CIRCUIT TECHNIQUES FOR COGNITIVE RADIO RECEIVER FRONT-ENDS

A DISSERTATION

SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL

OF THE UNIVERSITY OF MINNESOTA

BY

BODHISATWA SADHU

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

Doctor of Philosophy

June, 2012

Acknowledgements

This thesis is hardly a lone effort. I would like to express my gratitude to my advisor, Prof. Ramesh Harjani, whose guidance, motivation, friendship, and support has made this thesis feasible, and shaped my life and career over the last five years. I am indebted to him for his unfailing support and his unbiased advice at all times during my PhD. Looking back, I now feel astonished at the immense patience with which he guided me when I first joined the lab with my embarrassing lack of knowledge about circuit design. Over the years, our discussions, ranging from circuit theory to Sartre and Hesse, have become a dear part of my life. In years to come, I will remember and cherish all our animated conversations, the joyous parties at his home, and Savita ma'am's wonderful company.

Along with Prof. Harjani, the lab exposed me to an infectious zeal for circuits research to which I owe a large part of my enthusiasm. During my PhD, I had the privilege of working with some of the most wonderful people I now know. I was introduced into research in circuit design by Narasimha Lanka and Satwik Patnaik. I learned one of

the most important lessons of my PhD from Satwik, that of unswerving perseverance. For that, as for a lot of other things, I will be ever grateful to him. I am indebted to Sudhir Kudva for his unfailing friendship and support every single day of my PhD. With him around, hard times never seemed that hard, and work became enjoyable. I was privileged to work with Martin Sturm on the CRAFT project. I am grateful to him for sharing the intense work-load that the project often demanded. He has taught me a lot during the course of the project, and my PhD. I could not have learned about chip testing from a better engineer. I am also grateful to Sachin Kalia, Ashutosh Mehra and Mohammad Elbadry for their friendship. Our discussions have been a great learning experience for me. Also, with them around, there was never a dull moment in the lab. I am thankful to Jaehyup Kim for sharing his expertise in analog layout, and to Taehyoun Oh for his help and friendship during the last few years. I will also fondly remember the invigorating (and sometimes unending) discussions I have shared with Rakesh Palani during the final year of my PhD.

I am thankful to Prof. Anand Gopinath for his unconditional help and precious advice throughout my PhD. I have thoroughly enjoyed his stories filled with wit, inspiration, and humor. I am grateful to Prof. Chris Kim and Brian Sadler for their advice and support. Also, I would like to express my sincere thanks to Prof. Anand Gopinath, Prof. Chris Kim, Prof. Richard Moeckel and Brian Sadler for serving on my final defense committee.

My experience during internships at IBM Watson Laboratories was of great learning and stimulating research. I am deeply thankful to Alexander Rylyakov and J-O. Plouchart for their patient and encouraging mentorship. I am also grateful to José Tierno, Mark Ferriss, Scott Reynolds, Daniel Friedman, Arun Natarajan, Alberto Valdes-Garcia, Mihai Sanduleanu, Ankur Agrawal and Soner Yaldiz for making my internships some of the most enjoyable and productive periods of my graduate studies.

Carlos Soria and Chimai Nguyen have been incredible in their support of all the computer systems that we rely on for our livelihood. I am grateful to them for stretching their support beyond the stipulated work hours and work days, and treating our deadlines as theirs. Jamie Halleckson, Linda Jagerson, Parker Smith, Kyle Dukart, Dan Dobrick, Gina McCabe and Becky Colberg have worked hard to shade us from all the bureaucracy, and often do thankless jobs to make all this research possible. Thank you.

I am grateful to MOSIS, TAPO and IBM for providing fabrication of prototype chips. I would also like to express my gratitude toward the funding agencies: Center for Circuits and Systems Solutions (C2S2), Semiconductor Research Corporation (SRC), the U.S. Army Research Laboratories (ARL) and Defense Advanced Research Projects Agency (DARPA) who made this research possible.

Life as a graduate student would not have been as fulfilling without the company of several friends. I am grateful to Sonia Das and Ajay Yekkirala for their friendship, and their unfailing concern for my well-being (especially during those tape-outs). My room-mates, Parthiv Daggolu and Reetam Chakrabarti have had to endure me during

my deadlines as well as their aftermath, and borne both equally well. I will cherish the numerous refreshing discussions I have had with Gaurab Samanta and Shameek Bose, over dinner, coffee, or long drives to Chicago. I am thankful to Vidhi Zalani for her cheerful companionship, and her exemplary patience. I feel blessed to have the love and generosity of such wonderful people as Vivian Fu and Douglas Ewing touch my life. My undergraduate years at BITS, Pilani have shaped my life and career in numerous ways. I am indebted to Anurup Mitra and Prof. L. K. Maheshwari for introducing me to analog electronics. I am deeply grateful to Raghavendra Rao, Shibanka Das Biswas, Ramya Prayaga, Raghav Mehta, Ravish Patel and Naveen Karnam for their love, friendship, and support. I have been fortunate to have mentors such as Tania Roy and Madhurjya Banerjee who have always had the time for me when I needed it most. I am indebted to Samik Bhar for his love, and unfailing friendship since our childhood days.

Finally, but most importantly, I am grateful to my parents, grandparents, and my sister for everything they have given me. Fairytales do not boast of happier families. I am grateful to my father for unraveling to me the world of science, teaching, and research, and making my life much richer than it would have otherwise been. My mother has been a pillar of support, patience, and understanding through all my endeavors. I could not have done all this without her. Soumi does a brilliant job of tolerating her elder brother, of patiently listening to his rants, and yet adoring him (and even proof-reading his entire thesis overnight). I owe all of this to them.

Dedicated to my parents,
and my little sister

Abstract

This thesis discusses the design of the receiver front-end for software defined radio (SDR) based cognitive radio applications. Two aspects of SDRs for cognitive radios are distinguished: signaling and spectrum sensing. Narrowband wide tuning signaling architectures and instantaneous wideband spectrum sensing architectures are identified as candidates for feasible SDR implementations. Several architectures and circuit implementations are reviewed. Wide tuning range, low phase noise frequency synthesizers for signaling, and RF samplers and signal processors for spectrum sensing are identified as critical circuit design blocks.

A number of voltage controlled oscillator (VCO) techniques for wide-tuning range, and low phase noise frequency synthesis techniques are developed. Wide-tuning range techniques based on switched inductors are proposed as a way to design inductor-capacitor (LC) VCOs with wide-tuning ranges that maintain a good phase noise and power dissipation performance over the entire tuning range. Switched inductor VCOs are analyzed in detail, and a design framework is developed. Optimized capacitor array design techniques for wide-tuning ranges are discussed. Based on these techniques, measurements from two prototype designs are presented, that achieve tuning ranges of 87% and 157% in measurement. They also maintain good phase noise, power consumption, and figure of merit (FOM) over the entire tuning range.

In addition, a new family of VCOs that achieve superior phase noise is introduced. This set of novel topologies are based on linearized transconductance using capacitive feedback techniques. They achieve higher amplitudes of oscillation, and consequently, a superior phase noise performance. A wide tuning range is also maintained. The VCOs are analyzed, and detailed measurement results from a design prototype are presented.

For spectrum sensing, the design of CRAFT (Charge Re-use Analog Fourier Transform): an RF front-end channelizer for software defined radios (SDR) based on a 16 point analog domain FFT is described. The design relies on charge re-use to achieve 47dB average output SNDR on a 5GS/s input, and consumes only 12.2pJ/conv. These numbers represent orders of magnitude improvement on the work reported previously in literature. The thesis also briefly discusses the modeling of circuit non-idealities in CRAFT, and outlines circuit techniques for mitigating these. These design principles enable this implementation to achieve a large dynamic range even at high speeds. Additionally, these techniques can be easily extended to improve the performance of other passive switched capacitor designs.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Wireless Growth

Wireless technology has been evolving at a breakneck speed. The total number of cell-phones in use was over 4.6 billion (for a 6.8 billion world population) in 2009. Additionally, with user convenience becoming paramount, more and more functions are being implemented wirelessly (Fig. 1.1). For example, the U.S. army utilizes 40 different radios for its communications. Moreover, there is a considerable effort toward integrating all this wireless functionality in a single device. For example, smartphones today use as many as a dozen independent radios inside them.

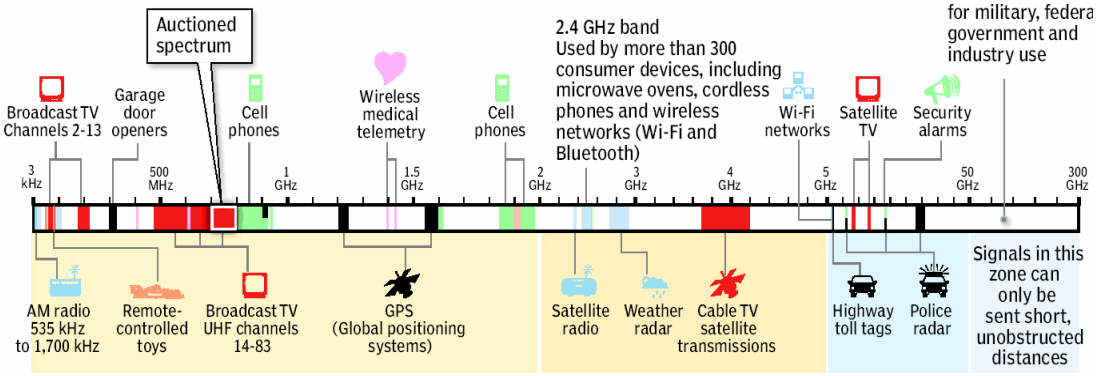

Figure 1.1: Frequency spectrum allocation in the United States (Source: New America Foundation, MCT, howstuffworks.com; Graphic: Nathaniel Levine, Sacramento Bee)

1.2 Spectral Congestion

Unlike wired communications which use dedicated connections, wireless communications use particular frequencies of electromagnetic waves in space, thus sharing a common connection medium. Consequently, with the growth of the wireless industry, the spectral congestion caused by wireless user traffic has become a significant concern, and has begun to threaten further growth of the technology [1].

However, this congestion is a result of sub-optimal frequency usage arising from the inflexibility of the spectrum licensing process. An example spectrum licensing by the Federal Communications Commission, USA, is shown in Fig. 1.2 [2]. Presently, each frequency is strictly allocated for a particular application (for e.g. 850MHz & 1.9GHz for cell phones in the USA). When this application experiences excessive usage, the corresponding frequency becomes congested (causing dropped calls and busy lines for cell-phone users for example). An example frequency usage snapshot taken at Berkeley,

California, USA [3], in Fig. 1.3 [3], clearly shows heavy usage and congestion at a few frequencies (e.g. 0-1, 1.9, 2.4 GHz) and sparse usage at others (e.g. 0.5% utilization in the 3−6 GHz band).

This inefficiency in spectrum allocation can be solved by allowing unlicensed use of spectrum. However, this will necessitate new spectrum sharing protocols to be developed in order to allow multiple users to utilize a single spectrum without causing harmful interference. A couple of approaches to spectrum sharing are being pursued: ultra-wideband (UWB), using a spectral underlay approach, and cognitive radio (CR), using a spectral overlay approach. In the ultra-wideband scenario (IEEE 802.15.3a, IEEE 802.15.4a), a secondary user is allowed to transmit in occupied bands, but with a power spectral density (PSD) so low that it is not deemed as harmful interference. The low PSD is compensated by the usage of a very wide bandwidth (>500MHz) to allow a significant transmit rate. In the cognitive radio scenario (IEEE 802.22), dynamic spectrum access (DSA) is utilized for spectrum sharing. This work focuses on the cognitive radio scenario for efficient spectrum utilization.

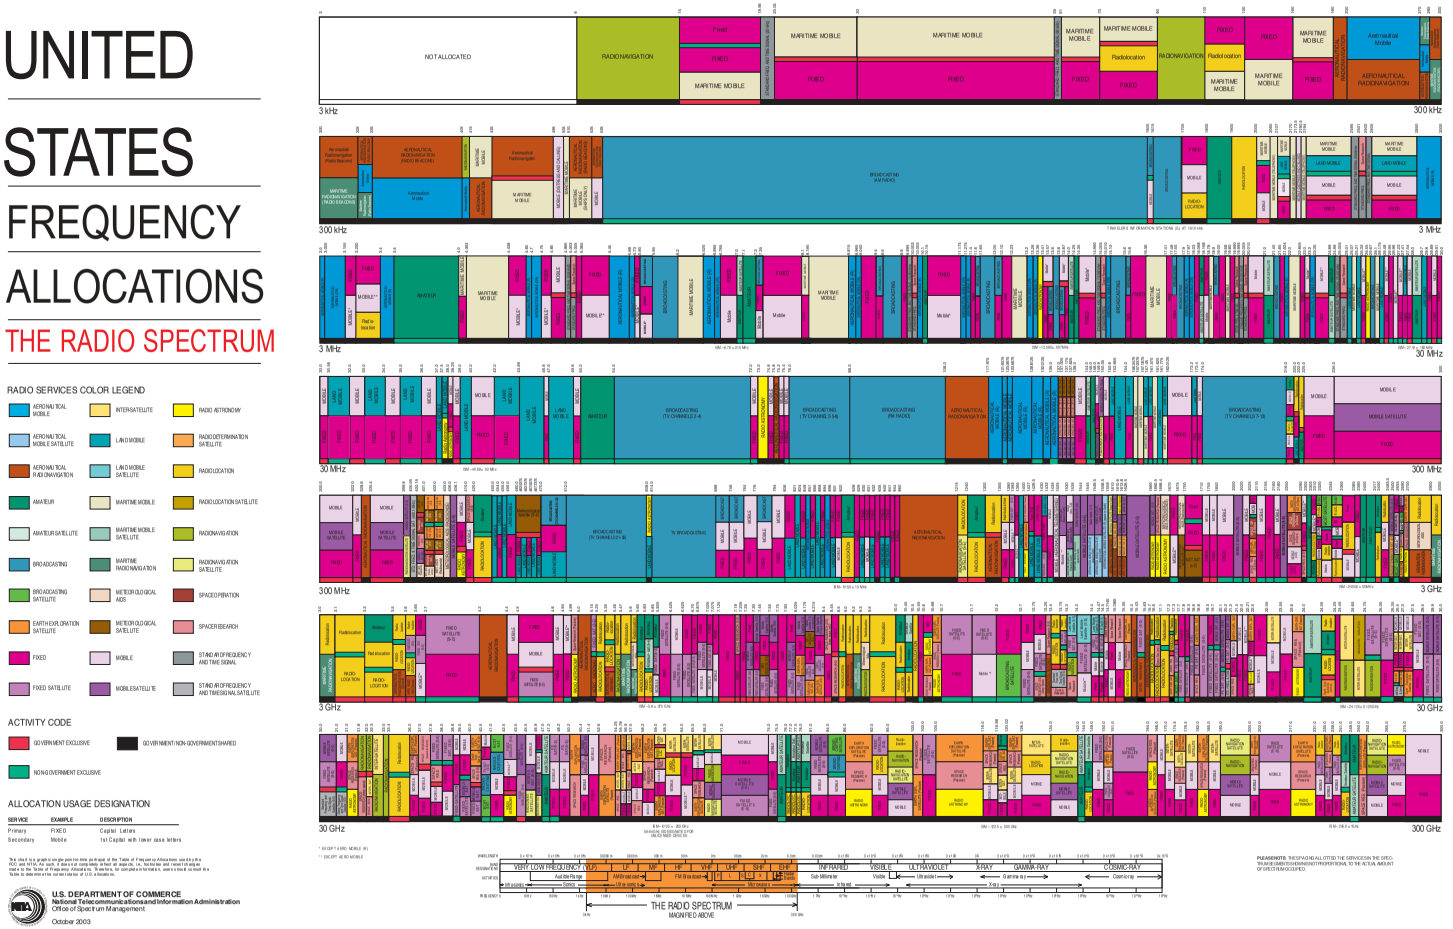

Figure 1.2: FCC spectrum allocation in the USA

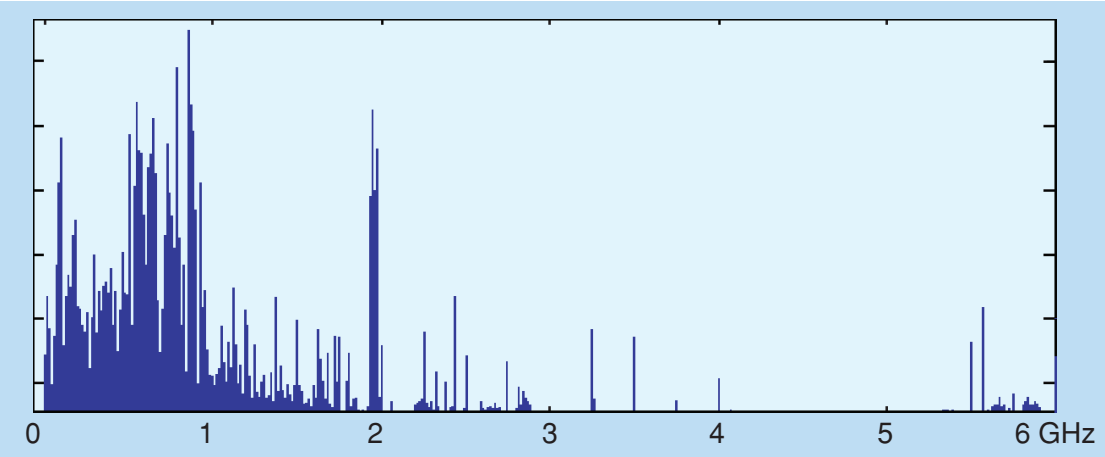

Figure 1.3: Spectrum utilization in downtown Berkeley showing 30% usage below 3GHz and 0.5% between 3−6 GHz [3]

1.3 Dynamic Spectrum Access

In dynamic spectrum access, a secondary user is allowed to utilize other allocated but temporarily unused frequencies, to solve the problem of spectral congestion. Fig. 1.4 shows a diagrammatic representation of the concept. An energy detector detects the power levels at different frequencies on a real-time basis (*spectrum sensing*) and determines spectral occupancy. The different colored cuboids in the figure represent spectral occupancy by different protocols, with different bandwidths, at different power levels, and for varying amounts of time. Frequencies not being used temporarily give rise to spectrum holes as shown using green discs. Dynamic spectrum access strives to seek out these spectral holes, and operate at these unused frequencies. This can greatly improve spectrum usage efficiency, and drastically reduce the problem of spectral congestion. For

example, if the users of the 1.88−1.9 GHz frequency (cordless phones) are allowed to use the higher under-utilized frequencies, congestion in these devices could be avoided. In fact, the different spectrum governing bodies around the world are exploring new standards to allow this kind of communication (IEEE 802.22 in the USA).

