHz at the output. At a clock rate of approximately 6MHz, only a 12-bit accumulator will be sufficient. In a classical DDS-driven PLL, a DDS operating at the same clock will have to have an 18-bit accumulator. This helps reduce the size and power consumption of the DDS.

An important benefit of subsampling is that the reference frequency is high with respect to other PLL architectures. A reference of 1-2MHz as that shown leads to a wide loop BW. The wide loop BW results in a fast settling time and helps attenuate the VCO phase-noise. This may help in integrating the VCO.

Finally, in our architecture, there is no high-speed prescaler. This offsets the power needed for the sampler and LPF.

All of these benefits are achieved by using a fixed- and low-frequency signal from the system crystal. Hence, no extra hardware, or associated power consumption is needed to implement our architecture. Since f_s is obtained from a crystal oscillator, it does not degrade the phase-noise performance.

6.4 Design Considerations

6.4.1 Effect of Subsampling on Noise

Noise in Sampling Signal

The effect of subsampling on the noise in the sampling signal (f_s) is similar to the effect of frequency multiplication. In other words, any noise in the sampling signal will be increased by $20log(N_s)$ at the output, where N_s is the subsampling ratio. This can be explained as follows. If the VCO signal is given by $Acos(2\pi f_o t)$, the value of a sample at t_s will ideally be

$$V_s = Acos(2\pi f_o t_s) \tag{6.2}$$

If the signal f_s contains phase-noise, the sampling instant will exhibit jitter in the time domain. If a jitter of magnitude Δt occurs, then the output of the sampler will now become

$$V_c = Acos(2\pi f_o(t_s + \Delta t)) \tag{6.3}$$

Hence, a shift of Δt in the sampling instant is exactly equivalent to a shift of Δt in the VCO signal. However, since the time period of the sampling signal is N_s times larger than the VCO frequency, the equivalent phase-shift is N_s times larger. In other words, the shift Δt represents a $2\pi f_s \Delta t$ phase shift in the sampling signal, but is equivalent to $2\pi f_o \Delta t$ in the VCO signal. Hence the phase-noise of f_s will be increased by $20log(N_s)$ when mapped to the output.

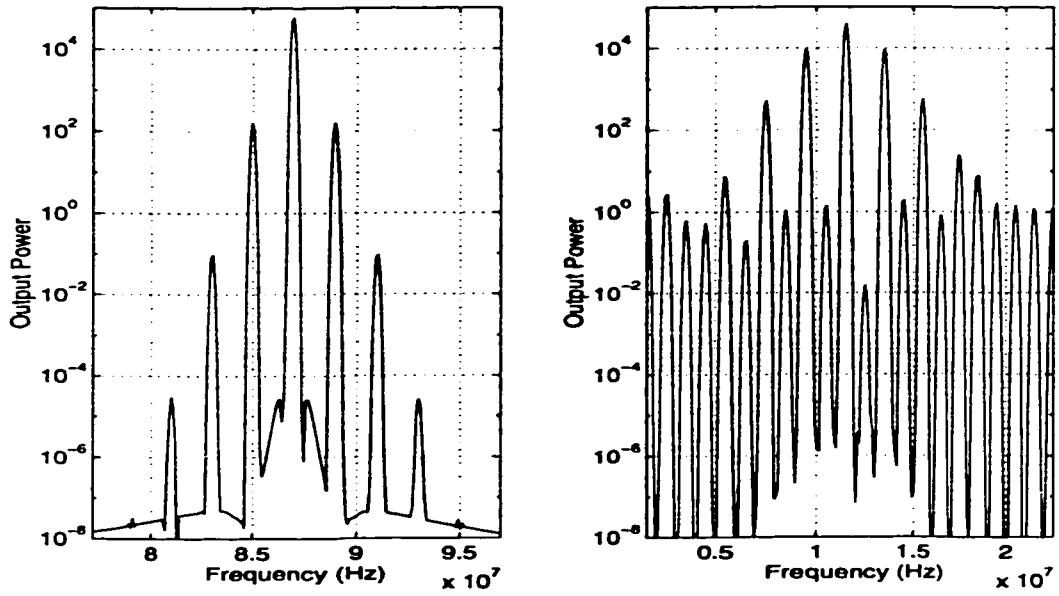

Figure 6.4: Effect of subsampling on noise in sampling signal.

This effect is demonstrated by Matlab simulations in Fig. 6.4. A free-running VCO running at 881.5MHz is sampled by an 87MHz signal, resulting in an subsampling of $N_s = 10$. The spectrum of the sampling signal, f_s, is shown on the left. A tone of -26dBc was injected at an offset of 2MHz in f_s to simulate its noise. The output of the sampler is shown on the right. The fundamental is at the beat frequency of 11.5MHz. As shown, the tone at the 2MHz offset has increased by 20dB which corresponds to $20log(N_s)$.

The value of N_s will usually be in the 10-20 range. This is a relatively small value as compared to division ratios in most PLLs. Nonetheless, the above analysis demonstrates the importance of having a 'clean' f_s signal. Hence in our architecture, f_s is derived from the crystal oscillator. The phase-noise level of such signals is

extremely low and even when increased by 20-30dB will not affect the output phase-noise.

We note here that the above analysis holds for phase-noise that is within the loop BW. Any noise that is at offsets larger than the loop BW will be attenuated by the loop response. Hence the noise-floor of the sampling signal will not affect the output.

VCO Noise

Another important consideration when dealing with subsampling is the VCO noise aliasing on top of itself, raising the noise floor. This is demonstrated in Fig. 6.5, which shows the output of a sampler with a sampling frequency f_s and beat frequency f_B. While the fundamental lies at f_B, sampling will create replicas of the signal at spacings of f_s. In addition negative-frequency components alias to the positive-frequency. Hence we get tones at $f_s \pm f_B$, $2f_s \pm f_B$, and so on (only the first image shown in figure). It is clear from the figure that the noise floors of those tones will alias on top of each other. This results in increasing the noise floor by the ratio of the sampler's input BW to the sampling frequency [82], approximately 16dB for the example of Fig. 6.3. Due to this fact, the noise figures of subsampling mixers are prohibitively high.