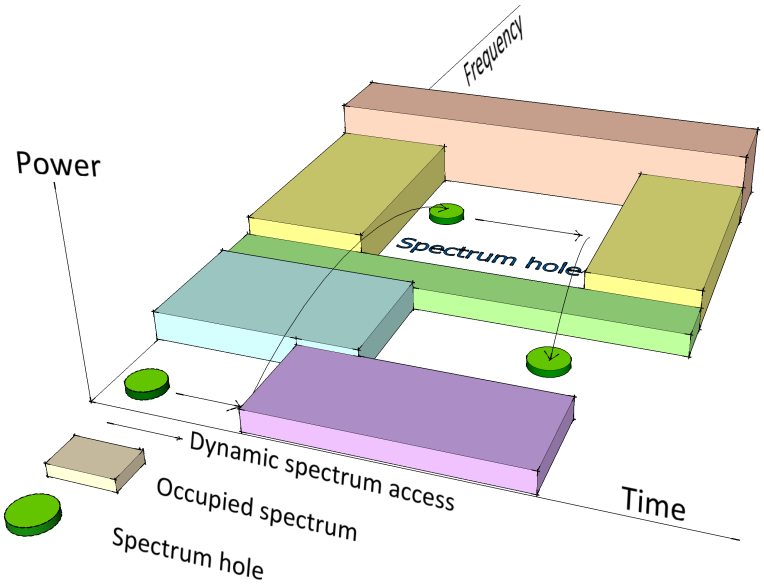

Figure 1.4: Concept of DSA

1.4 Cognitive Radio

The use of dynamic spectrum access to intelligently improve communication efficiency has been the driving force behind the cognitive radio concept. Unfortunately, the definition and scope of what constitutes a cognitive radio has remained unclear to this day (an example definition by the FCC: [4]). In the narrow sense, a cognitive radio is an intelligent device that is able to dynamically adapt and negotiate wireless frequencies

and communication protocols for efficient communications[1] . For this, each participating device will need to have many capabilities such as: determination of location, sensing the spectrum used by its neighboring devices and analyzing the external environment, to changing frequency and bandwidth, adjusting its output power level, and even altering transmission parameters and protocols [5]. Interestingly though, if these devices are capable of such functionality, they would also be able to function as multiple wireless radios and therefore function as an integrated and complete wireless solution. For example, the same wireless radio could work as a cell phone, a GPS receiver, a garage door opener, and even as a remote control for other devices as shown in Fig. 1.6. Note that this is different from multiple different radios *assembled* in a smart-phone.

Fig. 1.5 provides an indication of the growth of cognitive radios as a research area in the recent past. The figure shows results for the number of publications with different keywords published in a given year in IEEE publications. Many of the keywords represent growing research areas in wireless, where as other popular keywords such as 'VLSI' and 'DSP' have been included for comparison. The first cognitive radio paper was published in 1999; however, research in this area was relatively dormant till 2006. Since then, cognitive radios have seen a tremendous growth in research activity, and is now one of the most researched areas in wireless.

A cognitive radio can be structurally and functionally separated into (1) a software

[1] It is envisioned, that in the future, the cognitive radio functionality will also include spacio-temporal intelligence, and significant learning capabilities heralding a new era in wireless communications.

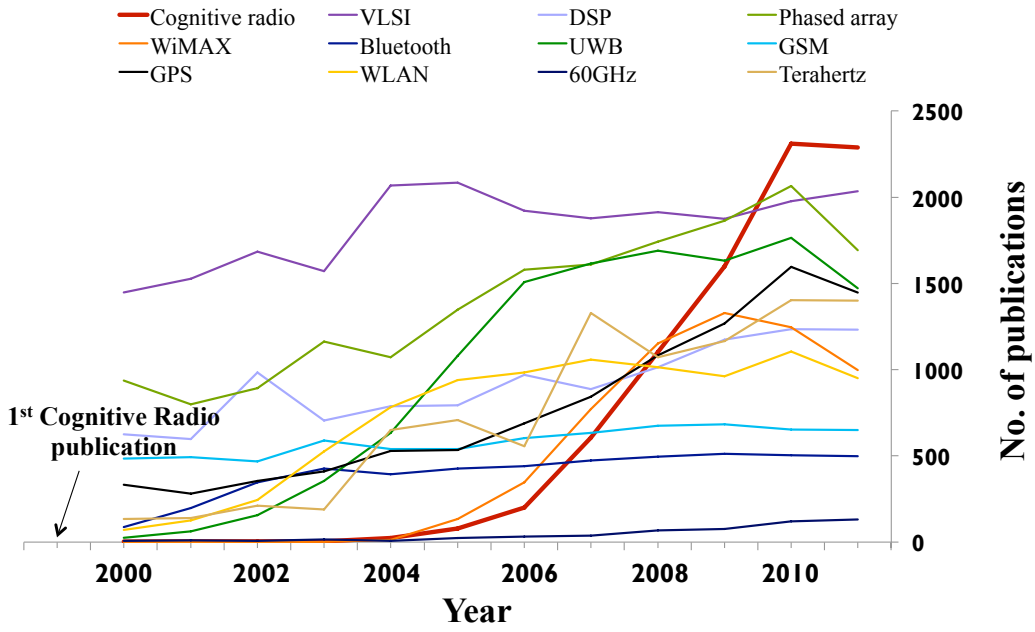

Figure 1.5: Number of publications on different research areas in the last decade

defined radio (SDR) unit that forms the hardware of the cognitive radio, and an intelligence unit, that provides the required software based intelligence (cognition) to the radio. In this thesis, the SDR unit, and more specifically, the receiver front-end of the SDR, will be emphasized.

Like any communication device, the SDR needs a transmitter and a receiver. However, unlike a traditional wireless tranceiver, a cognitive radio not only needs to perform signal transmission and reception (*signaling*) but also needs to monitor the spectrum in real-time (*spectrum sensing*) in order to execute dynamic spectrum access. Spectrum sensing provides a wideband frequency snapshot helping the cognitive radio identify currently unused (potentially usable) frequencies. The cognitive radio then decides

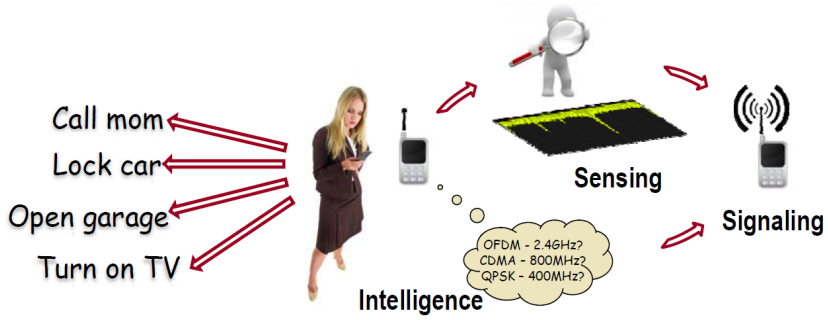

Figure 1.6: Functionality of an envisioned cognitive radio

which unused frequency and protocol to use, and starts signaling at that frequency [6] (Fig. 1.6).

1.4.1 Spectrum Sensing

The dynamic spectrum access relies on dynamic spectrum monitoring using a spectrum sensor. Among other features, this continuous wideband monitoring makes the cognitive radio unique in its hardware. From the hardware perspective, the spectrum sensor remains extremely challenging. Even for narrowband (small frequency range, < 100MHz) spectrum sensors, limiting the power consumption is a challenge. The cognitive radio spectrum sensor, on the contrary, needs to detect signals at all frequencies of interest instantaneously. Additionally, it needs to detect signals more unfailingly (about 100 times better sensitivity) than narrowband receivers to overcome the hidden-terminal problem (Fig. 1.7), shadowing, channel fading, multi-path, etc., lest it causes interference to other users due to incorrect sensing [3].

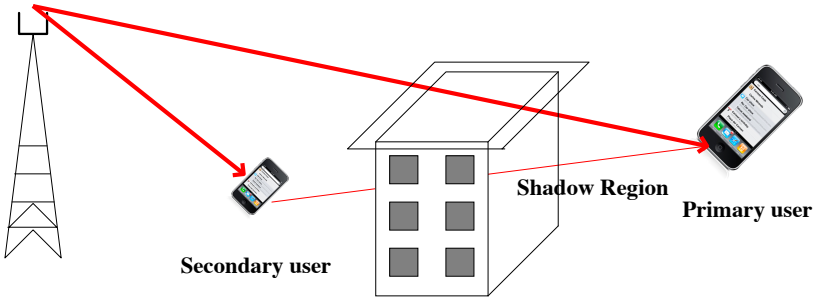

Figure 1.7: The hidden terminal problem in cognitive radios

1.4.2 Signaling

The signaling in a cognitive radio is somewhat similar to that in traditional radios, in the sense of being relatively narrowband. However, the cognitive radio signaling transceiver needs to be far more flexible. It should have the capability of operating over a very wide frequency range, changing its bandwidth of operation, altering the transmitted power, as well as using multiple standards for communications.

1.5 Organization

This dissertation focuses on the architecture and circuit design for cognitive radio receiver front-ends.

Chapter 2 explores the different types of architectures for signaling and sensing in software defined radios for cognitive radio applications. A number of competing architectures are reviewed and discussed. Individual blocks and circuit requirements for these architectures are identified, and candidate implementations in literature are described.

For spectrum sensing, the need for analog signal processing prior to digitization to lower the total power consumption, is emphasized. Two circuit blocks: the wide-tuning frequency synthesizer for signaling, and the analog signal processor for sensing, are identified as particularly challenging circuit blocks for SDR realization. These circuit blocks are discussed in later sections.

Chapter 3 discusses the design of wide tuning range, low phase noise VCOs for use in SDR signaling applications. The challenges to wide tuning range alongside low phase noise and superior power performance are discussed, followed by an identification of a viable solution using inductor switching. Two prototype designs are discussed that provide excellent power and phase noise performance over a very wide tuning range. Measurement results of tuning range, power dissipation, and phase noise across the tuning range are presented.

Chapter 4 describes a number of novel architectures that can be utilized to achieve superior performance to even the ubiquitous cross-coupled LC VCO topology utilized in Chapter 3. These topologies are based on capacitive feedback techniques, and provide a larger voltage swing by virtue of two techniques: 1) transconductance linearization of the active devices, and 2) decoupling the tank from the active devices, thereby achieving oscillation amplitudes that are not limited by the breakdown voltages of active devices. Additionally, the feedback ensures less noise injection into the tank. Consequently, low phase noise performance is obtained. A prototype design based on this approach in a frequency synthesizer for 60GHz applications is described. Measurement results

achieving very low phase noise alongside wide tuning range are presented.

Chapter 5 explores the concept of RF sampling followed by discrete time signal processing prior to digitization for spectrum sensing applications. Specifically, frequency domain discrimination techniques are emphasized to reduce digitizing power. An example architecture is utilized for system level comparison based on analytical power estimates.

Chapter 6 discusses the design of a discrete Fourier transform based RF signal processor prior to digitization. The concept, implementation details, and handling of non-idealities in a Charge Reuse Analog Fourier Transform (CRAFT) engine prototype is described. Measurement results are presented.

Chapter 7 summarizes the thesis, and draws several conclusions regarding the design of SDR based cognitive radios.

Chapter 2

Cognitive Radio Architectures

2.1 Introduction

The architecture design for the SDR analog/RF is significantly different from that of traditional narrowband radios. In the original software radio proposal by Joseph Mitola in 1992 [7], he envisioned an architecture that digitized the RF bandwidth (no down-conversion), and performed spectrum analysis and demodulation in the digital domain. Similarly, the RF signal is synthesized in the digital domain, converted to analog and transmitted. The conceptual transceiver architecture is shown in Fig. 2.1. The Mitola architecture provides the maximum amount of flexibility through increased software capability in the digital domain. However, this architecture imposes impractical requirements on the analog-to-digital and digital-to-analog converters. For example, as discussed in [8], a 12GHz, 12-bit ADC that might be used for this purpose would

dissipate 500W of power! As a result, the ideal goal of being able to communicate at any desirable frequency, bandwidth, modulation and data rate by simply invoking the appropriate software remains far from realizable.

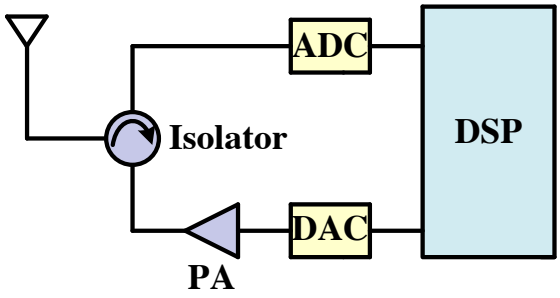

Figure 2.1: Functionality of an envisioned software radio by J. Mitola [7]

A number of SDR[1] architectures for cognitive radios currently being pursued provide only a subset of the functionality of the original software radio proposal. Specifically, in a practical interpretation, a large number of features of the waveform are defined in software, while the SDR hardware provides re-configurability to alter the waveform within certain bounds defined by the actual system. This re-configurability is commonly expected to encompass at least multiple radio access technologies (RAT) using a single transceiver IC over a wide frequency bandwidth [9]. However, in this thesis, we will focus on a much broader scope of cognitive radios that does not limit itself only to multi-standard, multi-band communications.

In order to co-exist with currently employed transceivers without causing them

[1] Note that a software-defined radio (SDR) is not the same as a software radio as prosposed by J. Mitola which relies on RF digitization. In this thesis, we reserve the term SDR for all radio architectures that provide adequate flexibility for cognitive radio functionality.

harmful interference, the SDR often needs to incorporate the most stringent specifications among all the radio access technologies being employed in the frequency range of interest.

For cognitive radio functionality, the SDR architecture can be considered a combination of a spectrum signaling transceiver system, and a sensing receiver system. Architecture options for these two systems are discussed below.

2.2 Signaling

For signaling, the cognitive radio is expected to utilize only a narrow band of frequencies [8]. However, this narrow bandwidth might occur anywhere in the entire frequency bandwidth of interest. As a result, the signaling transceiver is expected to provide high-frequency agility in conjunction with frequency selectivity over a large bandwidth. Once a particular band is selected for transmission/reception, message signal is frequency translated for communication. As observed in [8], for most civilian applications of cognitive radios, only a particular narrow band of RF frequencies are of interest at a particular time for transmission of the message signal. However, this narrow frequency band may occur anywhere along the entire bandwidth of operation of the SDR. In order to cover the entire range, multiple receiver front-ends may be used, each dedicated to its narrow band of operation [10,11]. However, it is easy to see that such an approach is power and area inefficient. Instead, it is desirable to use a wide-tuning range frequency

synthesizer in conjunction with a wideband/wide-tuning[2] front-end for covering the frequency range as shown for a receiver in Fig. 2.2. The number of such front-ends required will only be equal to the maximum number of non-contiguous frequency channels to be simultaneously used (typically assumed to be a single channel at present).

2.2.1 Receiver

An example wide-tuning receiver architecture is shown in Fig. 2.2. Each component in this architecture is discussed below.

RF bandpass filters:

For the receiver architecture, co-existence with employed technologies necessitates the use of RF bandpass filters. As shown in Fig. 2.2, these filters precede the LNA, and therefore their insertion loss directly compromises the receiver noise figure. Also, in order to remove large out-of-band blockers typical in the RF environment, high quality factors are mandated. Consequently, traditional RF architectures commonly use off-chip surface acoustic wave (SAW) filters to fulfill the front-end filtering requirements.

Unfortunately, SAW filters do not offer frequency tuning, and therefore, are not suitable for the SDR signaling architecture desired. As a result, the RF front-end filter is an important criteria for the realization of efficient SDRs. There has been significant

[2] In this thesis we make the following distinction between wideband and wide-tuning circuits. We use the term wideband for circuits that work over a large instantaneous bandwidth in RF. Wide-tuning circuits, on the other hand, function over a narrow instantaneous bandwidth, but can be tuned over a large frequency range.

research on tunable RF bandpass filters in recent years. A number of MEMS-based tunable filters, based on switched capacitors, have been developed [12–16]. A variety of varactors have also been utilized to realize frequency tuning [17–19].

More recently, driven by the vision of complete integration, a number of on-chip tuned filters are being developed. After initially working with the idea of on-chip switched capacitor LC filters [20], these filters have more recently been based on a revived idea originally published in 1960 [21] called the N-path filter. Recent work on this concept, using frequency-translation of complex impedances to a local oscillator frequency, is promising [22–26].

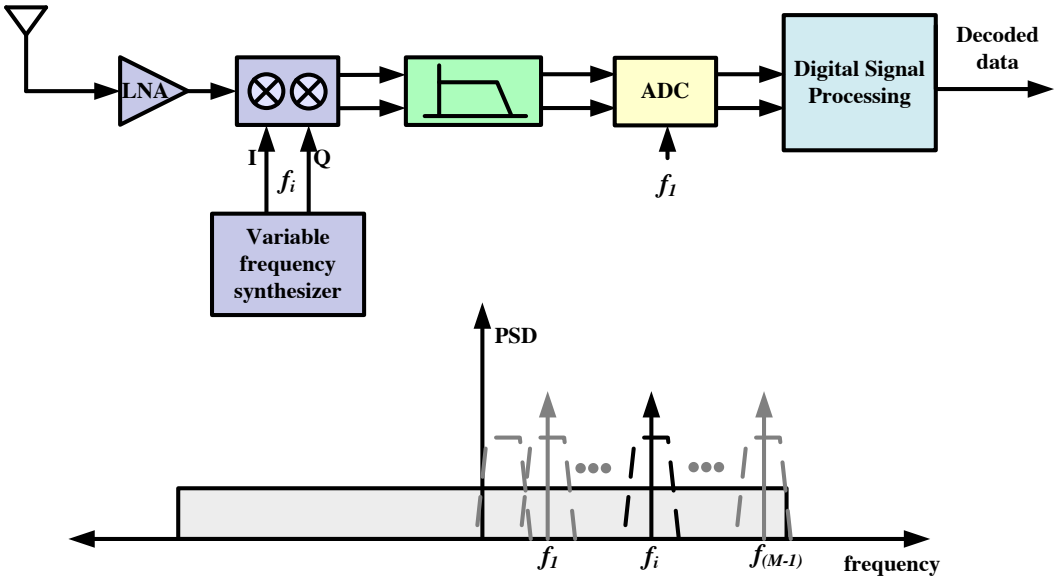

Figure 2.2: A narrowband, wide-tuning signaling/scanning approach for signal reception and decoding

LNA:

Traditionally, LNAs have been designed as narrow-band solutions to attain a required performance over a small bandwidth. For the purpose of signaling in SDRs, these architectures need to be extended to one of three possibilities:

1. *Multiple narrowband LNAs in parallel:* As discussed already, this is a wasteful idea in terms of area; moreover, it requires the implementation of RF switches that are challenging to design, and increase parasitics in the RF path.

2. *A wideband LNA:* Wideband LNAs can be designed either by preceding a narrowband design with a wideband LC-ladder filter [27, 28], or using broadband architectures [29]. The former uses multiple inductors using up valuable real-estate. The latter use MOS transistors and resistors (broadband elements) but typically suffer from a noise figure higher than 5dB when power matched at the input. Such a high noise figure does not satisfy the stringent requirements of an SDR. Using global feedback improves the noise figure, but at the expense of stability [30]. Noise canceling schemes have been proposed to improve the noise figure below 2.5dB for wideband architectures [8, 31], and are a promising candidate for SDRs.

   However, wideband architectures provide little suppression for out-of-band blockers. Consequently, their broad input bandwidth increases the linearity requirements significantly.

3. *A wide-tuning narrowband LNA:* For signaling purposes, narrowband LNAs that are tunable are the preferred option. They not only provide suppression for out-of-band blockers but are expected to provide a lower noise figure (<2dB) compared to their broadband counterparts. However, tuned LNAs have so far either provided multi-band (discrete frequency switching) operation [32, 33] or low tuning range (e.g.: 23% in [34]). For wider tuning range, switched inductors [35] may be utilized.

Frequency Synthesizer:

A wide-tuning range frequency synthesizer is a critical building block for the signaling receiver [36]. The following requirements can be identified for the synthesizer:

- Wide-tuning range, at least in excess of 67% (f to $2\times f$) such that lower frequencies can be obtained by integer division
- Low phase noise, such that the phase noise skirts do not cause down-conversion of spurious signals onto the desired signal (reciprocal mixing)
- Fast-settling behavior, to enable fast band-switching in the SDR
- Fine-frequency resolution determined by the minimum channel spacing desired

The first two requirements stem from the limitations in the voltage controlled oscillator (VCO). Unfortunately, these two performance features of VCOs: phase noise, and tuning range, strongly trade-off with one another. On-chip VCOs are based on two popular architectures: LC tank based, and ring based. Of these, LC tank VCOs are

traditionally well suited for low phase noise applications, but suffer from a low tuning range. On the contrary, ring oscillators provide a wide tuning range, but suffer from poor phase noise. Consequently, there has been a significant effort toward realizing wide-tuning range LC VCOs for SDR applications.

Varactor based tuning typically provides $10 - 15\%$ tuning range [37, 38]. Use of switched capacitors in conjunction with varactors can provide a larger tuning range, but is still limited to about 50% [39,40]. Fig. 2.3 shows the frequency tuning range and center frequency of published LC VCOs in the last decade. As seen, only one VCO (out of about 100) covers the tuning range required (66.67%) to obtain all lower frequencies by division. Moreover, the tuning range requirements are made more challenging by PVT variation.

In Chapter 3, a new technique for obtaining a very wide-tuning range (up to 160%) in LC VCOs while providing low phase noise throughout the tuning range is discussed. The oscillators described are based on switched inductors, and provide a viable option for the SDR signaling architecture.

Mixers:

For complete on-chip solutions, homodyne architectures have been the preferred choice. Homodyne architectures ease the requirements on the IF filter; however, they suffer from a number of issues including flicker noise, dc offsets, and second order non-linearity in

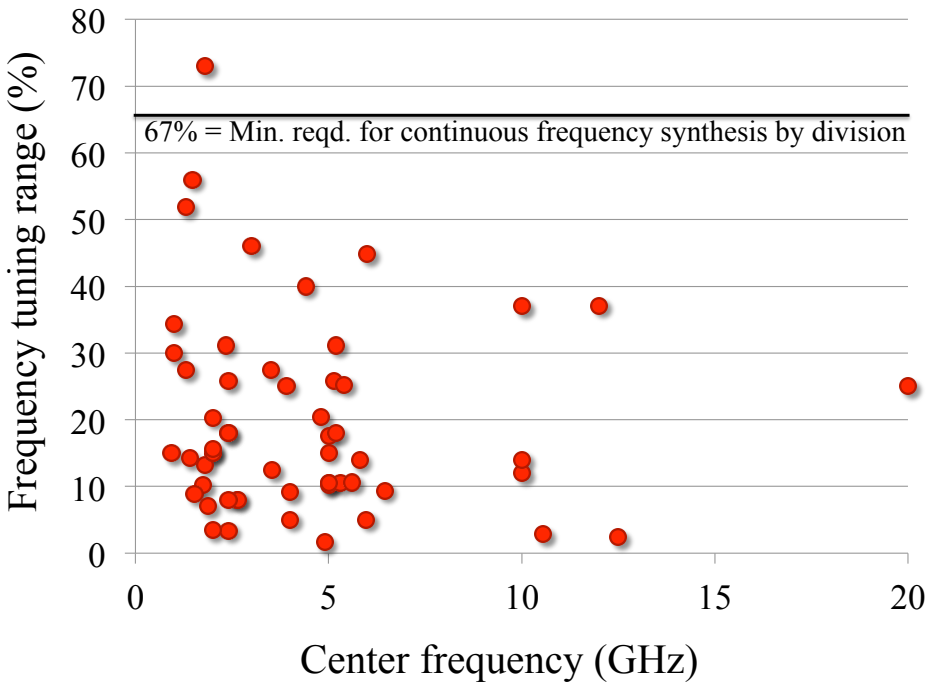

Figure 2.3: A survey of frequency tuning range of LC VCOs published in IEEE journals and conference proceedings appearing between 1992 and 2010

circuits[3] . A high-pass filter may be used to filter out the low frequency components; however, this contributes to information loss and is not the preferred solution for very narrow-band architectures (e.g. GSM). Consequently, homodyne architectures have gained popularity for moderate to wide bandwidths, especially in SDR applications. Due to their lower flicker noise and excellent linearity, passive mixers have often been the preferred choice for SDR applications [8, 41].

Considering that the RF tuned filter being used does not provide as effective out-of-band rejection as fixed RF filters, and that SDR applications often require more stringent out-of-band rejection specifications [41], harmonic reject mixers become critical. To understand the problem of harmonics, consider the scenario shown in Fig. 2.4(a). Despite the use of a variable bandwidth, band-select filter, large out-of-band blockers are not sufficiently suppressed. Subsequent mixing with a square wave (most mixers perform hard-switching) down-converts the desired signal, but also the interferers around the LO harmonics as shown in Fig. 2.4(b). Consequently, the IF (zero-IF in this case) signal is significantly corrupted as shown in Fig. 2.4(c).

To alleviate the problem of harmonic mixing, multi-phase harmonic rejection mixers can be employed [25]. Such mixers have been employed in a number of SDR architectures [8, 41], providing 30-40 dB harmonic rejection. Further harmonic rejection can be obtained using digital correction [41].

In addition, mixer first architectures [42], and a linear LNTA followed by a passive

[3] In case a single mixer is used, it also suffers from LO pulling, LO re-radiation, reciprocal mixing, etc. However, using a multi-step down-conversion mitigates these effects.

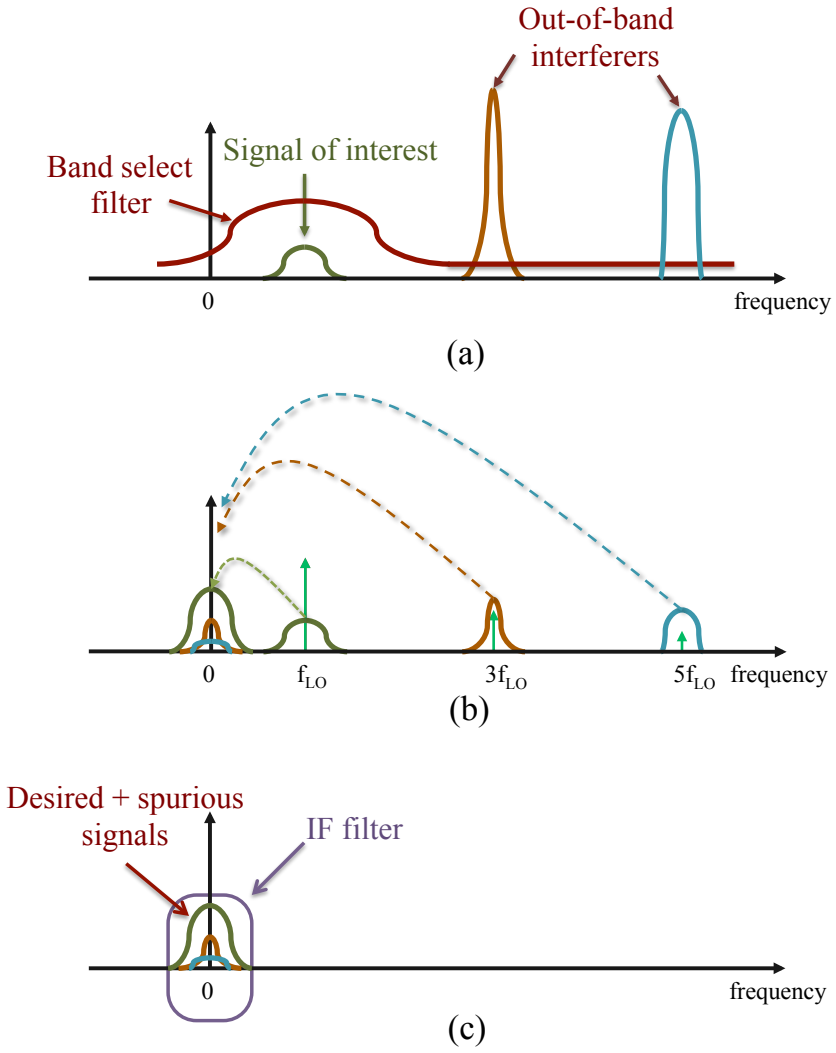

Figure 2.4: The problem of harmonic down-conversion in a receiver

mixer [41] architectures have gained popularity due to their improved linearity and subsequent capacity to handle large out-of-band interference.

Baseband Filters:

Prior to digitization by an ADC, a variable bandwidth analog baseband filter is required to attenuate interferers to reduce the ADC dynamic range. Additionally, the baseband filter acts as a variable pole anti-aliasing filter (AAF) prior to sampling. Various techniques can be used to implement the baseband filter, broadly classified into continuous time and discrete time varieties. Continuous time filters have the advantage of not suffering from aliasing issues, but typically have poorer accuracy due to matching limitations. These filters can be realized using op-amps or using Gm-C filters [43–45]. Programmability can be achieved by switching resistors and capacitors, or by tuning these elements. However, the accuracy of these filters in sub-micron processes, and their lack of easy programmability have made way for active research in discrete time variations for SDR applications.

Discrete-time baseband filters are based on switched capacitors and can be implemented with or without active elements [8,46–49]. The ease of programmability in these filters arises from the dependence of the filter notches on the clock frequency which can be tuned easily. Note that discrete-time filters require an anti-aliasing pre-filter that may be constructed using a combination of mixer output poles, current domain sampling, etc.

Analog to digital converter (ADC):

ADCs for signaling applications typically require a high dynamic range to accommodate a variety of access technologies, and the remaining in-band blockers. Additionally, a relatively large (tens of MHz) bandwidth is required to accommodate wide-band access technologies.

For high dynamic range applications, $\Sigma-\Delta$ ADCs [50] have been the popular choice. $\Sigma-\Delta$ converters are based on $\Sigma-\Delta$ modulators that perform noise shaping on an oversampled signal using negative feedback techniques. The feedback loop suppresses the non-idealities in the quantizer, improving the effective resolution despite the use of a lower resolution quantizer. The shaped noise is filtered in the digital domain to obtain the ADC bits; as a result of the large amount of digital processing involved in the $\Sigma-\Delta$ ADC, these converters scale well with technology. Typically, $\Sigma-\Delta$ converters have been bandwidth limited due to quantizer latency, and parasitic poles, and sampling speeds have been limited to 1GHz for 70dB dynamic range. This results in bandwidths in the tens of MHz. However, in recent years, wideband $\Sigma-\Delta$ architectures that provide both the high dynamic range of $\Sigma-\Delta$ ADCs, as well as wideband digitization have been explored [51, 52]. These circumvent the speed limitations to provide signal bandwidths above 100MHz alongside a high (>70dB) dynamic range.

For obtaining even larger bandwidths (usually at the expense of dynamic range), Nyquist ADCs are the more popular choice. Pipeline ADCs are the most popular choice for increasing the input bandwidth (>200MHz with $14-16$ bit resolution) [53, 54]. For

increasing the dynamic range, calibration techniques are popularly used [55]. However, the use of high-gain, wideband amplifiers and complex calibration techniques increase their power consumption, area, and design complexity.

Recently, successive approximation register (SAR) Nyquist ADCs have been used for moderate dynamic range (8−11 bit) at moderate speeds (<100MHz). SAR ADCs provide the advantage of very high power efficiency (low FOM) as they use just a single comparator for comparison [56–58]. However, they use N steps for providing an N-bit resolution, limiting their inherent speed. To overcome the speed limitation, time-interleaving is popularly used with the SAR architecture. Additionally, the SAR ADC resolution is limited by gain, offset, timing errors, and capacitor matching, and calibration techniques can be used for overcoming these [59].

A plot of the dynamic range and bandwidth of published ADCs [60] is shown in Fig. 2.5.

2.2.2 Transmitter

A generic transmitter architecture for SDR signaling applications is shown in Fig. 2.6. The critical blocks in this architecture comprise a wide-tuning frequency synthesizer, similar to one discussed in Section 2.2.1, a wideband mixer, similar to one discussed in Section 2.2.1, and a wideband, linear power amplifier (PA). Linearization techniques for wideband SDR architectures is discussed in [61]. Harmonic mixing [25] may be used; moreover, the mixer and power amplifier may be combined to obtain a linear power

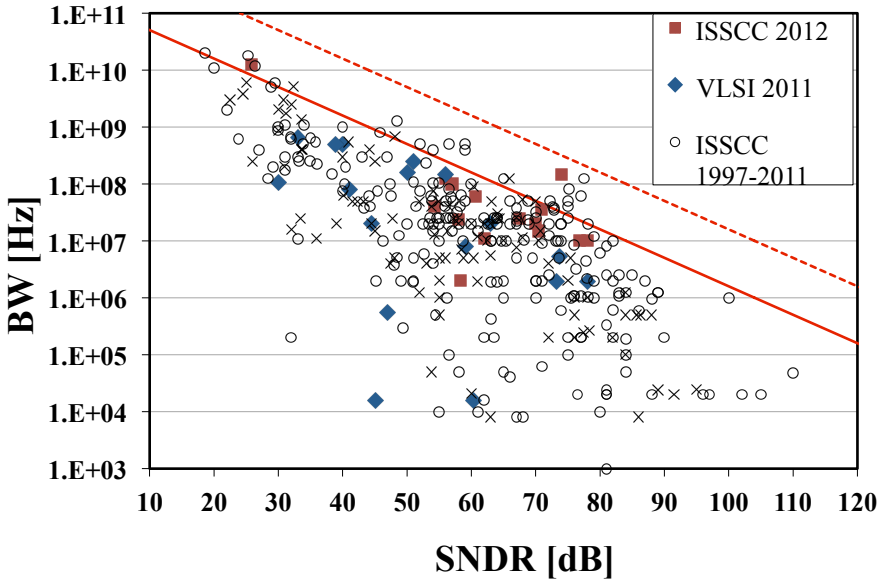

Figure 2.5: A survey of ADCs from 1997-2012 [60]

mixer [62].

2.3 Spectrum Sensing

Spectrum intelligence or knowledge about the spectral and spatial electromagnetic spectrum around us is of critical importance for dynamic spectrum access in cognitive radios. One of the primary bottlenecks in implementing dynamic spectrum access is reliably detecting primary users before deciding on the frequency, channel, access and modulations formats of transmission for the secondary user. The problem of primary user detection is further complicated by multi-path fading and hidden node / shadowing effects [3]. Some of these problems can be mitigated somewhat by using cooperative techniques across multiple sensors [63]. However, the issue of identifying primary user

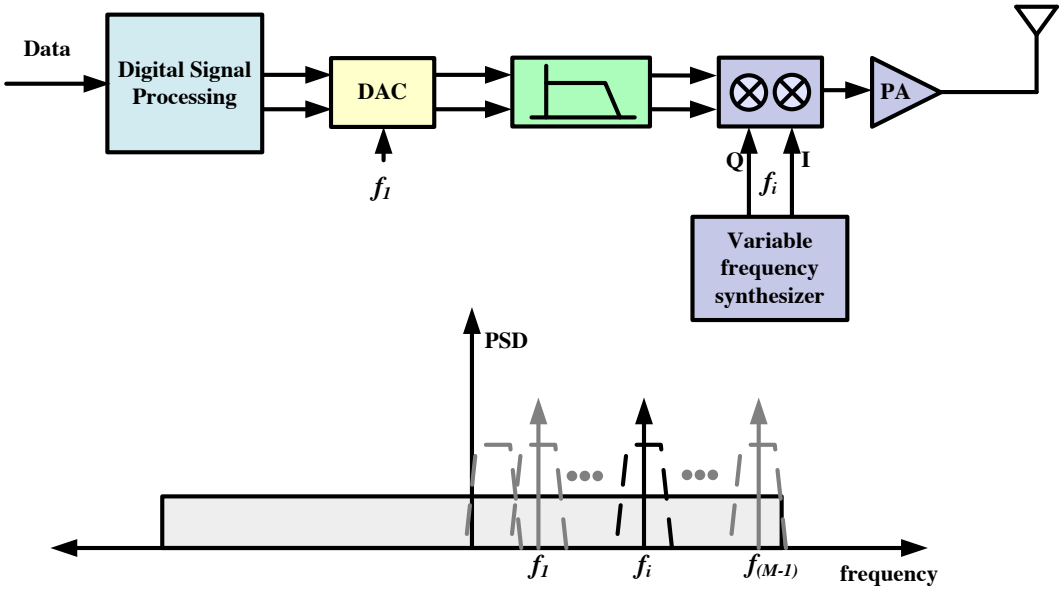

Figure 2.6: A narrowband, wide-tuning approach for signal transmission

activity with varying modulation formats and signal power levels, over a wide swathe of spectrum, using limited power continues to confound the industry.