A close look at Fig. 6.5 shows that in our case, this is not a problem. In the region around the carrier, the phase-noise (e.g. -95dBc/Hz at 10kHz) is much larger than the noise floor (e.g. -150dBc). Even when the noise floor is increased by 15dB, it will still be far below the VCO noise at the range of interest. As the VCO phase-noise starts to decrease at higher offsets, the noise-floor aliasing starts

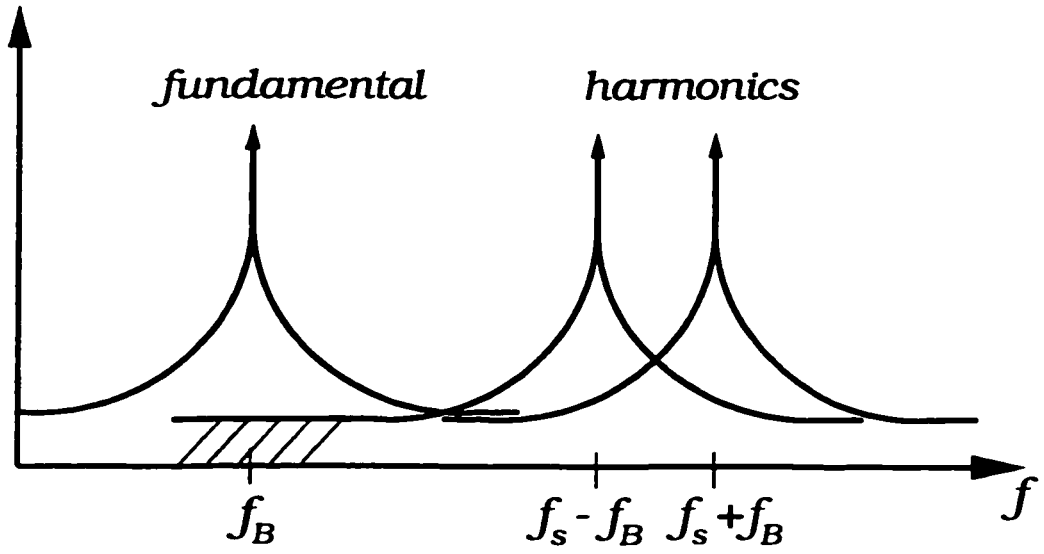

Figure 6.5: Noise folding due to sampling.

to become significant and affect the output noise level. However, this occurs at offsets much larger than the loop BW, and will be rejected by the loop low-pass filtering.

Matlab simulations, shown in Fig. 6.6, further illustrate this effect. Here the input of the VCO, which is running at 1GHz, is modulated with a noise source creating the phase-noise indicated by the dotted line. The noise floor of the VCO was set to -130dBc. The output of the VCO is sampled at 92MHz. The solid line shows the phase-noise of the sampler output. As shown, noise aliasing has no effect on the phase-noise up until 10MHz, far beyond the loop BW.

6.4.2 Harmonics and Spurs

The spurs created by the sampling process must be studied, as they influence the frequency planning and the requirements on the subsequent LPF. A diagram of a

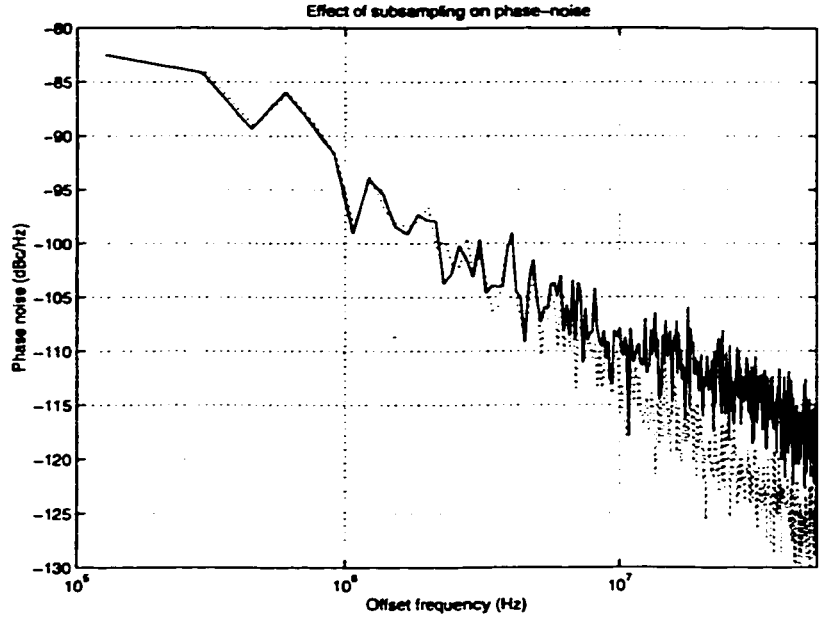

Figure 6.6: Matlab simulations of VCO noise aliasing. Dotted line represents the phase-noise of the VCO while solid line shows the phase-noise of the sampler output.

typical output of the sampler is shown in Fig. 6.7.

The VCO signal beats with the sampling signal, f_s, creating the fundamental at f_B and and the predicted sampling harmonics at $f_s \pm f_B$, $2f_s \pm f_B$ and so on. If the VCO contains a large second and third harmonics, then the additional spurs shown in Fig. 6.7 will also appear. The second harmonic of the VCO will beat with f_s creating a tone at $2f_B$ and images at $f_s \pm 2f_B$. The third harmonic will result in $3f_B$ and $f_s \pm 3f_B$, an so on. Due to the sample-and-hold function of the sampler, the output is filtered by the $sin(x)/x$ function.

The tones at $2f_B$, $3f_B$, $4f_B$, ...etc, are exact multiples of the fundamental frequency and do not cause any jitter in the zero-crossing. Hence, they do not need to be filtered out. However, the tones at $f_s \pm nf_B$ (the dashed ones in the figure)

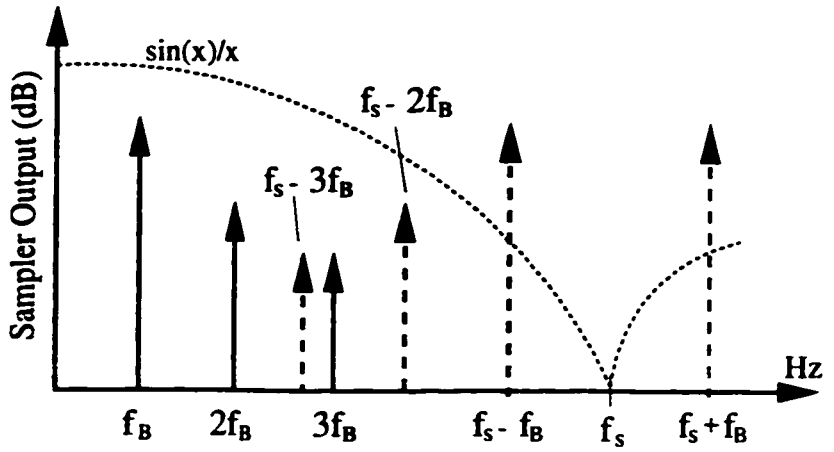

Figure 6.7: Spurs and harmonics at the sampler output.

are not a multiple of the fundamental and can cause zero-crossing jitter and hence output spurs.

The spur problem can be overcome by a combination of different remedies. The LPF following the sampler should be sharp enough to reject the unwanted spurs, while the VCO output could be filtered to reduce its harmonics. Careful frequency planning can also be exercised so that the undesirable spurs lie outside the LPF band and/or the loop cutoff frequency.