From an SNR perspective, the optimal method for signal detection is a matched filter. However, such coherent detection techniques require *a priori* knowledge of the primary user signal and modulation format. Worse still, a cognitive radio based on coherent detection with matched filters would require a dedicated radio for each primary user class. One method to address this problem is to employ a blind, non-coherent detection scheme using energy detection. Despite the implementation simplicity of this approach, non-coherent energy detection schemes require a large number of samples $[O(1/SNR^2)]$ and therefore, a longer sensing time. Moreover, such a scheme would be unable to distinguish noise and in-band interference from primary user signals, or work

for spread spectrum signals (direct sequence and frequency hopped). An intermediate approach would be to use feature detectors which rely on some knowledge (or feature) of the primary user. Cyclostationary feature detectors rely on detecting the built-in periodicity of modulated signals such as their carrier frequency, bit rate, repeated spreading, pulse strains, hopping sequences, and pilots. Moreover, even in the absence of clock timing or phase knowledge, cyclostationary feature detectors obtain considerable improvement in detection performance [64]. A brief description of the spectral correlation function on which cyclostationary feature detectors are based is provided in Appendix A.

Spectrum Sensing Architectures:

There are two fundamentally distinct options for realizing a spectrum sensing receiver for an SDR front-end: a scanner type, and a wide instantaneous bandwidth digitizer type

1. Scanner: In this scheme, a narrowband, wide-tuning receiver scans and digitizes the entire bandwidth (similar to a spectrum analyzer) for analysis. The digital backend processes each band sequentially and stitches the frequency domain outputs to obtain a spectral map of the environment. The architecture used for this scheme is very similar to that used in the signaling receiver discussed above (Fig. 2.2. However, in order to overcome issues such as multi-path, fading, hidden nodes, interference problems, etc [3, 65], the sensitivity and dynamic range

requirements of the architecture are more challenging than the signaling scheme. Moreover, note that sensing is a blind detection problem, as opposed to signaling where *a priori* knowledge of the transmitted signal is available.

Although the scanning architecture is able to re-use much of the signaling architecture (or vice versa), this detection technique suffers from multiple short-comings. These systems lack the agility to be able to detect any fast-hopping signals. Frequency domain stitching is power hungry in the digital domain due to the phase distortion of the analog filters. Moreover, stitching the frequency domain information from several scans is imperfect in the face of multi-path; consequently, signals spanning across multiple scan bandwidths are imperfectly reconstructed. Due to these and other reasons, it is desirable to construct a real-time instantaneous bandwidth digitizer (similar to J. Mitola's original software radio idea) in the spectrum sensor.

2. Wide instantaneous band digitizer: Unlike the scanning type architecture, a wide instantaneous band digitizer is expected to digitize the entire wide RF bandwidth instantly, and perform the digitization in a real-time fashion. Understandably, the wideband digitizer has widely been considered as the bottleneck to the realization of the SDR based cognitive radio. A number of efforts in recent years have focused on wider bandwidths, broadband matching, higher front-end linearity, and most importantly, wideband analog to digital converters.

Several architectures have been proposed for the RF front-end. Of these, the most popular is the extension of the traditional receiver architecture as shown in Fig. 2.7 effectively performing an RF to digital conversion (R-to-D). As shown, the front-end requires a wideband LNA prior to the digitizer. Moreover the front-end needs to handle a very large dynamic range due to the large peak to average power ratio (PAPR) of wideband signals. The increase in PAPR for wide bandwidths is described in Fig. 6.4. As shown, the PAPR for the narrowband signals is only 2, while that for the wideband signal (5 times the bandwidth) with multiple signals, all having similar powers, is 10. As a result of the large PAPR of the wideband inputs, a very linear front-end is required. The linearity requirements of the LNA have been addressed in [66–68]. Another approach using an LNTA followed by mixers is discussed in [41]. Moreover, passive mixer first topologies have been proposed for high IIP_3 performance [42].

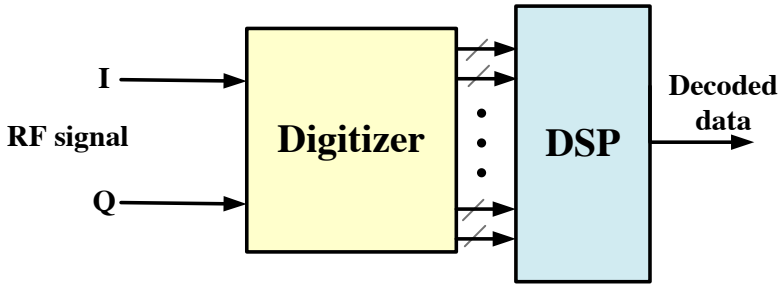

Figure 2.7: A wideband RF to digital conversion architecture for spectrum sensing

The digitizer block shown in the figure is essentially an ADC with performance specifications beyond that capable using state-of-the-art converters discussed in

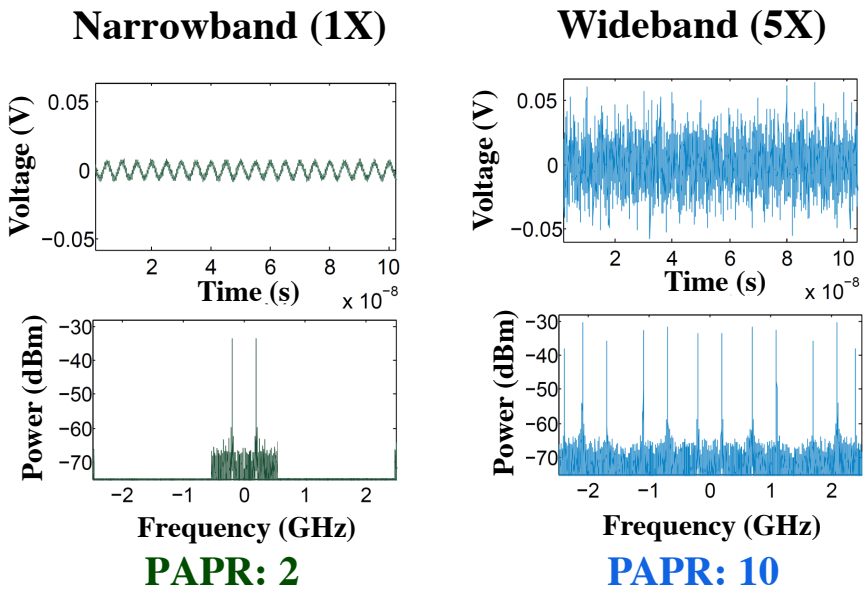

Figure 2.8: PAPR increase in wideband signals compared to narrowband signals

Section 2.2.1. This wideband digitizer can be implemented in multiple ways, all based on some form of multiplexing in order to ease the requirements on the ADCs. A multiplexed broadband approach using time-interleaving as shown in Fig. 2.9 was proposed in [69]. This scheme reduces the sampling rate of ADCs. However, all the ADCs still see the full bandwidth, and therefore, still require high dynamic range capability. Also, the sample and hold circuitry remains difficult to design.

In order to reduce the dynamic range requirements on the ADCs, it is possible to transform the signal to a different domain prior to digitization [70]. Specifically, a frequency domain transform is particularly attractive [71]. A frequency domain transform can be approximated in practice using band-pass filters for channelization, as proposed in [72]. This reduces the dynamic range requirements of

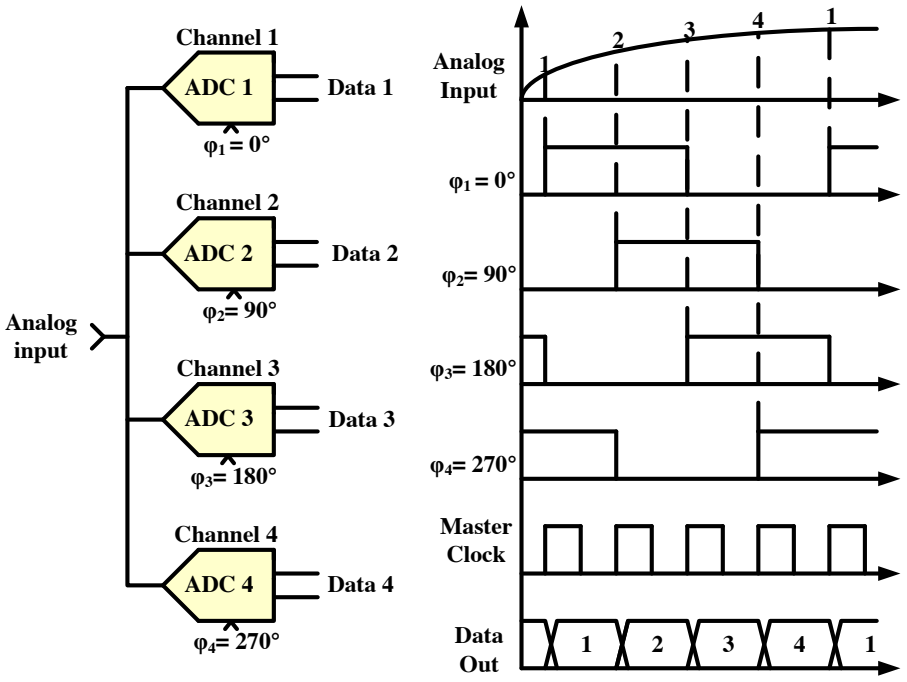

Figure 2.9: Time interleaved ADCs for broadband channelization

the ADCs but introduces the problem of designing impractically sharp band-pass filters. In [73], replacing sharp band-pass filters by frequency down-converters followed by sharp low-pass filters eliminates this problem as shown in Fig. 2.10. However, these are based on PLLs, mixers and low-pass filters [74], or on injection locked oscillators [75][4] , and can be power hungry. Moreover, harmonic mixing of signals within the SDR input bandwidth severely corrupts the channelized base-band signals. Additionally, signal reconstruction from the digitized filter-bank outputs is challenging.

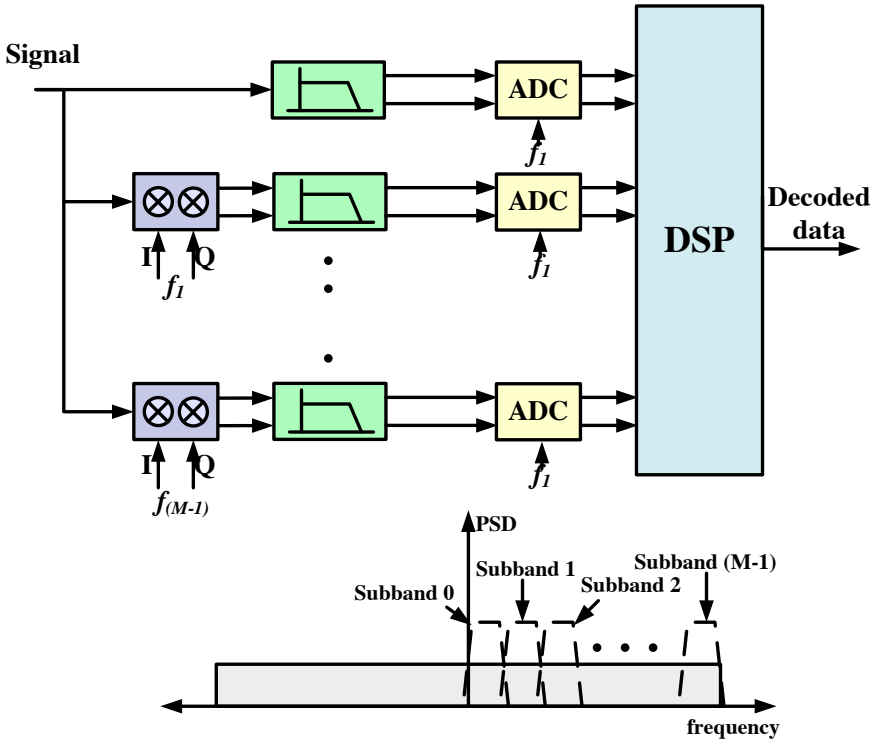

Figure 2.10: Low pass filterbank approach for channelization

[4] Note that injection locked oscillators have the advantage of a larger noise suppression bandwidth (≈ lock range) [76] and provide better reciprocal mixing robustness compared to PLLs (assuming the reference phase noise is better than the VCO phase noise).

In this thesis, we propose a digitizer approach based on frequency discrimination using an analog domain DFT followed by ADCs (Fig. 6.2). This technique enjoys the advantages of analog domain frequency discrimination: channelization and consequent ADC dynamic range reduction. In addition, it also enjoys an ultra-low power implementation, signal spreading benefits, no harmonic mixing problem, and simple reconstruction in the digital domain. Details on the architecture and implementation is included in Chapter 6.

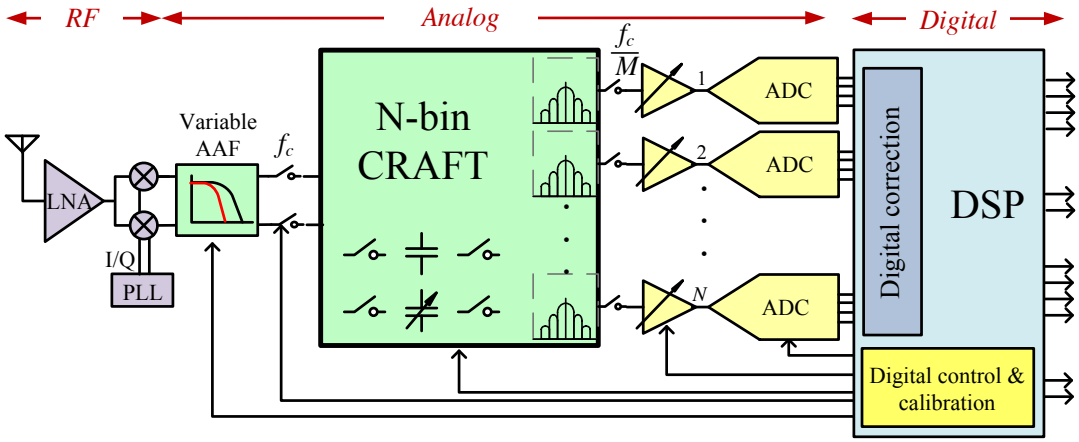

Figure 2.11: An envisioned SDR architecture enabled by CRAFT

2.4 Conclusions

In this chapter, several software defined radio architectures for cognitive radio applications were discussed. The SDR system was divided into two functional blocks: signaling and sensing. Novel architectures for each of the blocks were reviewed. Additionally, new circuit topologies that are potential candidates for satisfying the required specifications

were reviewed. Specifically, wideband, and wide-tuning range, and programmable circuits for RF and baseband filters, low noise amplifiers, mixers, frequency synthesizers, power amplifiers, and analog to digital and digital to analog circuits were discussed. Existing architectures for tackling the extremely challenging problem of wideband spectrum sensing was explored, and a new architecture for the same was proposed.

Chapter 3

Wideband Voltage Controlled Oscillator

3.1 Introduction

In this chapter, we explore techniques to increase the tuning-range of VCOs while maintaining adequate phase noise and power dissipation performance over the tuning range. Two popular oscillator architectures: LC tank based, and ring based, are considered for SDR signaling and spectrum scanning applications. Of these, LC tank VCOs are traditionally well suited for their superior phase noise and low power consumption at radio frequencies, and are therefore the preferred choice. However, LC tank VCOs are notorious for their lower tuning range as compared to ring oscillators. In this work, we

select the LC tank oscillator and devise a scheme based on switched inductors, and capacitor array optimization, to extend the desirable power and phase noise properties of LC VCOs over a wide tuning range for use in SDR signaling (and sensing) applications.

3.2 Frequency Tuning Options in LC Tank VCOs

The frequency of oscillation of an LC tank is given by:

$$f_0 = \frac{1}{2\pi\sqrt{LC}} \tag{3.1}$$

Frequency variation can be realized by changing the capacitance and/or the inductance. For wide tuning ranges, traditional tuning by varying the capacitance with a fixed inductance is disadvantageous. For such wide ranges, switched-inductor resonators offer distinct advantages [77–80] as discussed below. The popular cross-coupled LC VCO topology has been used for the discussion.

3.2.1 Tuning range

Theoretical Analysis The tuning range of LC oscillators can be estimated using Barkhausen's startup condition as shown below[1] :

$$\begin{aligned} & g_m R_p \geq 1 \text{ from the startup condition} \\ \Rightarrow \quad & g_m Q^2 R_s = g_m \frac{\omega^2 L^2 R_s}{R_s^2} = g_m \frac{L}{R_s C} \geq 1 \end{aligned} \tag{3.2}$$

[1] Here, we assume that the tank Q is dominated by the Q of the inductor at our frequencies of interest.

where g_m is the transconductance of the active circuitry, and R_p is the equivalent parallel resistance of the tank. To the first order, the parasitic resistance, R_s, of the inductor is proportional to its inductance, L, and the expression reduces to

$$\frac{g_m}{C} \geq K_L \tag{3.3}$$

where K_L is the constant of proportionality, i.e. $R_s = K_L L$.

Insight As seen from the analysis above, the tuning range of capacitively tuned oscillators is limited by the parasitic capacitance (say, C_{min}) at the highest frequency, and the startup criterion ($\frac{g_m}{C_{max}} = K_L$) derived in (3.3) at the lowest frequency. This range of capacitances, C_{min} to C_{max}, determines the tuning range that can be obtained in capacitively tuned oscillators:

$$f_{min} = \frac{1}{2\pi\sqrt{LC_{min}}} \text{ to } f_{max} = \frac{1}{2\pi\sqrt{LC_{max}}} \tag{3.4}$$

However, if a different inductance is used with the same variable capacitors (C_{min} to C_{max}) in parallel to it, a new set of f_{min} and f_{max} are obtained from equation (3.4), and therefore a new frequency range can be covered. This is because, assuming a constant K_L, equation (3.3) remains valid for all values of $C_{min} \leq C \leq C_{max}$ even for the new inductor. Therefore, switching between different inductance values enables the oscillator startup criteria to be fulfilled over different sets of frequency ranges. *Consequently, for*

non-overlapping bands, the total tuning range obtained becomes the sum of the tuning ranges of each frequency band obtained from switching between inductance values.

Additionally, inductance variation ensures that the tuning range is no longer theoretically limited (by the startup condition for instance). A practical limit is imposed by the range of inductance values implementable without degrading the L/R_s ratio. Consequently, the tuning range can be increased significantly, limited only by the constraints in the physical implementation of the specific variable inductance scheme used. This provides a convenient means for obtaining extremely wide tuning ranges in LC VCOs.