It is also beneficial to have a large sampling frequency as this will move the spurs to higher offsets and relax the requirements on the LPF. From Nyquist criterion, the minimum value of f_s is twice the desired tuning range at the output. The value of f_s will need to be larger than that, however, for adequate filtering of the spurs. The value of f_s will be limited by the available crystal, though. This is a drawback of the subsampled PLL. If the tuning range at the output is very large, then an overtone crystal has to be used. Alternatively, frequency doublers can be used to multiply the crystal frequency. However this will dissipate power and can

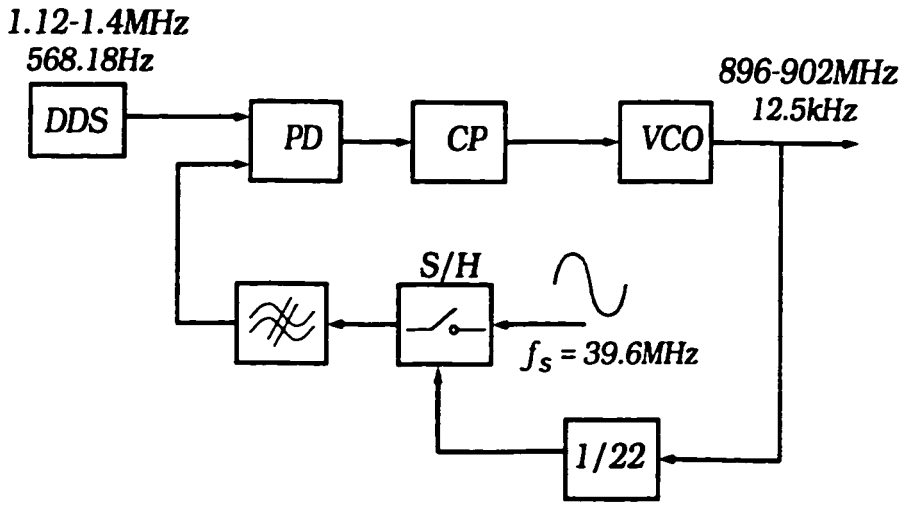

Figure 6.8: Alternative Sampled PLL.

contribute to the noise.

We also note that the LPF that follows the sampler operates at a very low frequency, where a sharp active filter is feasible and will not consume a lot of power.

6.5 Alternative Sampling Architecture

Fig. 6.8 shows an alternative sampled-PLL architecture. Here the VCO frequency is divided down first, and then is mixed with the fixed reference signal via a sampling mixer. Since the output of the frequency divider is a square-wave, it must act as the clock to the sampler, whose input is the sinusoidal signal derived from the crystal. Since the frequencies of the sampler inputs are close, this is not sub-sampling anymore.

The advantage of this modification is that the harmonics generated by sampling are much further from the beat frequency. This greatly relives the spur problem and reduces the necessary order of the subsequent LPF. In fact, the PLL low-pass response might be sufficient in rejecting the spurs. The price paid for this is that the division ratio seen by the DDFS (which is operating at about the same frequency as Fig. 6.3) increases from 8 to 22. Thus there is a trade-off between the LPF requirements and the DDFS spurious requirement. In addition, a high-frequency counter is needed to divide the VCO signal.

6.6 Simulations

The system of Fig. 6.3, with the shown frequency bands, was simulated using Simulink from Matlab. The VCO gain was 10MHz/V, the natural frequency 75kHz and the damping factor 0.707. The LPF used was a 4^{th}-order elliptic LPF with a cutoff frequency of 15MHz. For a 3MHz step, the VCO control is shown in Fig. 6.9. This is approximately identical to the ideal second-order response depicted by the dashed line in the figure.

The spectrum of the VCO output is shown in Fig. 6.10. The output frequency was 896MHz giving a beat frequency of 14.8MHz, which is the worst case separation from the image tone. The simulation step was 12ps and 4M simulation points used for FFT. As shown, the spurs are more than 65dBc below the carrier.

The effect of the LPF order was also examined. When a 3^{rd}-order LPF was used, the spurs increase by about 10dB as shown in Fig. 6.11. Slight improvement can be observed in Fig. 6.12 where a 6^{th}-order LPF was used. For higher LPF

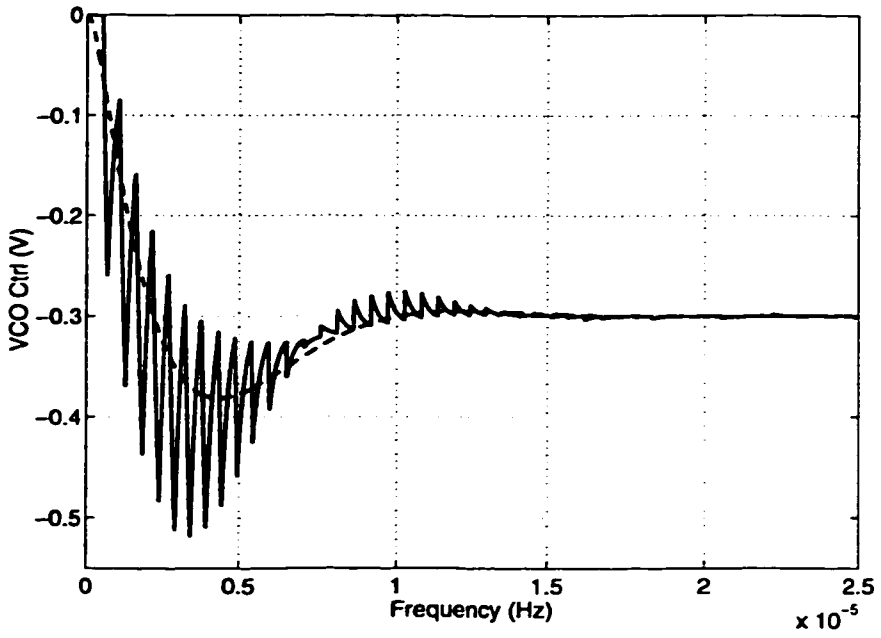

Figure 6.9: Subsampled PLL transient response (solid) vs. ideal second-order system (dashed).

orders, the spurs levels start to be dominated by the simulator accuracy.

The suppression of the VCO phase-noise was also simulated. For a bandwidth of 75kHz, it is necessary to use an excessively large number of points for FFT in order to get a sufficient resolution. Thus the loop parameters where changed in order to increase the loop bandwidth. The VCO output was around 881.5MHz, sampled at 87MHz. The LPF was a 4th-order filter with a cutoff frequency of 32MHz. The bandwidth was set to 1.8MHz. Three equal tones were injected at the output of the VCO at offsets of 0.9, 1.8 and 5MHz. Fig. 6.13 shows the VCO output when configured in the loop. As shown the tone within the loop bandwidth was attenuated by approximately 17dB, while the one outside the bandwidth was not attenuated. As expected, the tone at 1.8MHz was attenuated by only 3dB.

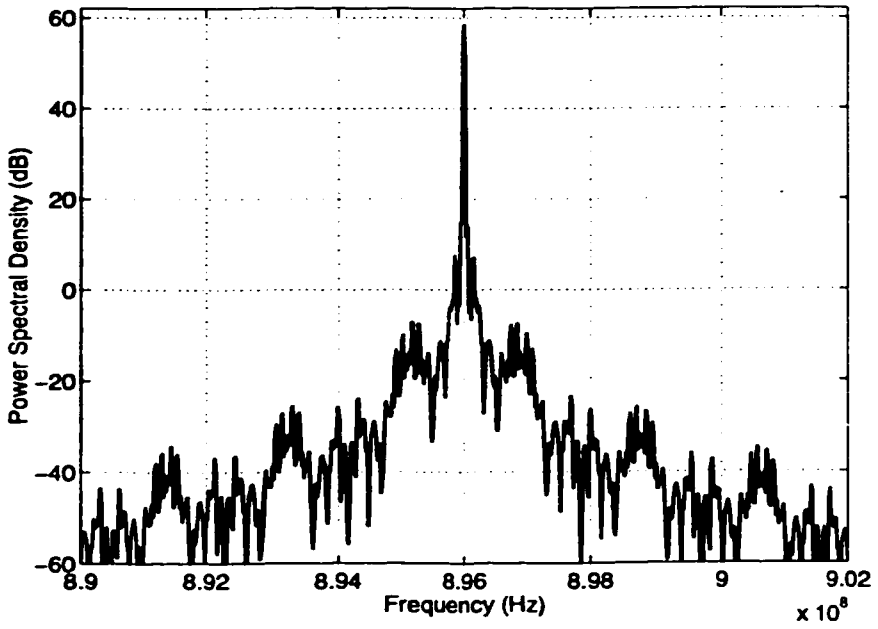

Figure 6.10: PLL output spectrum for a 4^{th}-order LPF.

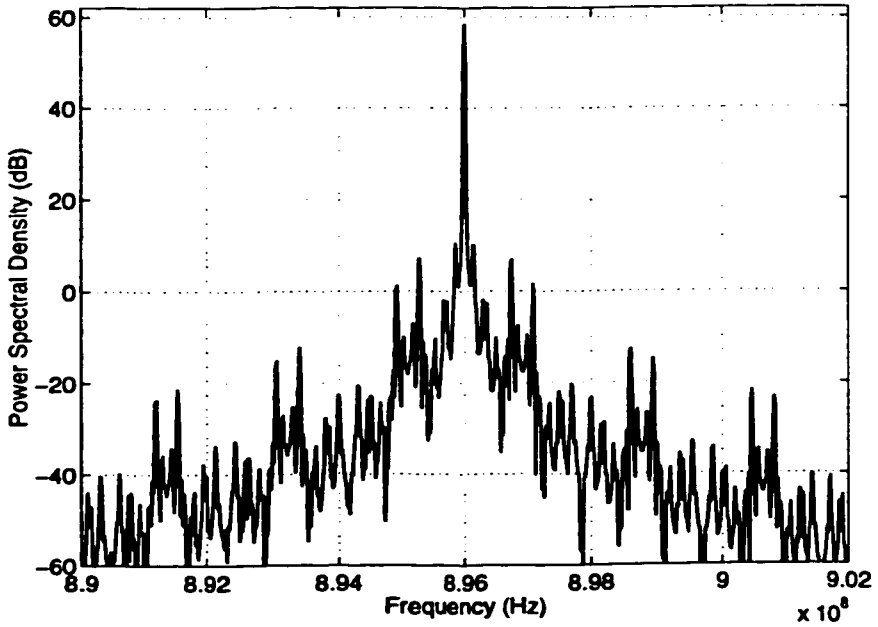

Figure 6.11: PLL output spectrum for a 3^{rd}-order LPF.

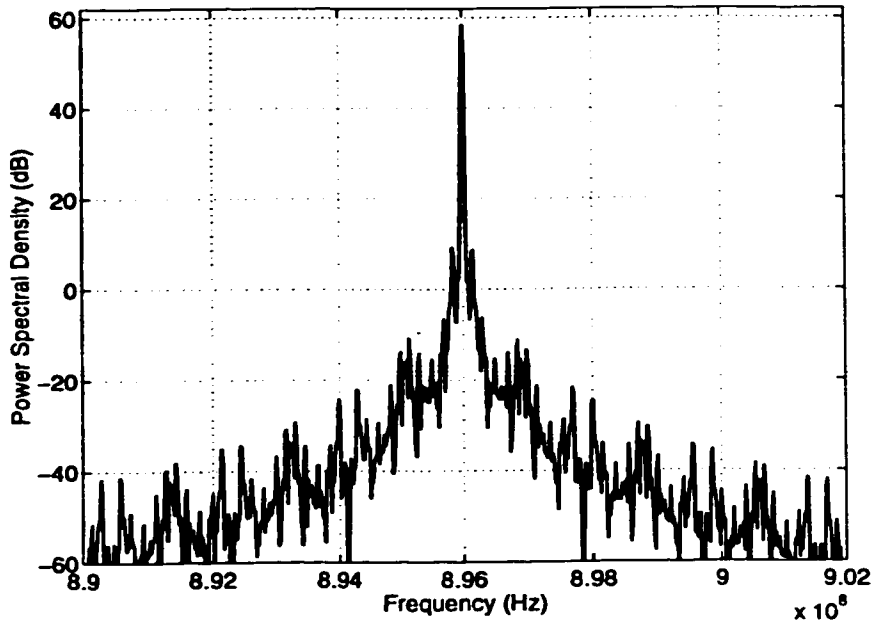

Figure 6.12: PLL output spectrum for a 6^{th}-order LPF.

6.7 Conclusion

A novel PLL architecture employing subsampling in the feedback loop greatly reduces the division ratio, thus attenuating the phase-noise due to the phase detector and input reference, without the use of extra loop or mixers. Also, if a DDS is used to provide the reference, it will have significantly fewer bits. In addition, the reference frequency to the PLL is relatively high allowing for a large loop BW and thus lower VCO phase-noise contribution.