3.2.2 Phase Noise

Theoretical Analysis The approximate phase noise of an LC oscillator is derived from Leeson's model[2] in (3.5). In resonators dominated by the inductor Q, and for $\omega_0 \gg Q\Delta\omega$, equation (3.5) reduces to (3.6).

$$L(\Delta\omega) = 10\log\left[\frac{2FkT}{P_{sig}}\left\{1+\left(\frac{\omega_0}{2Q\Delta\omega}\right)^2\right\}\right] \tag{3.5}$$

$$\approx 10\log\left\{\frac{kTFR_s}{2V_{rms}^2}\left(\frac{\omega_0}{\Delta\omega}\right)^2\right\} \tag{3.6}$$

Here $\Delta\omega$ is the frequency offset, F is the excess noise factor, R_s is the inductor's parasitic series resistance, P_{sig} is the signal power, ω_0 is the oscillation frequency, and Q is the inductor's quality factor.

[2] Note that this equation holds even if the active devices (and not the resonator) are the main contributor towards the phase noise in the circuit [81].

To the first order, R_s is proportional to L, and the expression approximately reduces to (3.7)

$$L(\Delta\omega) \approx 10 \log \left\{ \frac{kTFK_L L}{2V_{rms}^2} \left(\frac{\omega_0}{\Delta\omega} \right)^2 \right\} \tag{3.7}$$

where K_L is the constant of proportionality.

Insight As seen in the analysis above, the phase noise becomes *directly proportional* to $\omega^2 L$ as $\frac{kTFK_L}{(\Delta\omega V_{rms})^2}$ is a constant. For traditional LC VCOs with constant inductance and employing capacitive tuning, the phase noise is proportional to ω^2. Consequently, even if the phase noise of the oscillator is carefully optimized for the lowest frequency, it is bound to degrade at the higher frequencies as shown by the solid blue line in Fig. 3.1(a). However, from (3.1), for a constant capacitance, $\omega^2 L$ is a constant. Therefore, if the resonator can be tuned using inductance tuning with a constant capacitance, the phase noise would remain constant with frequency (to the first order) as shown by the dotted green line in Fig. 3.1(a). Since inductance tuning is not realizable in bulk-CMOS processes, we use a combination of inductance switching and capacitance tuning to obtain a suitable compromise as shown by the dashed red line in Fig. 3.1(a). Consequently, the worst case phase noise shows considerable improvement for switched-inductor resonators.

The improvement in phase noise by the introduction of a switch might at first seem counter-intuitive. For an LC tank dominated by the inductor Q, introducing a switch

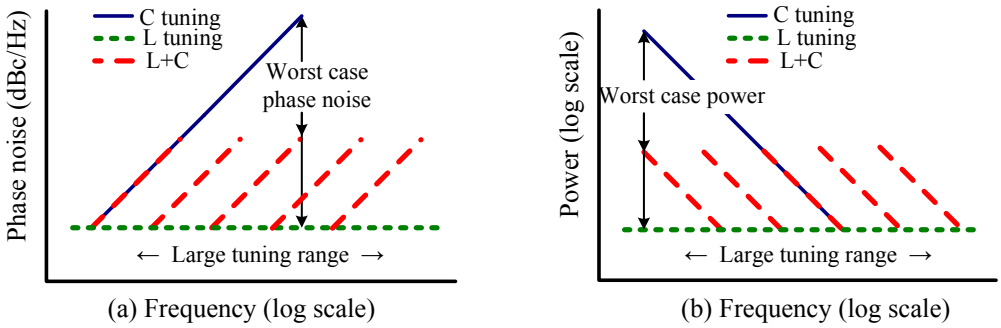

Figure 3.1: VCO power and phase noise for alternate frequency tuning options

degrades the quality factor of the tank. However, as expressed in equation (3.6), the phase noise expression depends on the parasitic series resistance of the inductor, rather than its quality factor. Therefore, as long as the parasitic resistance introduced by the switch is lower than the parasitic resistance of the inductor switched out, there is a net phase noise improvement due to the switching.

Literature Survey In this regard, it is instructive to review the past literature on LC VCOs to gain further insight based on our analysis. Fig. 3.2 shows the phase noise versus frequency of LC VCOs operating between 900MHz and 10GHz published in IEEE journals and conferences between 1992 and 2010. Assuming that the VCOs were optimized at each frequency for overall performance with complete freedom over the choice of a suitable inductance, the plot is expected to roughly signify inductance tuning over the frequency range. Consequently, a flat overall trend is expected (instead of the commonly accepted 20dB/decade thumbrule). The trendline shows a 8.8dB/decade increase with frequency. The deviation from a flat trend can be attributed to a number

of factors including:

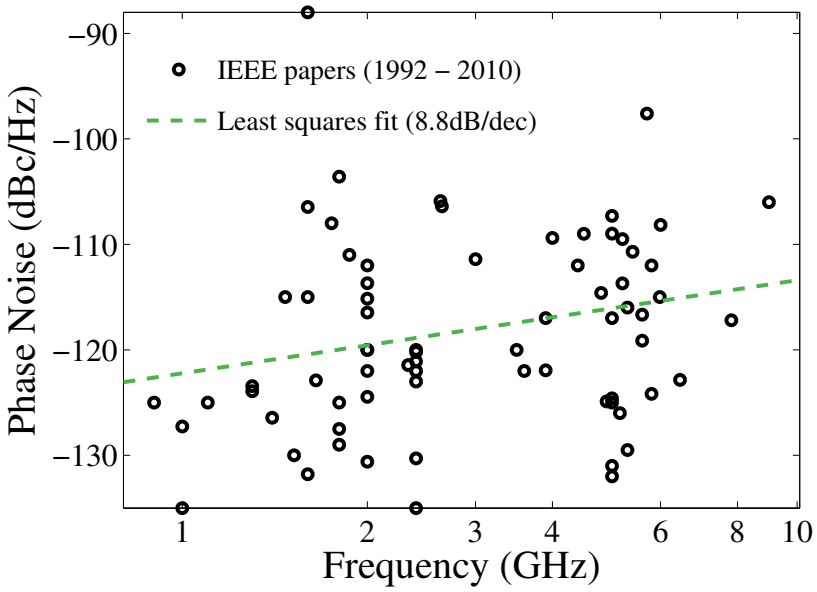

Figure 3.2: A literature survey of LC VCO phase noise versus frequency (log scale) in IEEE journals and conference proceedings appearing between 1992 and 2010

1) Design requirements: The phase noise requirements in the different standards are relatively more relaxed at the higher frequencies, encouraging optimization of some other performance specifications (eg. power dissipation) at the cost of phase noise.

2) Higher order effects: Skin and proximity effects at higher frequencies degrade the phase noise at the higher frequencies.

3) Implementation issues: Due to implementation issues, designs at higher frequencies are not necessarily optimized to use the smallest inductance possible. Smaller inductors in IC technologies tend to have lower quality factors as a result of lower number of turns and/or smaller footprints. Small inductors also suffer from lower modeling accuracy and

larger process variation. Additionally, they are less robust to startup in the face of process variation. As seen in equation (3.7), writing $\omega^2 = \frac{1}{LC}$, we note that the phase noise is inversely proportional to the capacitance. Therefore, using a larger inductance than is optimal implies a smaller capacitance at a particular frequency, and consequently, a degradation in phase noise.

These effects therefore combine constructively to degrade the phase noise slope with frequency as observed in literature.

3.2.3 Power

Theoretical Analysis The power dissipation, for a constant voltage swing (V_{sw}), can be derived as,

$$\begin{aligned} V_{sw} &\approx R_p I_{bias} = Q^2 R_s I_{bias} = \frac{(\omega L)^2}{R_s} I_{bias} = \frac{(\omega_0 L)^2}{R_s} \frac{Power}{V_{DD}} \\ \Rightarrow Power &= \frac{R_s V_{sw} V_{DD}}{(\omega_0 L)^2} \end{aligned} \tag{3.8}$$

where V_{sw} is the oscillator signal swing, and I_{bias} is the bias current in the oscillator. Again, to the first order, R_s is proportional to L, and therefore, the power dissipation becomes *inversely proportional* to $\omega^2 L$ as given by

$$Power = \frac{K_L V_{sw} V_{DD}}{\omega^2 L} \tag{3.9}$$

where K_L is the constant of proportionality.

Insight For the traditional capacitive tuning scheme using a constant inductance, the power is inversely proportional to ω^2. Consequently, even if the power dissipation is optimized at the highest frequency, it is bound to degrade at the lower frequencies as shown by the solid blue line in Fig. 3.1(b). Again, in the case of pure inductance tuning the power dissipation remains constant (to the first order) versus frequency as shown by the dotted green line in Fig. 3.1(b). The tuning technique using switched inductance and variable capacitance provides a suitable compromise between these two methods with improved power dissipation, in comparison to the traditionally used fixed inductor, variable capacitance scheme, as shown by the dashed red line in Fig. 3.1(b). Therefore, as seen through the simple analysis above, *inductor switching provides a combination of improved phase noise and improved power alongside a wide tuning range providing an optimal basis for wideband applications.*

As seen from the first order analysis above, inductance switching in VCOs provides a powerful technique for significant increase in the tuning range of LC VCOs without affecting their phase noise and power performance. Consequently, this technique can be used for extending the optimization frameworks developed for LC VCOs ([82], [83]) to wide tuning ranges without adversely affecting their performance.

Literature Survey Again, a review of the past literature on LC VCOs is insightful. Fig. 3.3 shows the power dissipation versus frequency of LC VCOs operating between 900MHz and 10GHz published between 1992 and 2010. As explained above, the

plot is expected to approximately signify inductance tuning over the frequency range, and consequently, a flat overall trend is expected (instead of the commonly accepted −20dB/decade thumbrule). The trendline shows a 0.06dB/decade decrease with frequency. In this case, the same factors, described while discussing the phase noise trend, reappear. However, instead of adding up, their effects oppose and nullify each other. These effects are further explained below:

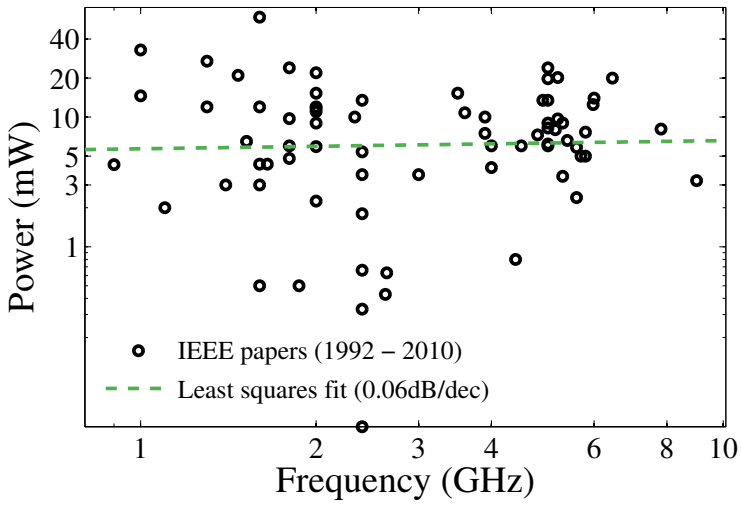

Figure 3.3: A literature survey of LC VCO power versus frequency (log log plot) in IEEE journals and conference proceedings appearing between 1992 and 2010

1) Design requirements: As discussed earlier, the relaxed phase noise requirements in the relatively higher frequencies allows optimization of power dissipation at these frequencies. So power dissipation reduces slightly with increasing frequency.

2) Higher order effects: Skin and proximity effects at higher frequencies increase the power dissipation at the higher frequencies.

3) Implementation issues: As discussed earlier, designs at higher frequencies are not necessarily optimized to use the smallest inductance possible. As seen in equation (3.9), writing $\omega^2 = \frac{1}{LC}$, we note that the power is directly proportional to C. Therefore, using a larger inductance than is optimal implies a smaller capacitance at a particular frequency, and consequently, a reduction in the power consumption.

These opposing effects therefore nullify each other to provide a relatively flat power dissipation slope with frequency.

3.3 Resonator Synthesis

We consider the problem of trying to meet a certain tuning range and a phase noise specification over that range while staying within a fixed power budget. In this section, we develop a simple scheme for designing a resonator, using a switched-inductor if necessary, to cover the required frequency range. As described in Section 3.2, apart from a larger tuning range, inductor switching can also provide phase noise and power benefits. In this paper, we will design the resonator with a primary emphasis on the tuning range. We will then verify the worst case power dissipation and phase noise, to the first order, against a given set of specifications in Section 3.4.1, and modify the design, if necessary, in order to meet these specifications. This provides a starting solution with the minimum number of inductor switches, and is therefore not only the easiest to work with, but also provides a solution fairly close to the optimum possible. This will be further elaborated using a specific design example in Section 3.4.1. In

this design example, we shall assume that we start with a large inductor and short out portions of it to reduce the inductance. However, it is possible to think of starting with a large inductor and then using a smaller portion of it, i.e., leaving the excess open, as an alternate method to reduce inductance.

For the present design, let the tuning range specification for the oscillator be f_{min} to f_{max}. We start with a resonator comprising an inductor L_1, a switched capacitor bank, and a varactor. Let these combine to provide f_{min} as the minimum frequency of oscillation. For this, the inductor and capacitor bank are selected, such that, for the smallest inductor, Barkhausen's magnitude criterion for oscillation is just met at the lowest frequency of interest, f_{min} [82]. Next, we estimate the total capacitive parasitics of (a) the varactor, (b) the switched capacitor bank when switched off, (c) the transconductance cells, and (d) the inductor. Let the estimate of these combined parasitic capactiances be called C_{par}. Using this estimate, we calculate the maximum frequency that can be obtained using this resonator: $f_{1,max} = 1/(2\pi\sqrt{L_1 C_{par}})$. If the maximum frequency obtainable, $f_{1,max}$ is greater than f_{max}, we have fulfilled the requirements of the desired resonator. If this frequency is less than f_{max}, an inductor switch needs to be introduced.

We now find a second inductor value, $L_2 < L_1$ to cover the remaining frequency range. L_2 can be obtained by switching out a portion[3] of the inductor L_1. L_2 is calculated as the smallest inductor that provides a minimum frequency of oscillation

[3] Although a switched inductor example is used in this section, this analysis is valid for other techniques that may be used to vary the tank inductance.

$f_{2,min} < f_{1,max}$ with some frequency overlap. The frequency overlap is designed to account for the capacitive parasitics of the inductor switch, the uncertainty in determining the parasitics, and process and temperature variation. For obtaining $f_{2,min}$, L_2 and a part of the capacitor bank is selected[4] so that Barkhausen's criteria for oscillation are just met at this frequency. Now, using L_2 and C_{par}, the maximum frequency of oscillation $f_{2,max}$ is determined. If $f_{2,max} > f_{max}$, the resonator design is complete. The switched out portion of the inductor equals $L_1 - L_2$, and this denotes the position at which the switch needs to be added. In case $f_{2,max}$ is not high enough, a second switch needs to be added. We find another inductor value $L_3 < L_2$ following a similar procedure as had been used for determining L_2. We check if $f_{3,max} > f_{max}$, and, if not, we continue adding inductor switches until the entire frequency range is covered. The design methodology described above is summarized for a single switch case in Fig. 3.4.

3.4 Prototype I

In this section, we will focus on a design example for building further insight, discussing implementation issues, and developing a prototype design as a proof of concept. For this, we use a frequency floorplan spanning 3.3-8.3 GHz and target a CMOS only implementation in the IBM 130nm SiGe BiCMOS technology. For our particular application, since phase noise and power are the important criteria being considered, the

[4] It is assumed that switching the inductor lowers its quality factor. Therefore, for the same g_m, Barkhausen's criteria can no longer be fulfilled by the entire capacitor bank in combination with the switched-inductor (eqn. 3.2).

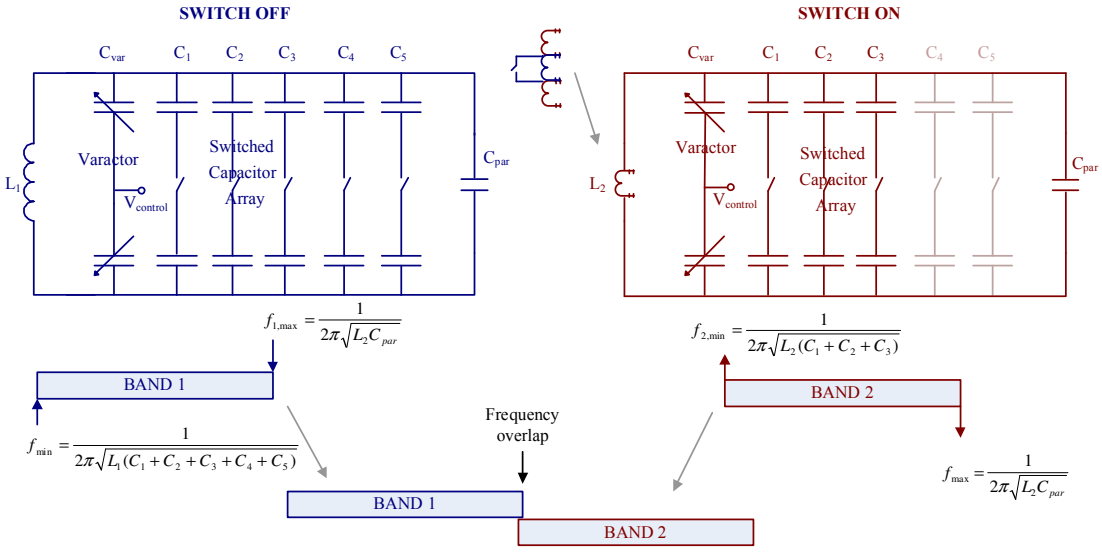

Figure 3.4: Summary of design methodology described in Section 3.3

cross-coupled nMOS-pMOS topology in Fig. 3.5 is an appropriate choice [82].