The noise-aliasing caused by subsampling does not constitute a problem in this architecture as the noise-floor is much lower than the phase-noise in the range of interest. The linearity of the sampler and the low harmonics of the VCO are important to keep the spur level down. These spurs should be adequately filtered

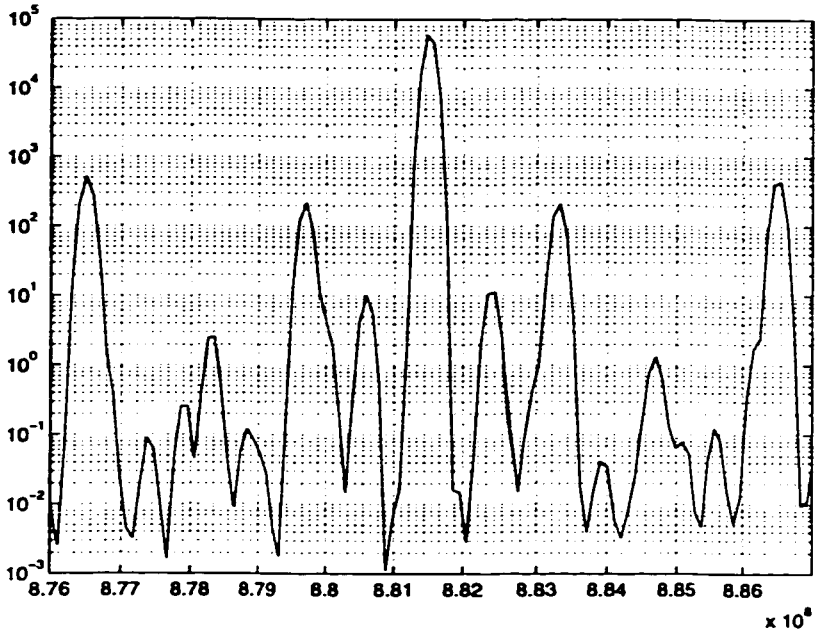

Figure 6.13: VCO phase-noise suppression.

by the LPF and/or the loop response. If a wide tuning range is desired at the output, an overtone crystal or frequency doublers might be needed.

Chapter 7

Conclusion and Future Work

From a system designer point-of-view, the simplest way to integrate the VCO on-chip would be if a low-noise and low (or moderate) power oscillator that meets the requirements by itself is available. Not only does this shift the burden onto the circuit designer, but also leads to a relatively straightforward design of the PLL without the need for extra components as in the architectures presented in this work. Thus the great deal of technology- and circuit-level research that is directed towards this objective is certainly important.

However, in current technologies, the inductor quality factors remain very limited and the noise levels much higher than the requirements of many communication standards. The incremental improvements in this path are relatively small. The monolithic VCO that meets the stringent phase-noise requirements over all temperature, process and operation conditions, while having enough tuning range and tolerable power consumption does not seem to be eminent. In this situation, the only way to achieve a completely monolithic VCO is through the use of a wide-

bandwidth PLL. In this work, we have presented three PLL architectures with wide-BW that enhance the performance of the frequency synthesizer and allow the integrating the VCO.

The first wide-bandwidth PLL architecture is the nested-loop PLL, where a wide-bandwidth loop is configured inside another slow outer PLL. The outer PLL has the same reference frequency as an ordinary PLL and thus achieves the same spurious and speed performance. The inner loop however provides a very wide BW that greatly attenuates the phase-noise of the VCO.

The nested-loop architecture is very simple and efficient and allows for very wide bandwidths. The BW achievable can be as large as 10-20MHz where the VCO noise is down approximately to the noise-floor. This results in a low phase-noise level at all offsets from the carrier. A method of incorporating the nested-loop PLL in the superheterodyne transceiver, resulting in no need for any extra components, has been proposed. This however involves undesirable consequences and a restructuring of the radio architecture. It is thus probably best to implement the nested-loop synthesizer independently. Yet the extra overhead is small. The IF VCO, although needing an off-chip tank, consumes little power and so does the frequency counter.

The phase-noise levels achievable with the nested-loop PLL are very low. With a 65MHz reference, we have designed a bipolar PLL with monolithic VCO that achieves -100dBc/Hz phase-noise at 10kHz offset from 1GHz. This is more than an order of magnitude better than typical monolithic VCOs, and was limited only by the available signal generator. Due to its simplicity, we see the nested-loop PLL as a very attractive method of VCO integration.

An improvement to the nested-loop architecture was the closed-loop voltage-to-frequency conversion PLL. A novel voltage-locked loop was developed that eliminated the need of the IF VCO. The CLV/F-based PLL achieves the same benefits as the nested loop PLL, while eliminating the IF VCO and its associated off-chip tank and power consumption. The achievable BW of the CLV/F might not be as large though, due to the possible difficulty in designing an efficient frequency-to-voltage converter at frequencies in the 100-200MHz range.

The low-noise operation of the CLV/F converter is contingent upon having a low-noise opamp and a low-noise, high-gain frequency-to-voltage converter. The former should not be a problem given the fact that many of the other parameters of the opamp are relaxed, and since the field of low-noise opamps is well developed. On the other hand, F/V converters are not readily available. We have proposed an F/V converter that achieves a large gain. However, in order to keep the noise contribution of the F/V low, the integrating capacitor must be increased leading to larger bias current. In addition, extra digital circuitry is needed for generating the different phases of the clock, adding to the power consumption and to the substrate noise. An implementation of the CLV/F has shown a 20dB improvement in the phase-noise of a ring oscillator. However, if an LC-VCO with lower noise levels was used, the noise of the F/V might start to dominate.

The idea of using F/V converters in frequency synthesis is novel and seems to be very promising. It does need further development though, primarily to simplify the design of the F/V and reduce its power. Until then, the nested-loop PLL could prove easier and of better performance in many situations.

Finally, we have presented a PLL architecture that uses subsampling in the feedback path to downconvert the VCO signal. The subsampled PLL primarily addresses the problem of noise and spurs at the PLL input. Using the subsampled architecture, a great reduction in the division ratio can be achieved without using any extra PLLs, VCOs or mixers. Instead, the system crystal was utilized. The reduction in the division ratio, led to a lower input-contribution to the phase-noise and lower spur content. This architecture was particularly suitable for DDS-driven PLLs. The constraints on the DDFS are greatly relaxed leading to a feasible, low-power design. Such a architecture can make the DDS-driven PLL a practical solution.

The BW of the loop in the subsampled PLL is moderately increased (e.g. 100kHz-1MHz). This leads to a faster PLL and a reduced VCO noise close to the carrier. A larger increase in BW would necessitate running the DDFS at a higher frequency which might not be desirable.

The two main limitations of the subsampled PLL is the sampling spurs and the value of the sampling frequency. The former problem can be alleviated by careful frequency planning and circuit design. The later however can be a limitation in some cases. The sampling frequency has to be at least twice the required output tuning range, preferably larger. In system where the frequency band is large, the required value of the system crystal frequency can be too large. In these cases, the use of overtone crystals or frequency doublers might be necessary.

Table 7.1 compares the three main PLL architectures presented in this work to a classical PLL with respect to key performance parameters.

	Nested-Loop PLL	Closed-Loop V/F PLL	Subsampled-PLL
VCO noise	greatly attenuated	greatly attenuated	moderately attenuated
Input noise	same	same	greatly reduced
Integration	Highly integrated (external tank for IF VCO)	completely integrated	external VCO
Complexity	very simple	needs careful F/V design	needs low-frequency synthesizer (e.g. DDFS)
Spurious tones	easily filtered	easily filtered	critical
Power	small	depends on F/V	much lower than hybrid DDFS/PLL

Table 7.1: Comparison of the proposed PLL architectures.

There is several areas of extension to the work presented here. A major area is the improved implementation of the closed-loop V/F loop. The objective is to implement a CLV/F loop with an integrated LC-VCO and a low input-noise level. This necessitates circuit-level improvement and further architectural investigations. A better circuit implementation can lead to an improved F/V converter with lower noise and power consumption. In addition, other architectures for frequency-to-voltage conversion should be investigated. A simple F/V architecture would be very beneficial for the CLV/F PLL. The CLV/F loop would then be incorporated in a outer PLL so that a full frequency synthesizer is built. This will fully demonstrate the potential of this architecture.

The subsampled PLL also warrants further work, primarily on the implemen-

tation level. With several pitfalls detailed in this work, a careful circuit implementation should follow. The harmonics of the VCO should be minimized and a linear sampler designed. Careful frequency planning is also necessary. Initially, only the subsampled loop, apart from the DDFS, would be designed, fabricated and tested with an ordinary signal generator or an off-the-shelf DDFS. When the design is refined, the DDFS would be designed in the second stage. An off-chip VCO is seen appropriate in the first stage in order to isolate the problems of on-chip VCOs. After the subsampled PLL is fully tested, the achievable BW is to be examined and the use of an integrated VCO is to be evaluated.

Further, an eventual goal is to explore the possibility of a combination of the ideas presented here. For example, a nested-loop PLL can be built, with the outer PLL being a subsampled one. This would provide an inner wide-BW for VCO integration while providing a low input phase-noise and spur level. Such a solution addresses both ends of the frequency synthesizer problem: the VCO phase-noise, and the input phase-noise and spurs. In such an architecture, a fast settling time will result, and the use of the DDFS as the input reference will allow the use of the synthesizer in direct modulation.

In conclusion, this research has advanced the state-of-the-art in the area of monolithic frequency synthesizers for wireless transceivers through the following contribution:

- Presenting a novel nested-loop PLL architecture that achieves very low VCO phase-noise without affecting the spur or settling time performance [85, 86].
- Developing a novel means for voltage-to-frequency conversion via a wide-

bandwidth feedback loop [74, 87, 88].

- Proposing the employment of subsampling techniques in frequency synthesis in order to achieve very low division ratios while maintaining a low-frequency reference [83, 89, 90].

Bibliography

[1] T. Stetzler et. al., "A 2.7-4.5 single chip GSM transceiver rf integrated circuit," *IEEE J. Solid-State Circuits*, vol. 30, no. 12, pp. 1421–1428, Dec. 1995.

[2] C. Marshall et. al., "A 2.7V GSM transceiver ICs with on-chip filtering," in *ISSCC Dig. Tech. Papers*, pp. 148–149, Feb. 1995.

[3] T. Yamawaki et. al., "A 2.7-V GSM RF transceiver IC," *IEEE J. Solid-State Circuits*, vol. 32, no. 12, pp. 2089–2096, Dec. 1997.

[4] K. Hess and D. Fague, "Performance evaluation of a single-chip radio transceiver," in *Intl. Symp. Circuits and Systems*, pp. 1048–1051, June 1996.

[5] H. Sato et. al., "A 1.9-GHz single chip IF transceiver for digital cordless phones," *IEEE J. Solid-State Circuits*, vol. 31, no. 12, pp. 1974–1980, Dec. 1996.

[6] C. Hull, J.L. Tham, and R. Chu, "A direct-conversion receiver for 900MHz (ISM band) spread-spectrum digital cordless telephone," *IEEE J. Solid-State Circuits*, vol. 31, no. 12, pp. 1955–1963, Dec. 1996.

[7] W. Veit et. al., "A 2.7V 800MHz-2.1GHz transceiver chip-set for mobile radio applications in 25GHz f_t Si-bipolar," *Bipolar/BiCMOS Circuits and Tech. Mtng.*, pp. 175–178, 1994.

[8] J. Sevenhans et. al., "An analog radio front-end chip set for a 1.9ghz mobile radio telephone application,," in *ISSCC Dig. Tech. Papers*, pp. 44–45, Feb. 1994.

[9] T. Manku, "Microwave CMOS- device physics and design," *IEEE J. Solid-State Circuits*, vol. 34, no. 3, pp. 277–285, March 1999.

[10] D. Shaeffer and T. Lee, "A 1.5-V, 1.5-GHz CMOS low noise amplifier," *IEEE J. Solid-State Circuits*, vol. 32, no. 5, pp. 745–759, May 1997.

[11] A. Rofougaran et. al., "A 1GHz CMOS RF front-end IC for a direct-conversion wireless receiver," *IEEE J. Solid-State Circuits*, vol. 31, no. 7, pp. 880–889, July 1996.

[12] J. Crols and M. Steyaert, "A 1.5-GHz highly linear CMOS downconversion mixer," *IEEE J. Solid-State Circuits*, vol. 30, no. 7, pp. 736–742, July 1995.

[13] P. Kinget and M. Steyaert, "A 1-GHz upconversion mixer," in *IEEE Custom Integrated Circuits Conf. Dig. of Tech. Papers*, pp. 197–200, May 1996.

[14] J. Rudell et. al., "A 1.9-GHz wide-band IF double conversion CMOS receiver for cordless telephone applications," *IEEE J. Solid-State Circuits*, vol. 32, no. 12, pp. 2071–2088, Dec. 1997.

[15] K. Yamamoto et. al., "A GaAs RF transceiver IC for 1.9-GHz mobile digital communication systems," in *ISSCC Dig. Tech. Papers*, pp. 340–341, Feb. 1996.

[16] M. Muraguchi et. al., "1.9GHz-band low voltage and low power consumption RF IC chip-set for personal communications," in *Intl. Symp. Circuits and Systems*, pp. 504–507, June 1994.

[17] P. Gray and R. Meyer, "Future directions in silicon ICs for RF personal communications," in *IEEE Custom Integrated Circuits Conf. Dig. of Tech. Papers*, pp. 83–90, May 1995.

[18] S. Heinen et. al., "A 3V 2-GHz transmitter IC for digital radio communication with integrated VCOs," in *ISSCC Dig. Tech. Papers*, pp. 146–147, Feb. 1995.