3.4.1 Simple Model for Design Space Exploration

In accordance to our approach in Section 3.3, we begin the design by calculating the minimum inductance that enables us to obtain $f_{min} = 3.3GHz$. The design of the VCO core has been explored in detail for this topology [82], [83], and we base our design on the same framework. The inductance L_1 is fixed at 1200pH, and a 5 bit capacitor array is designed accordingly. $f_{1,max}$ is found to be $6.1GHz$ from circuit simulation. Since this value is less than the required $f_{max} = 8.3GHz$, we add a switch to the inductor. An nMOS transistor switch is considered by virtue of its lower on-resistance, and an estimate of the parasitics of this switch is obtained from circuit simulation. Now, a simplified model of the switched-inductor resonator is constructed as shown in

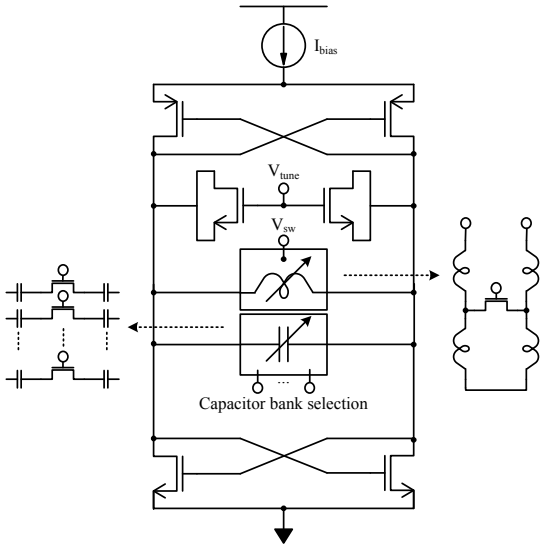

Figure 3.5: Cross-coupled nMOS pMOS topology used for the prototype design

Fig. 3.6(a). Using the switch, two frequency bands are obtained as shown in Fig. 3.4. Here L_1 and L_2 are the effective inductances when the switch is off and on respectively, C_{par} is the minimum capacitance attainable, and $C_1, C_2, ..., C_5$ form the capacitor bank, as shown in the figure. We also define a switching ratio $k = (L_1 - L_2)/L_1$, where(ideally) $L_1 - L_2$ is the fraction of the inductor switched out.

We begin with simple models for the switch in Fig. 3.6(a) based on its region of operation. When off (transistor in cut-off), the switch can be modeled as a capacitor ($C_{sw} \approx (C_{gd} + C_{db})/2$), as shown in (Fig. 3.6(b)). When on (transistor in triode), it can be modeled as a resistor in parallel with the switched out part of the inductor, as shown in Fig. 3.6(c). Since the switch widths being considered result in a very low on-resistance, the parasitic capacitance appearing in parallel with this small resistance can be ignored when the transistor is on (operating in the triode region). Using parallel

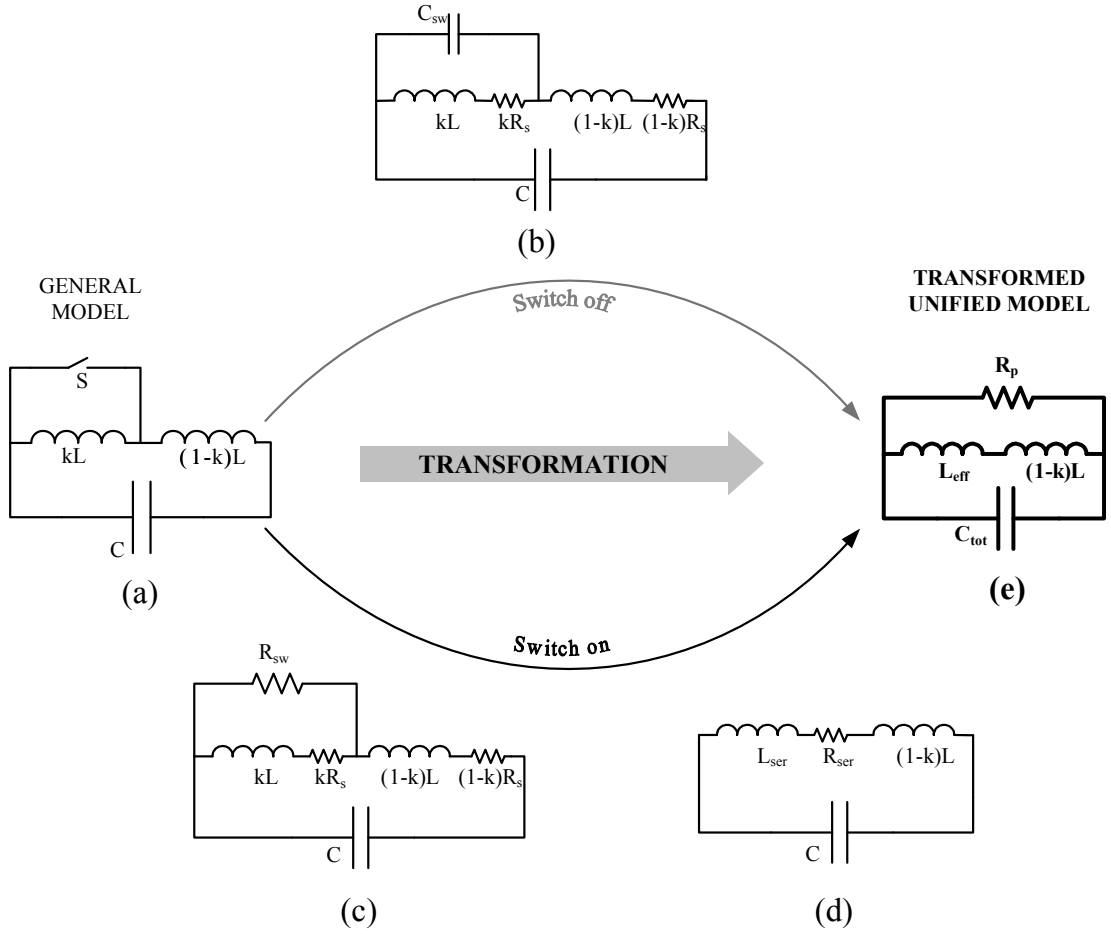

Explanation of variables used			
(b)		(e)	
C_s	Switch parasitic capacitance when off	**On case**	
R_s	Inductor parasitic resistance	L_{eff}	$\approx L_{ser}$
		C_{tot}	C
(c)		R_p	$(1+Q^2)R_{ser}$
R_{sw}	On resistance of switch	**Off case**	
		L_{eff}	$\approx kL$
(d)		C_{tot}	$\approx C + kC_s$
R_{ser}	$(1-k)R_s + R_{sw} - \frac{R_{sw}^2(R_{sw}+kR_s)}{(R_{sw}+kR_s)^2+(\omega kL)^2}$	R_p	$(1+Q^2)R_s$
L_{ser}	$\frac{R_{sw}^2 kL}{(R_{sw}+kR_s)^2+(\omega kL)^2}$		

Figure 3.6: Modeling of the switched resonator, and parallel series transformations: (a) Simple model for switched resonator (b) Resonator when the switch is off (c) Resonator when the switch is on (d) All parasitic resistances absorbed into Rs when the switch is on (e) Parallel RLC equivalent for both switch on and off cases

to series transformations, we can absorb the switch parasitics into the inductor branch. An equivalent parallel RLC tank can then be constructed using series to parallel transformations of the combined inductor branch parasitics (Fig. 3.6(a) to Fig. 3.6(e)). As evident, the inductor switch introduces parasitics that impact the performance. This comes as a price paid for the wider frequency range obtained, and as we will see, the technique is not beneficial for narrowband tuning.

Now, from an exploration of the design space using this simple model we will determine: 1) the range of possible k values based on the tuning range specifications, 2) for what values of k, switching the inductor is beneficial from a phase noise view-point, and 3) an approximate optimal switch size.

We plot the variables $f_{1,max}$ (marked ideal), $f_{2,min}$, and $f_{2,max}$, as described in Section IV, versus varying k in Fig. 3.7(a). Since a particular switch width has not been determined yet, we analyze the tuning range with an ideal switch ignoring the capacitive parasitics for the time being. Later, we will revise the analysis with the switch capacitance included in our calculations. Based on our specifications, $f_{2,min}$ has to be lower than $f_{1,max}$ which provides an upper limit of 0.58 for k (*condition* 1) as shown in Fig. 3.7(a). Also, we obtain a lower limit of 0.47 for k (*condition* 2) in order to ensure $f_{2,max} > 8.3GHz$ as shown in Fig. 3.7(a).

Now from equation (3.6), observe that for resonators dominated by the inductor Q, the phase noise reduces with a reduction in the effective series resistance, R_{ser} (as defined in Fig. 3.6(d)). Fig. 3.7(b) shows a plot of R_{ser} vs. k, calculated analytically,

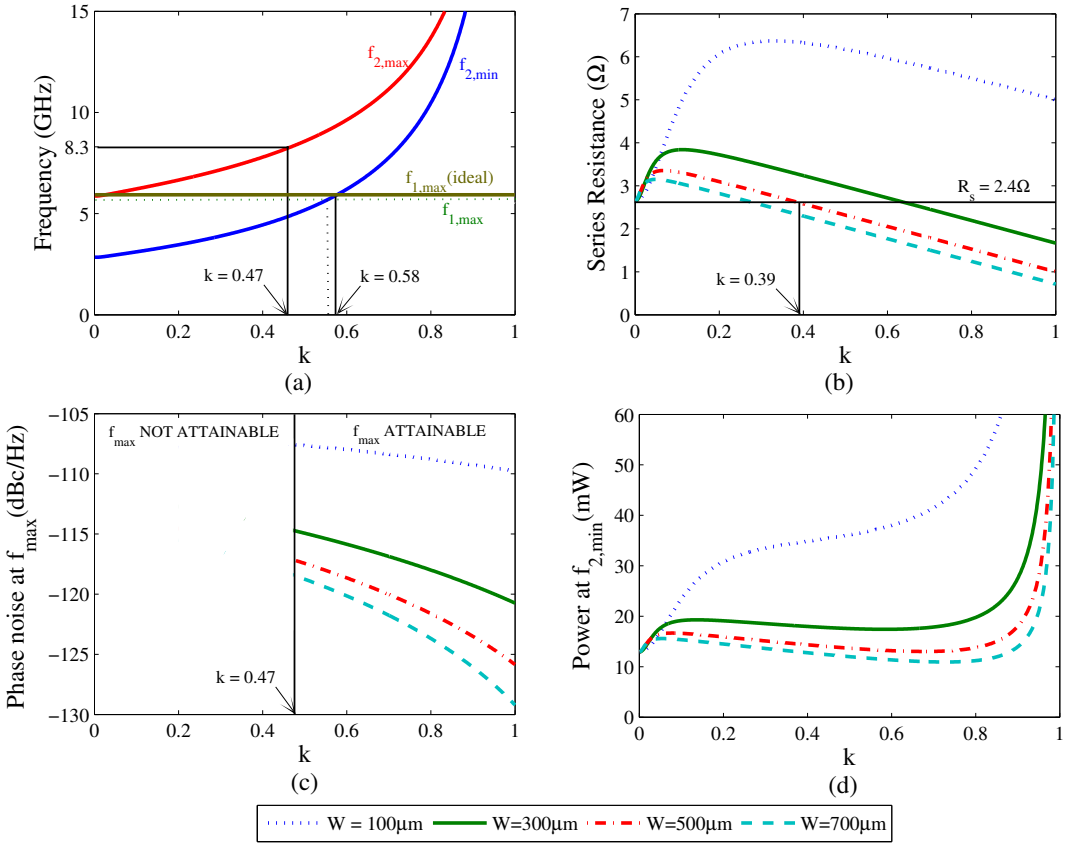

Figure 3.7: Calculated values vs. switching ratio (k) and nMOS switch widths (W for a constant L = 120nm) (a) Frequencies $f_{1,max}$, $f_{2,min}$, and $f_{2,max}$ which provides a lower and an upper limit for k based on the given specifications, (b) R_{ser} (see Fig. 3.6(d)); here r(Fig. 3.6(d)) = 2.64Ω, (c) Worst case phase noise calculated at f_{max}, and (d) Worst case power dissipation calculated at $f_{2,min}$

in the higher frequency band for different device widths. It is required that R_{ser} be lower than the effective parasitic resistance of an un-switched-inductor ($R_s = 2.64\Omega$) so as to obtain a phase noise benefit from switching. Note from Fig. 3.7(b) that a phase noise benefit cannot be obtained using a switch of width 100μm. For a switch of width 500μm, a phase noise benefit will be obtained for values of $k > 0.39$ as shown. This criterion provides another lower limit for k (*condition* 3).

Let us now plot the worst case phase noise and power dissipation for varying k and switch widths based on our simplified models. We had derived an expression for phase noise in (3.6). Due to the ω^2 term in the expression, the worst case phase noise occurs at f_{max}. The phase noise at f_{max} is plotted in Fig. 3.7(c). Our criterion for obtaining f_{max} (*condition* 2) is also incorporated in this figure, and the phase noise is plotted only for values of k for which *condition* 2 is fulfilled. Similarly, we had derived the power dissipation in (3.8). Note that the worst case value occurs either at $f_{1,min}$ or $f_{2,min}$. Since, in practical implementations, the inductor Q inevitably degrades on the addition of a switch, the worst case power performance is expected to occur at the frequency $f_{2,min}$. The power dissipation at $f_{2,min}$ is plotted in Fig. 3.7(d).

Both these plots (Fig. 3.7(c, d)) suggest an improvement in performance with an increasing value of k. Also, we note that the performance benefits from increasing the transistor width diminish beyond 500μm. However, the parasitic capacitance continues to increase linearly with increasing width, reducing the maximum frequency obtained. Also, in a practical implementation, using a larger switch causes a degradation of the

inductor quality factor due to various higher order effects as will be introduced in Section 3.4.2. Based on these considerations, we identify an approximate width of 500μm as the optimum transistor width for switching.

Using this value of switch transistor width, we re-estimate a more realistic value of $f_{2,min}$, and plot it in Fig. 3.7(a) (shown using dotted lines). As a result, *condition* 1 is revised to $k < 0.55$. Also, *condition* 3 is fixed at $k > 0.39$ for a transistor width of 500μm. Based on these three *conditions* obtained in our example, we identify $0.47 < k < 0.55$ as possible values that fulfill the given tuning specifications, and also provide a phase noise benefit from switching. Also, from Fig. 3.7(c, d), using the largest k within the given range provides the best power and phase noise performance. However, a lower k value would safeguard a continuous tuning range in the face of PVT variations by ensuring overlap of the two frequency bands. Therefore, both these factors need to be considered when selecting a particular value of k for a given application.

3.4.2 Proposed Switched Inductor Implementation

The new switched-inductor implementation as first proposed in [77] is shown in Fig. 3.8. An octagonal inductor with a single switch positioned within and underneath it is designed. For the switching ratio ($k = (L_1 - L_2)/L_1$) obtained in Section 3.4.1, the inductor switch has to be physically connected between the arms of the inductor's inner coil. A 500μm wide nMOS transistor (low resistance) is used as the switch, and is positioned directly underneath and within the inductor coil. When the switch is off,

the entire inductor is included in the circuit. When the switch is on, the inner coil is switched out, and only the outer coil inductance is included. This provides the two inductance values L_1 and L_2 required for our design.

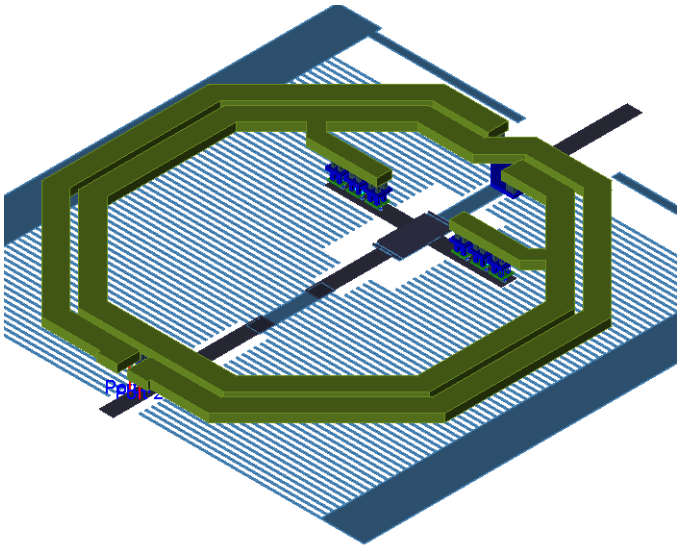

Figure 3.8: 3-D view of the switched inductor with the switch shown at a lower tier in the center of the inductor

As compared to placing the switch outside the spiral, this scheme reduces the effects of interconnect parasitics significantly. Higher order electromagnetic (EM) effects of placing the transistor underneath the inductor were considered, and the switch is placed away from the traces in order to reduce unpredictable capacitive coupling effects which dominate over magnetic effects at the frequencies of interest. Also, in this switching scheme, the relatively lower quality inner coil [84] of the inductor is switched out, as compared to previous designs [85], thereby promising a better phase noise performance. The large switch size adds parasitic capacitance to the LC tank. However, by placing

the switch close to the midpoint (small signal ground) of the differential inductor, the effect of these parasitics is considerably reduced. Further, a high resistance substrate underneath the inductor in conjunction with a patterned ground shield [86] are used to improve the quality factor of the inductor.

The Effect of Mutual Coupling

So far, we have not considered the mutual inductance between the two inductor coils. However, in the practical implementation shown in Fig. 3.8, mutual coupling causes the net inductance of the inductor to increase when the switch is turned off. The total inductance in this case is $L_{coil_1} + L_{coil_2} + k'\sqrt{L_{coil_1} L_{coil_2}}$ where L_{coil_1} and L_{coil_2} are the self inductances of the individual coils and k' is the coupling coefficient between them. When the switch is turned on, only the outer coil is functional. However, a parasitic current that effectively opposes the field due to the outer coil is induced in the shorted inner coil. This current reduces the net inductance of the switched-inductor below L_{coil_1}, and degrades its quality factor. The effective inductance and series resistance of the switched-inductor can be calculated from electromagnetic simulations. These can then be included in the analysis discussed in Section 3.4.1 above to incorporate higher order effects for greater accuracy.

The degradation of the inductor quality factor from the parasitic current in the inner coil can be avoided by placing another large switch (Fig. 3.9) at the center of the differential inductor[5] . A second prototype was designed based on this concept, and is

[5] Note that since this device is placed at the center of the differential inductor (a.c. ground), the

discussed in Section 3.5.

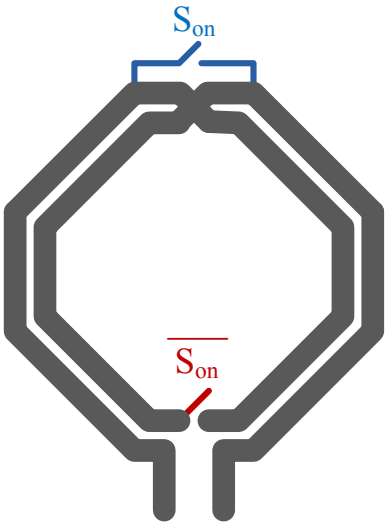

Figure 3.9: A scheme to eliminate the parasitic current in the inner inductor loop for a switched inductor

3.4.3 Detailed Prototype Design and Simulation

Using the analysis above, approximate values for the switch size and switching ratio k were obtained. In this section, we will describe the design and efficient simulation of resonator components for the prototype design based on the values of switch size and k.

parasitic capacitance of this transistor does not adversely affect the tuning range. Consequently, it is possible to make this device large enough to have negligible resistance.