[19] C. Takahashi et. al., "A 1.9GHz Si direct conversion receiver IC for QPSK modulation systems," in *ISSCC Dig. Tech. Papers*, pp. 138–139, Feb. 1995.

[20] A. Karanicolas, "A 2.7V 900MHz CMOS LNA and mixer," in *ISSCC Dig. Tech. Papers*, pp. 50–51, Feb. 1996.

[21] J. Sevenhans et. al., "An integrated Si-bipolar RF transceiver for a zero-IF 900MHz GSM digital mobile radio front-end of a hand portable phone," in *IEEE Custom Integrated Circuits Conf. Dig. of Tech. Papers*, pp. 7.7.1–4, May 1991.

[22] Q. Huang, P. Orsatti, and F. Piazza, "GSM transceiver front-end circuits in 0.25-μm CMOS," *IEEE J. Solid-State Circuits*, vol. 30, no. 3, pp. 292–303, March 1999.

[23] B. Razavi, "Challenges in portable RF transceiver design," *The Optoelectronics Magazine*, vol. 12, no. 5, pp. 12–25, Sept. 1996.

[24] A. Abidi, "Direct-conversion radio transceivers for digital communications," *IEEE J. Solid-State Circuits*, vol. 30, no. 12, pp. 1399–1410, Dec. 1995.

[25] B. Razavi, "A 1.8GHz CMOS VCO," in *ISSCC Dig. Tech. Papers*, pp. 388–389, Feb. 1997.

[26] F. Brianti, "High integration CMOS RF transceivers," *Proc. of AACD*, April 1996.

[27] A. Bateman and D. Haines, "Direct conversion receiver design fr compact low cost portable mobile radio terminals," 39^{th} *IEEE Vehc. Tech. Conf.*, pp. 57–62, 1989.

[28] A. Abidi, "Low-power radio-frequency IC's for portable communications," *Proc. IEEE*, vol. 83, no. 4, pp. 544–569, April 1995.

[29] J. Crols and M. Steyaert, "A single-chip 900MHz CMOS receiver front-end with a high-performance low-IF topology," *IEEE J. Solid-State Circuits*, vol. 30, no. 12, pp. 1483–1492, Dec. 1995.

[30] T. Tsukahara et. al., "A 2V 2GHz Si-bipolar direct conversion quadrature modulator," *IEEE J. Solid-State Circuits*, vol. 31, no. 2, pp. 263–267, Feb. 1996.

[31] P. Davis, "RF ICs for cellular and cordless telephones," in *1996 ISSCC Tutorial*, Feb. 1996.

[32] K. Pahlavan and A. Levesque, *Wireless Information Networks*, John Wiley & Sons, New York, 1995.

[33] B. Razavi, "Challenges in the design of frequency synthesizers for wireless applications," in *IEEE Custom Integrated Circuits Conf. Dig. of Tech. Papers*, pp. 395–402, May 1997.

[34] B. Razavi, "A study of phase-noise in CMOS oscillators," *IEEE J. Solid-State Circuits*, vol. 30, no. 12, pp. 1421–1429, Dec. 1995.

[35] T. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*, Cambridge University Press, Cambridge, UK, 1998.

[36] R. Best, *Phase-Locked Loops*, McGraw Hill, New York, 1993.

[37] W. Rhee, "Design of high-performance CMOS charge-pumps in phase-locked loops," in *Intl. Symp. Circuits and Systems*, vol. II, pp. 545–548, June 1999.

[38] Philips Semiconductors, *SA8025 Fractional-N Synthesizer for 2GHz Band Applications*, Application Note AN1891.

[39] Texas Instruments, *Low-Voltage 1.2GHz Fractional-N/Integer-N Synthesizer*, Datasheet TRF2050.

[40] Peregrine Semiconductors, *1.2GHz/550MHz Dual Fractional-N FlexiPower PLL for Frequency Synthesis*, Datasheet PE3291.

[41] T. Riley, M. Copeland, and T. Kwasniewski, "Delta-sigma modulation in fractional-N frequency synthesis," *IEEE J. Solid-State Circuits*, vol. 28, no. 5, pp. 553–559, May 1993.

[42] J. Crawford, *Frequency Synthesizer Design Handbook*, Artech House, Norwood, MA, 1994.

[43] B. Goldberg, *Digital Techniques in Frequency Synthesis*, McGraw Hill, 1996.

[44] A. Edman, A. Bjorklid, and I. Soderquist, "A 0.8μm CMOS 350MHz quadrature direct digital frequency synthesizer with integrated D/A converters," in *Symp. VLSI Circuits Dig. of Tech. Papers*, pp. 54–55, June 1998.

[45] J. Vankka et. al., "A direct digital synthesizer with an on-chip D/A-converter," *IEEE J. Solid-State Circuits*, vol. 33, no. 2, pp. 218–227, Feb. 1998.

[46] L. Larson, "Integrated circuit technology options for RFIC's- present status and future directions," *IEEE J. Solid-State Circuits*, vol. 33, no. 3, pp. 387–399, March 1998.

[47] B. Kim et. al., "PLL/DLL system noise analysis for low jitter clock synthesizer design," in *Intl. Symp. Circuits and Systems*, pp. 31–34, June 1994.

[48] K. Negus et. al., "Highly-integrated transmitter RFIC with monolithic narrowband tuning for digital cellular handsets," in *ISSCC Dig. Tech. Papers*, pp. 38–39, Feb. 1994.

[49] B. Razavi, "High performance frequency synthesizer design," in *1999 Custom Integrated Circuits Conf. Tutorial*, May 1999.

[50] K. Kurokawa, "Injection locking of microwave solid-state oscillators," *Proc. IEEE*, vol. 61, pp. 1386–1410, Oct. 1973.

[51] D. Wolaver, *Phase-Locked Loop Circuit Design*, Prentice Hall, 1991.

[52] N. Filiol, T. Riley, C. Plett, and M. Copeland, "An agile ISM band frequency synthesizer with built-in GMSK data modulation," *IEEE J. Solid-State Circuits*, vol. 33, no. 7, pp. 998–1008, July 1998.

[53] M. Margarit et. al., "A low-noise, low-power VCO with automatic amplitude control for wireless applications," *IEEE J. Solid-State Circuits*, vol. 34, no. 6, pp. 761–771, June 1999.

[54] Q. Huang, "Power consumption vs LO amplitude for CMOS colpitts oscillators," in *IEEE Custom Integrated Circuits Conf. Dig. of Tech. Papers*, pp. 255–258, May 1997.

[55] M. Zannoth et. al., "A fully integrated VCO at 2 GHz," *IEEE J. Solid-State Circuits*, vol. 33, no. 12, pp. 1987–1991, Dec. 1998.