Switched Inductor Simulation

The switched spiral inductor is shown in Fig. 3.8. The inductor was built as described in Section 3.4.2. For selecting a particular physical location on the inductor to correspond to the switching ratio k determined in Section 3.4.1, 4-port electromagnetic(EM) simulations were performed using ADS Momentum®. Two ports were placed at the two physical ends of the symmetric inductor, while the other two were placed at the positions where the switch was to be inserted. The switched out inductor value was then calculated from the 4-port S-parameter data obtained. The position of the switching point (represented by two out of the four ports) was then swept till the desired switched-inductor value of L_2 was obtained.

For the purpose of circuit simulations, an S-parameter 4-port network was used to mimic the switched-inductor. For this, the resultant S-parameter data from the EM simulations in ADS Momentum® were exported to the Cadence Virtuoso® design environment for circuit simulations. This provided a simple and efficient way to include EM effects for the custom-made inductor in circuit simulations.

Switch Implementation

The switch was implemented using a wide nMOS transistor within and underneath the inductor coil, as described in Section 3.4.2. The trade-off between the switch on-resistance and off-capacitance was analyzed in Section 3.4.1, and an approximate value of 500μm was obtained for the nMOS switch width. Simulations were performed using

the S-parameter black-box technique described in the previous subsection (3.4.3), and based on these the switching ratio k, as well as the switch width, were fine-tuned and finalized as 0.5 and 500μm respectively.

G_m cells

The W/L ratios of the g_m transistors were designed based on [82], but allowing for a maximum current of 10mA. This allows start-up at the lowest frequencies of interest in the two bands, but causes the oscillator to function in the current limited regime (non-optimal phase noise as discussed in [82]). Such a design technique offers two advantages: it reduces the parasitic capacitance offered by the g_m transistors, and it limits the worst-case power dissipation of the VCO. Although this causes a degraded phase noise performance at the lower frequencies, the overall worst case phase noise is still exhibited at the highest frequencies of each band (Fig. B.4). Therefore, this degradation does not affect the overall phase noise performance of the VCO. In effect, the lower phase noise at the lowest frequencies is traded off for obtaining an improved power performance and tuning range.

Capacitor Bank

For the capacitor bank, a 5-bit, binary-weighted, differential, switched, MIM capacitor array was constructed for coarse frequency tuning. For fine frequency tuning, two differentially connected MOS varactors were utilized. Apart from enabling a larger

tuning range, the combination of discrete and continuous tuning also ensures a desirable, low value of K_{VCO} for the oscillator. The switched capacitor array was constructed using multiple unit cells to provide good matching and ensure a monotonic frequency tuning characteristic.

3.4.4 Measurement Results

The design was implemented using CMOS transistors only in the IBM 0.13μm SiGe BiCMOS process. Fig. 3.10 shows a die photograph of the VCO. The total area including the bondpads is 0.87mm^2, and the area of the VCO core is 0.1mm^2.

Measurement results for the prototype design are discussed below. A table comparing the performance of this VCO with other recent publications is shown in Table 3.3.

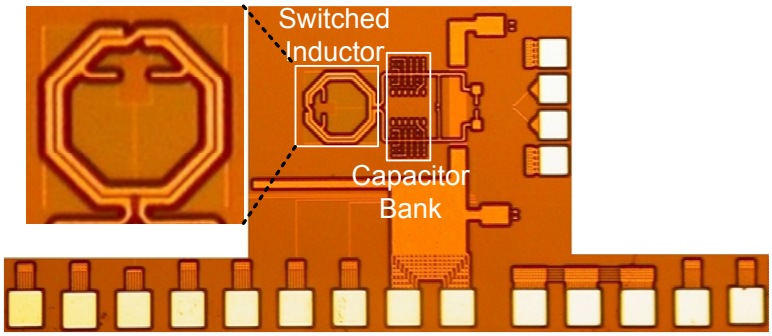

Figure 3.10: Die photograph of wide tuning range VCO prototype

Tuning

For measuring the tuning range, the leakage output was coupled out by an off chip antenna to avoid the effect of loading from the 50Ω buffer and the probe pads. Using the combination of discrete and continuous tuning described above, the VCO achieves a frequency tuning range (FTR) of 87.2%, from 3.28–8.35 GHz at room temperature. The FTR increases with a decrease in temperature to 88.6% (3.28–8.52 GHz) at 0°C. As seen in Table 3.3, the FTR obtained betters other previously reported single inductor wideband VCO solutions. When probed, the frequency tuning range reduces to 83.9%, from 3.2–7.82 GHz, due to the loading by a sub-optimally designed probe buffer, but is still better than prior designs.

Fig. 3.11 plots the frequency tuning of the VCO versus the capacitance added, where x is the capacitance of a unit capacitor in the capacitor bank. The frequencies are plotted on the y-axis against the capacitance on the x-axis. Note that all the binary weighted capacitors are not used for switching the frequency in the high frequency band. This is because, in the high band, the VCO fails the start-up condition for the larger capacitors due to the reduction of the inductor quality factor.

The two frequency bands obtained by inductor switching are shown, and achieve a 384MHz overlap (along the y-axis in Fig. 3.11). A considerable frequency overlap when switching capacitors is also ensured through the use of suitable varactors. This overdesign ensures a continuous frequency tuning in the face of PVT variations.

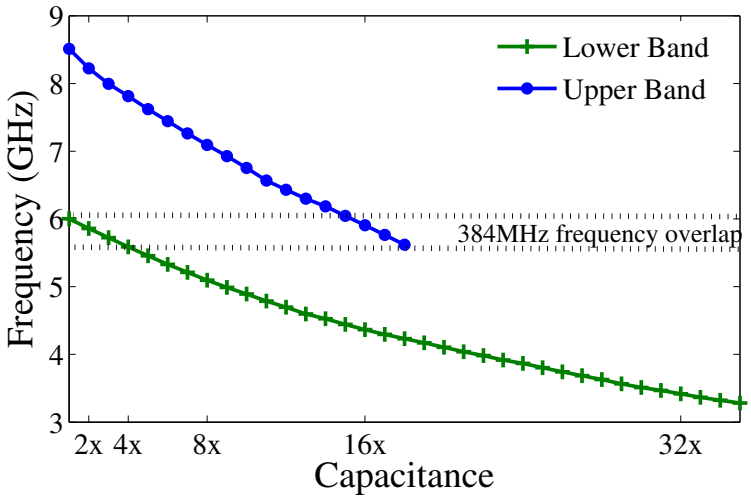

Figure 3.11: Measured tuning range versus capacitance in terms of the 'x', where 'x' is the capacitance of an unit switched capacitor

Power Dissipation

The measured variation in the power dissipation with frequency (log scale) is plotted in Fig. 3.12. The power requirement of the VCO, as predicted for a constant voltage swing, reduces as frequency increases in each individual band. The measured power dissipation for the core VCO varies from 6.6 to 14.1 mW in the upper band, and 6.5 to 15.4 mW in the lower band from a 1.6V supply. The trends expected in the variation of power dissipation with frequency are evident (compare with Fig. B.5). The improvement in worst case power dissipation through inductor switching, as was discussed in Section 3.2, is also obtained.

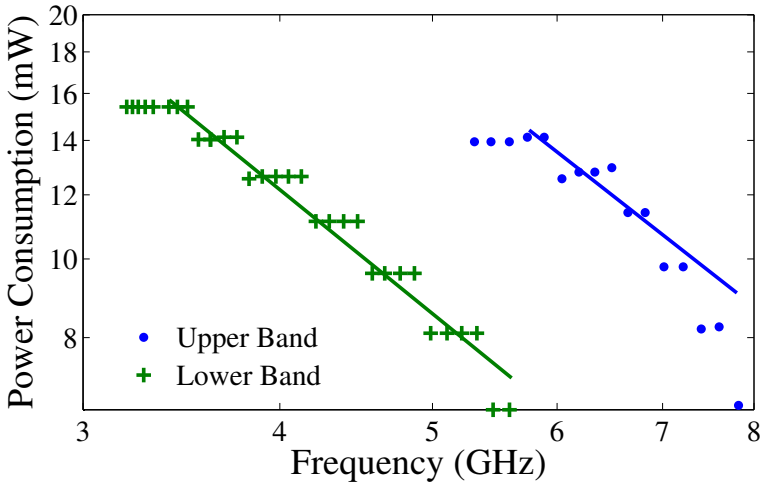

Figure 3.12: Measured power dissipation versus frequency (log log plot)

Phase Noise

For phase noise measurement, a buffer was included on chip in order to drive a 50Ω environment. The probed differential outputs were converted to a single-ended signal compatible with the measurement apparatus using an off-chip balun. Phase noise measurements were performed using an HP E4407B spectrum analyzer. A screenshot of the phase noise measurements at the lowest frequency point (3.28GHz) is shown in Fig. 3.13. The variation of the phase noise with the frequency (log scale) is shown in Fig. 3.14. The phase noise varies between −122 and −117.5 dBc/Hz at 1MHz offset in the lower frequency band, and −119.6 and −117.2 dBc/Hz at 1MHz offset in the higher frequency band. The trends expected in phase noise variation with frequency is evident (compare with Fig. B.4). Again, an improvement in the worst case phase noise performance through inductor switching, as was discussed in Section 3.2, is also

obtained.

In order to evaluate and compare VCO performances, the power-frequency-tuning normalized figure of merit (FOM_{PFTN}) as described in [82] was used.

$$FOM_{PFTN} = 10\log\left\{\frac{kT}{P_{dc}}\left(\frac{f_{tune}}{f_{off}}\right)^2\right\} - L(f_{off}) \tag{3.10}$$

where f_{off}, P_{dc} and $L(f_{off})$ denote relative offset frequency from the carrier, power dissipation and phase noise respectively.

A temperature of 300K was used for the FOM_{PFTN} calculations. Fig. 3.15 shows the VCO performance over the entire frequency range. The FOM_{PFTN} varies between 6.6 and 10.2 dB in the low frequency band, and 6.7 and 9.5 dB in the high frequency band. As shown in Table 3.3, this design provides one of the best tuning range performances reported to date. As seen from the excellent FOM_{PFTN} obtained, by using the switched-inductor technique, the phase noise and power are not compromised for obtaining this tuning range. Also notable is the lack of significant variation in the FOM_{PFTN} values over such a wide tuning bandwidth. This was ensured through a carefully designed trade-off between phase noise and power across the entire tuning range.

3.5 Prototype II

In this section, we describe another switched inductor VCO, using three inductor switches to extend the VCO tuning range to 158%. The VCO schematic is identical to that shown in Fig. 3.5.

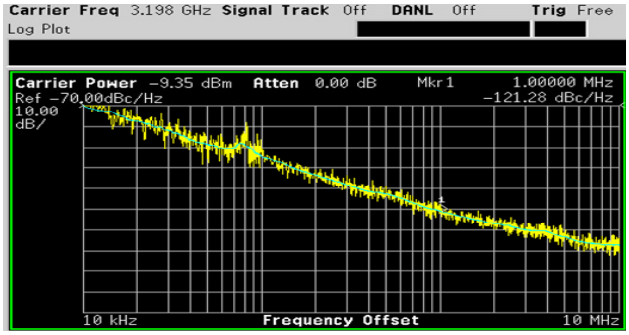

Figure 3.13: Measured VCO phase noise across offset frequencies from a 3.28GHz carrier

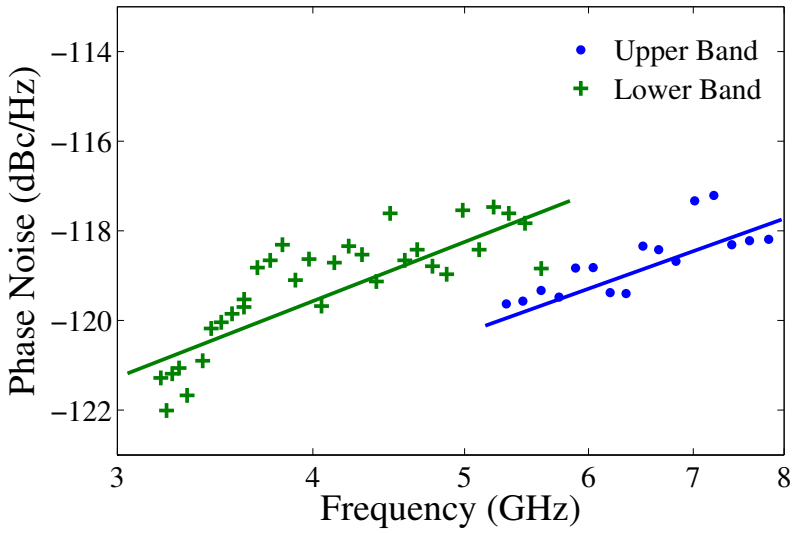

Figure 3.14: Measured phase noise at 1MHz offset versus frequency (log scale)

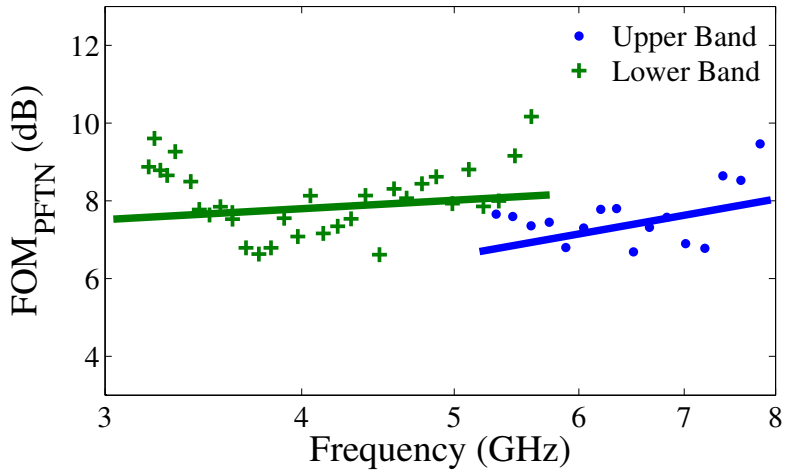

Figure 3.15: FOM_{PFTN} versus frequency (log scale)

3.5.1 Switched Inductor Implementation

The inductor is implemented with 3 switches, as shown in Fig. 3.17. As shown in the figure, only one of the switches are on at any given time. Depending on which of the switches is on, three different inductances are obtained. The switched spiral inductor is shown in Fig. 3.5.1. For selecting a particular physical location on the inductor to correspond to the switching ratio k, 4-port electromagnetic(EM) simulations were performed using Integrand EMX® [87] (see Fig. 3.16).

Two ports were placed at the two physical ends of the symmetric inductor, while the other ports were placed at the positions where the switch was to be inserted. The switched out inductor value was then calculated from the 8-port S-parameter data obtained. The position of the switching points (represented by six out of the eight ports) was then swept till the desired switched-inductor values of L_2 and L_3 were obtained. As

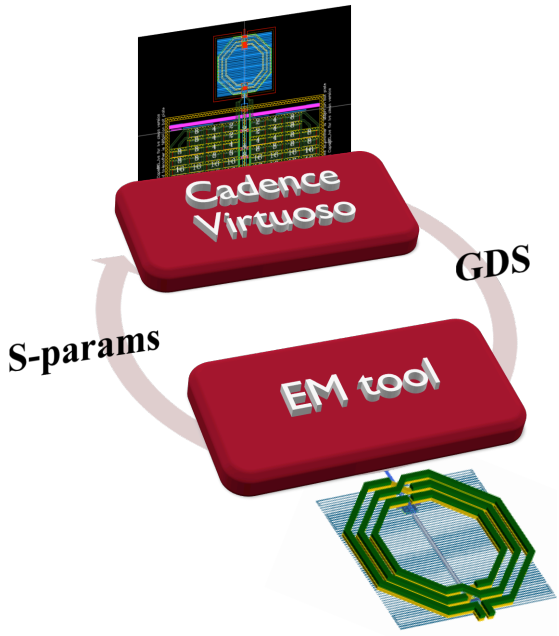

Figure 3.16: Summary of partial EM simulation techniques used for VCO design

shown, an attempt was made to design the inductor such that the switches are placed underneath the traces, (unlike in Section 3.4.2) for ease of implementation.

Simulations were performed by importing S-parameters from Integrand EMX® into Cadence Virtuoso® similar to the technique described in Section 3.4.2.