[56] J. Craninckx, M. Steyaert, and H. Miyakawa, "A fully integrated spiral-LC CMOS VCO set with prescaler for GSM and DCS-1800 systems," in *IEEE Custom Integrated Circuits Conf. Dig. of Tech. Papers*, pp. 403–406, May 1997.

[57] J. Craninckx and M. Steyaert, "A 1.8-GHz low-phase-noise CMOS VCO using optimized hollow spiral inductors," *IEEE J. Solid-State Circuits*, vol. 32, no. 5, pp. 736–744, May 1997.

[58] T. Ahrens and T. Lee, "A 1.4-GHz 3-mW CMOS LC low phase noise VCO using tapped bond wire inductances," *Intl. Symp. on Low Power Electronics and Design*, pp. 16–19, August 1998.

[59] F. Svelto, S. Deantoni, and R. Castello, "A 1.3-GHz CMOS VCO with 28% frequency tuning," in *IEEE Custom Integrated Circuits Conf. Dig. of Tech. Papers*, pp. 645–648, May 1999.

[60] M. Perrott, T. Tewksbury, and C. Sodini, "A 27-mW CMOS fractional-n synthesizer using digital compensation for 2.5-Mb/s GFSK modulation," *IEEE J. Solid-State Circuits*, vol. 32, no. 12, pp. 2048–2060, Dec. 1997.

[61] Analog Devices Inc., *CMOS, 125MHz Complete DDS Synthesizer*, 1999, Datasheet AD9850.

[62] A. Yamagishi et. al., "A 2-V, 2-GHz low-power direct digital frequency synthesizer chip-set for wireless communications," *IEEE J. Solid-State Circuits*, vol. 33, no. 2, pp. 210–217, Feb. 1998.

[63] J. Craninckx and M. Steyaert, "A fully integrated CMOS DCS-1800 frequency synthesizer," *IEEE J. Solid-State Circuits*, vol. 33, no. 12, pp. 2054–2065, Dec. 1998.

[64] Murata Inc., *Surface Mount Chip Inductors- LQN1H Series Datasheet.*

[65] Coilcraft Inc., *Chip Inductors-1008CS Series Datasheet.*

[66] F. Piazza and Q. Huang, "A 170 MHz RF front-end for ERMES pager applications," *IEEE J. Solid-State Circuits*, vol. 30, no. 12, pp. 1430–1437, Dec. 1995.

[67] S. Sanielevici et. al., "A 900-MHz transceiver chipset for two-way paging applications," *IEEE J. Solid-State Circuits*, vol. 33, no. 12, pp. 2160–2168, Dec. 1998.

[68] B. Jansen, K. Negus, and D. Lee, "Silicon bipolar VCO family for 1.1 to 2.2GHz with fully-integrated tank and tuning circuits," in *ISSCC Dig. Tech. Papers*, pp. 392–393, Feb. 1997.

[69] R. Ahola et. al., "Phase detector with no dead zone and a very wide output voltage range chargepump," in *Intl. Symp. Circuits and Systems*, vol. 1, pp. 155–159, June 1998.

[70] A. Hill and J. Surber, "The PLL dead zone and how to avoid it," *RF Design*, pp. 131–134, March 1992.

[71] W.-Z. Chen and J.-T. Wu, "A 2-V, 1.8-GHz BJT phase-locked loop," *IEEE J. Solid-State Circuits*, vol. 34, no. 6, pp. 784–789, June 1999.

[72] J. Parker and D. Ray, "A 1.6-GHz CMOS PLL with on-chip loop filter," *IEEE J. Solid-State Circuits*, vol. 33, no. 3, pp. 337–343, March 1998.

[73] B. Razavai, "A 2-GHz 1.6mW phase-locked loop," *IEEE J. Solid-State Circuits*, vol. 32, no. 5, pp. 730–735, May 1997.

[74] Amr N. Hafez and M. I. Elmasry, "Closed-loop voltage-to-frequency converter," Patent 09/394,185, Pending, 1998.

[75] F. Gardner, "Charge-pump phase-lock loops," *IEEE. Trans. Comm.*, vol. COM-28, pp. 1849–1858, Nov. 1980.

[76] National Semiconductors, *Low Noise, Low V_{os}, Single Op Amp*, 1999, Datasheet LMV751.

[77] Texas Instruments, *Low Noise Precision Rail-to-Rail Output Operational Amplifier*, 1997, Datasheet TLC2202.

[78] C. Menolfi and Q. Huang, "A fully integrated untrimmed CMOS instrumentation amplifier with submicrovolt offset," *IEEE J. Solid-State Circuits*, vol. 34, no. 3, pp. 415–420, March 1999.

[79] T. Holman and C. Alvin, "Compact low noise operational amplifier for a 1.2μm digital CMOS technology," *IEEE J. Solid-State Circuits*, vol. 30, no. 6, pp. 710–714, June 1995.

[80] Analog Devices, *Low Noise, Precision Operational Amplifier*, Datasheet OP-27.

[81] S. Norsworthy et. al., *Delta-Sigma Data Converters: Theory, Design and Simulation*, IEE Press, 1997.

[82] P. Chan, A. Rofougaran, and A. Abidi, "A higly linear 1-GHz CMOS down-conversion mixer," *European Solid-States Circuits Conf.*, pp. 210–213, Sept. 1993.

[83] Amr N. Hafez and M. I. Elmasry, "PLL architecture employing sub-sampling," Patent Pending, 1998.

[84] Qualcomm Inc., *Hybrid PLL/DDS Frequency Synthesizers*, 1997, Application Note.

[85] A. Hafez and M. Elmasry, "A fully-integrated low phase-noise nested-loop PLL for frequency synthesis," *submitted to 2000 IEEE Custom Integrated Circuits Conf.*, Nov. 1999.

[86] A. Hafez and M. Elmasry, "A fully-integrated low phase-noise nested-loop PLL for frequency synthesis," *submitted to IEEE. Journal of Solid-State Circ.*, Dec. 1999.

[87] A. Hafez and M. Elmasry, "Fully integrated low phase-noise PLLs using closed-loop voltage-to-frequency converter architectures," in *IEEE Custom Integrated Circuits Conf. Dig. of Tech. Papers*, pp. 353–356, May 16-19, 1999.

[88] A. Hafez and M. Elmasry, "Fully integrated low phase-noise PLLs using closed-loop voltage-to-frequency converter architectures," *submitted to IEEE. Journal of Solid-State Circ.*, Dec. 1999.

[89] A. Hafez and M. Elmasry, "Low power PLL architecture for wireless transceivers," in *Intl. Symp. Circuits and Systems*, vol. II, pp. 549–552, May 30-June 2, 1999.

[90] A. Hafez and M. Elmasry, "A novel low power low phase-noise PLL architecture for wireless transceivers," *Proc. of the 1999 Great Lakes Symposium on VLSI*, pp. 306–309, March 4-6, 1999.