Switch Implementation

The switches were implemented using wide nMOS transistors underneath the inductor traces, as shown in Fig. 3.5.1. For *Switch 1*, the switch appears exactly at the small signal ground of the inductor, and the capacitance does not affect performance. Therefore, a $\frac{W}{L} = \frac{1\text{mm}}{60\text{nm}}$ was used for this switch. The capacitance of *Switch 2* and *Switch 3* affect the tuning range more and more adversely. Therefore, these were made smaller: *Switch*

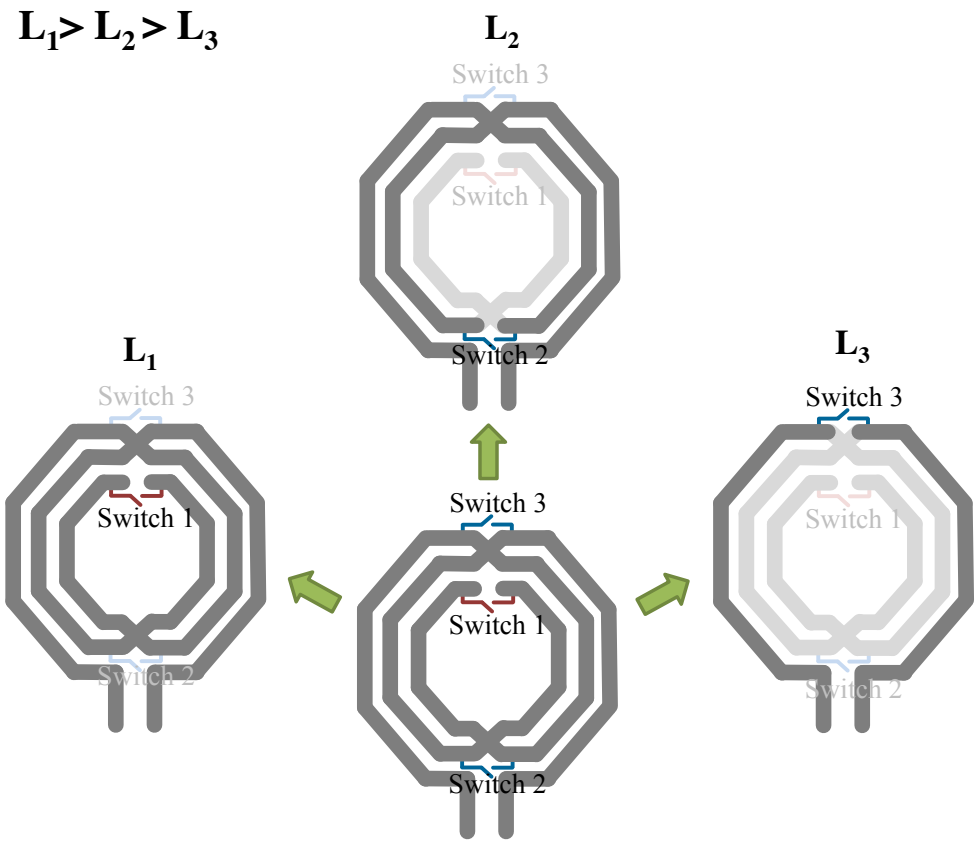

Figure 3.17: An inductor with 3 switches producing 3 inductances

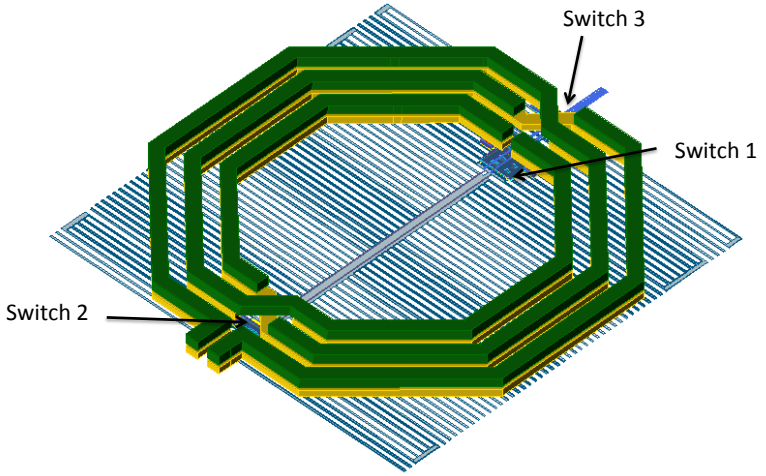

(a) Inductor implementation with three switches

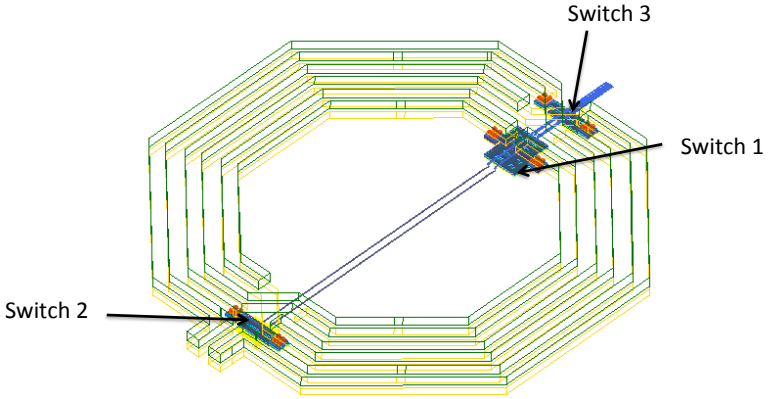

(b) Inductor implementation showing switches more clearly (transparent top-metals)

Figure 3.18: Inductor with three switches implementation: (a) shows the inductor implementation with the three switches, and (b) shows the switch locations more distinctly; top metals are made transparent and M1 ground shield is not shown in (b) for clarity

2: $\frac{W}{L} = \frac{500\mu\text{m}}{60\text{nm}}$, and *Switch 3:* $\frac{W}{L} = \frac{350\mu\text{m}}{60\text{nm}}$.

A summary of results for the different inductances along with that of the unswitched inductor (L_{orig}) is shown in Table 3.1.

Table 3.1: Summary of switched inductances

Case	f_0 (GHz)	L (nH) @ f_0	Q @ f_0	Peak Q	Peak Q freq.
L_{orig}	3.4	1.89	13.3	17.5	7.88
L_1	1.25	1.65	7.5	11.9	1.88
L_2	3.00	0.92	7.6	12.4	8.00
L_3	5.00	0.36	6.9	9.0	4.37

3.5.2 Capacitor Array Implementation

For obtaining a wide-tuning range per inductance value, a large switched capacitor array is necessary. The smallest frequency step can then be covered (with margin) using varactors for continuous tuning. In order to maintain a low VCO gain, K_{VCO}, the smallest discrete frequency step needs to be minimized necessitating a small unit capacitor. For a binary weighted capacitor array designed using multiples of unit capacitors, the foundry rules (design rule check rules, or DRC rules) impose minimum distance between unit capacitors causing the size of the array to grow extremely large. This array needs to be connected using interconnect, and the interconnect parasitics begin to dominate the characteristics of the tank. As a result, for large capacitor array design, efficient

techniques for design and layout are critical.

A top level layout screenshot of the second switched inductor VCO prototype is shown in Fig. 3.19. For obtaining a large tuning range, an 8-bit capacitor array, with a minimum switchable capacitance of 190fF, was designed. Therefore, $C_{max} = 48pF$ and the $\frac{C_{max}}{C_{min}}$ ratio of the capacitor array is 255. As shown in the figure, the size of the capacitor array is larger than the size of the inductor. Consequently, the interconnect parasitics begin to dominate the tank characteristics in a naive design.

The Problem of Interconnect Inductance

The parasitic inductance of interconnects in a large capacitor array can be significantly detrimental to the VCO performance. Some of these effects are discussed below.

1. *Altered frequency of oscillation:*

   As a result of parasitic interconnects, the overall inductance and capacitance of the tank is altered. As seen from the example in Fig. 3.19, the interconnects traverse a length almost equal to the length of the inductor coil. This results in an oscillation frequency different from the performance in schematic, or even extracted simulations[6] . This effect is shown diagrammatically in Fig. 3.20. The VCO on the right shows the alteration of the oscillation frequency due to the effect of interconnect inductance.

2. *Parasitic oscillation modes:*

[6] Most extractors do not model parasitic inductance. Even when they do, modeling accuracy is typically poor.

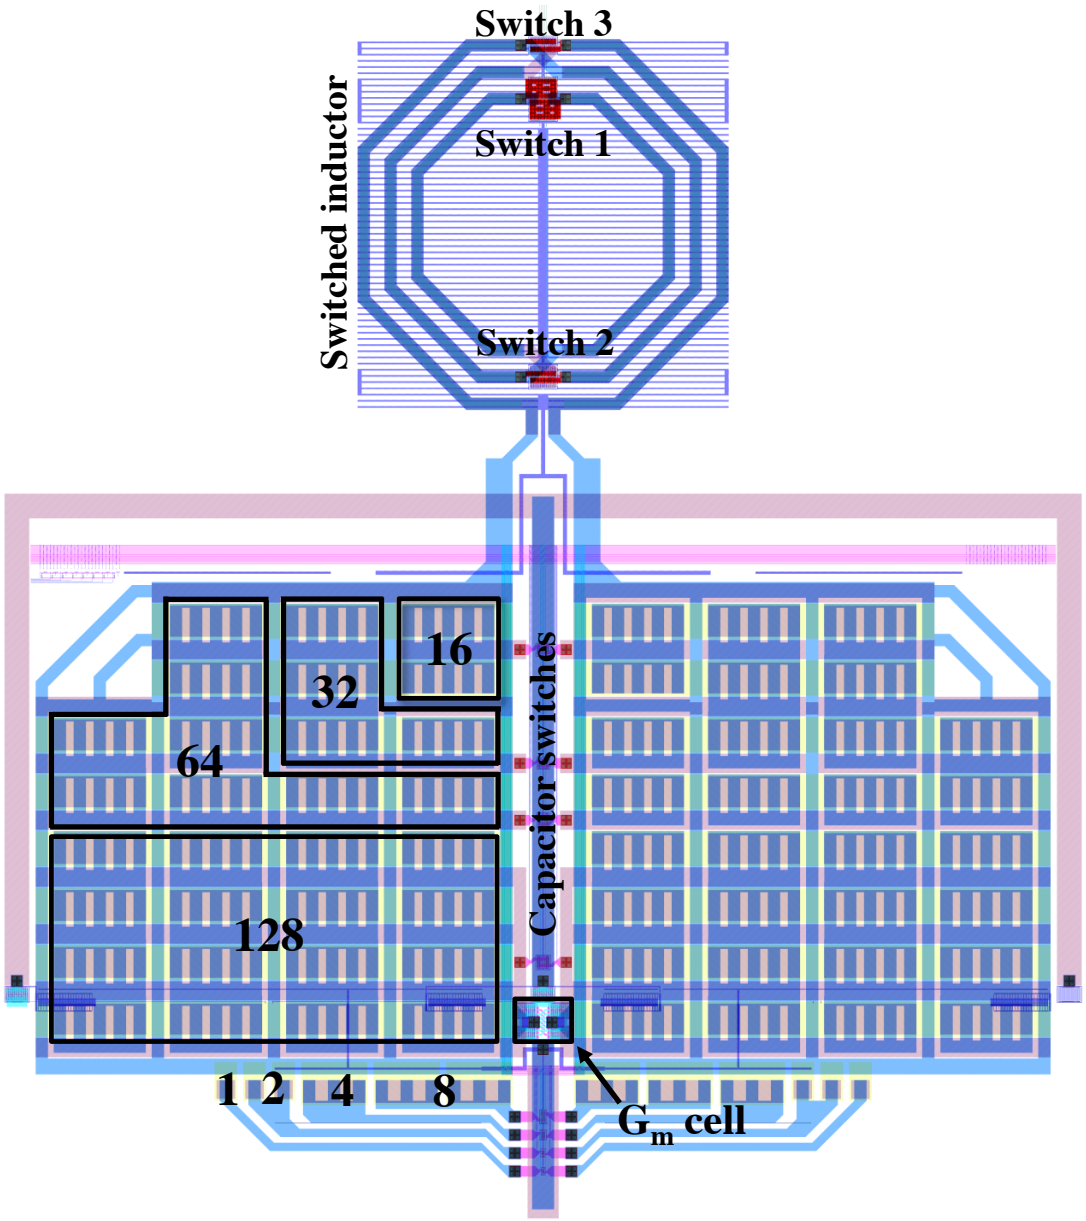

Figure 3.19: A layout snapshot of the VCO prototype II showing the switched capacitor array design

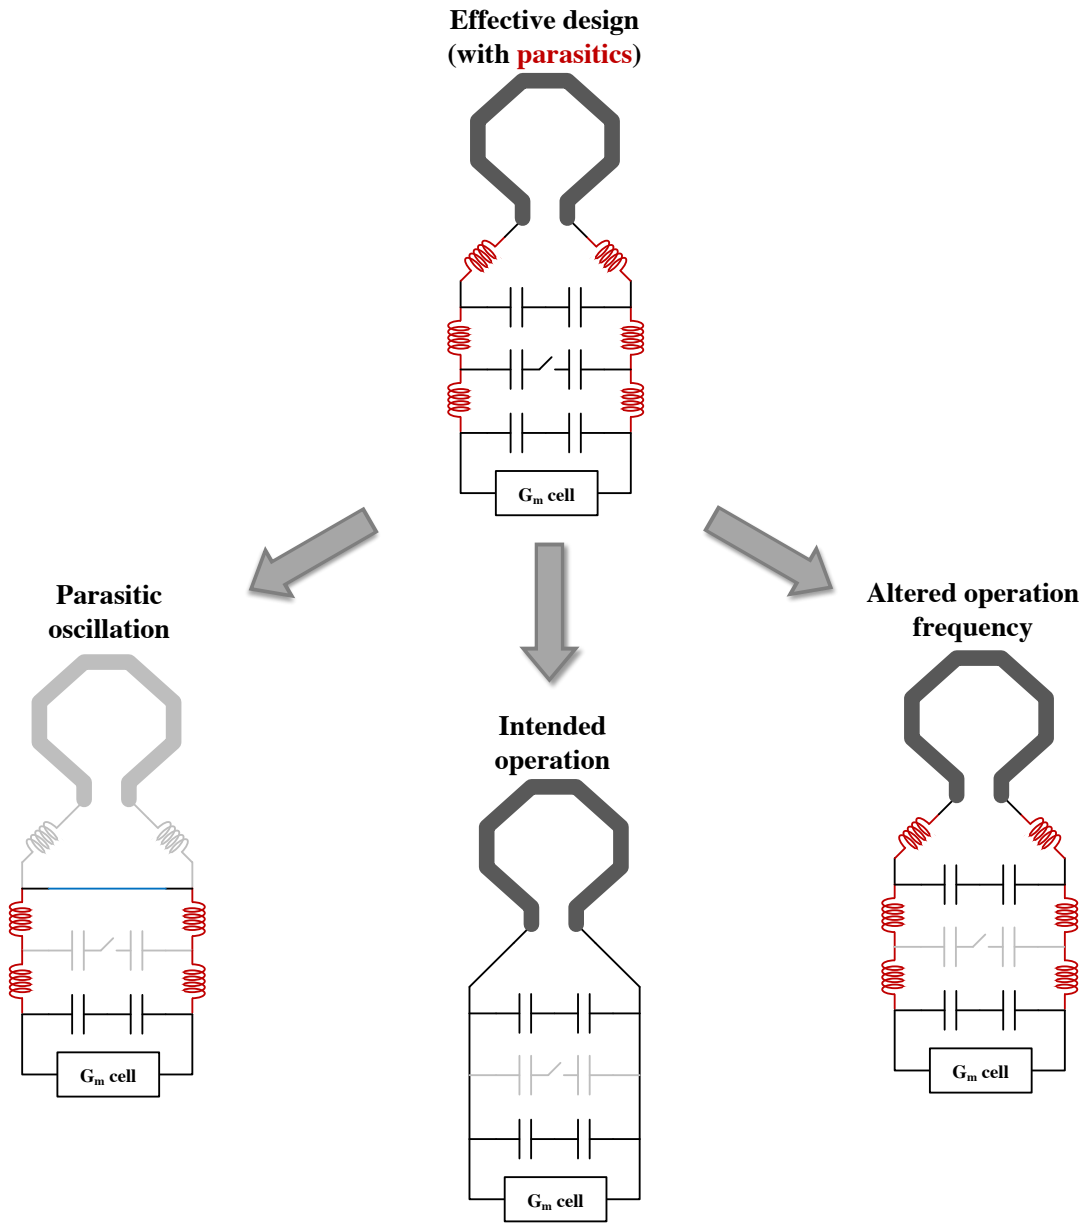

Figure 3.20: The adverse effects of interconnect inductance on VCO performance; parasitic inductors are shown in red

The parasitic inductance, being of the order of magnitude of the tank inductor, gives rise to a higher order tank with multiple modes of oscillation. As an example, consider the VCO shown in Fig. 3.20. As seen, the implemented VCO shown at the top comprises a higher order tank as a result of the parasitic interconnect inductances. The diagram on the left shows an unintended mode of oscillation, where, at the frequency of oscillation, one of the capacitors appears as a short. Consequently, at that frequency, the parasitic interconnect inductance, alongside another capacitor, forms a parasitic tank that oscillates at a completely different oscillation than intended.

3. *Phase noise degradation:*

   The parasitic interconnect inductance typically has a lower quality factor compared to the intended coil inductors in a given technology. As a result, in both the cases discussed above, the tank Q is degraded by the presence of the interconnects. This, in turn, severely affects the phase noise performance of the VCO.

4. *Difficulty in modeling:*

   The interconnect inductance is specifically difficult to model. It is typically necessary to solve Maxwell's equations in order to compute the inductance. Some extraction tools make an estimate of the partial self-inductance of the interconnects without solving Maxwell's equations; however, these estimates ignore the effect of the surrounding circuit elements, as well as the mutual inductance effects, and

can be a very inaccurate for modeling purposes. Full-wave electromagnetic simulators may be used for modeling; however, these simulators can handle only simple structures, and the simulation of the entire capacitor array is typically impossible.

5. *Frequency tuning difficulty:*

   Due to the additional interconnects, each time a capacitor in the switched capacitor array is turned on, a parasitic inductor gets added to the tank. This inductor affects the frequency of oscillation, and therefore, the desired frequency jump is not obtained. This can cause difficulty in designing a well-controlled and predictable frequency tuning characteristic.

A number of techniques can be used to minimize the adverse effects of the interconnect inductance, as well as model it accurately. These techniques, used for this prototype design, are discussed below.

1. *Techniques to reduce interconnect inductance:*

   (a) *Capacitor sizing:* Instead of selecting one unit capacitor for the 8-bit array, two unit capacitors were selected to minimize the total area occupied. In effect, the total capacitance was broken into two sets of 4 bits each. Each set uses its own unit capacitor: the smaller array uses a unit capacitance of 190fF, and the larger array uses a unit capacitance of 16×190fF. Using this technique, a small unit capacitor is utilized to lower the VCO gain, and a multiplicity of unit capacitors are used for good matching in each capacitor

set. However, by breaking the array into two parts, the area consumption is greatly reduced. These two sets, one comprising capacitor units $1-8$ and another, $16-128$, are shown in Fig. 3.19.

(b) *Capacitor placement:* In order to reduce the interconnects, it is necessary to optimize the placement of capacitors. First, the capacitors need to be placed such that their perimeter to area ratio is minimized. A circular arrangement is preferred. However, given the fabrication limitations, an octagonal/square placement is suitable. Fig. 3.19 shows such a low perimeter to area placement. Second, the capacitors need to be packed as close as allowed by the foundry rules (DRC rules). Third, the largest capacitors need to be placed closest to the G_m cell. This can be understood by observing the parasitic tank being formed in the left of Fig. 3.20. As seen in the figure, one of the capacitors is effectively shorted out (blue connector) at some (high) frequency to produce an effective inductance at that frequency. If the inductance at that frequency can oscillate with the remaining capacitance, a parasitic mode oscillation will occur. However, if the interconnect inductance is low, or the series capacitance is low, the frequency at which this occurs is very high, and the start-up criteria may not be fulfilled at these high frequencies. Moving the larger capacitors toward the G_m cells reduces the interconnect inductance associated with these capacitors, while the smaller capacitors that are further away have larger interconnect inductances associated with them, but get

shorted out at a much higher frequency. For the current prototype, the G_m cell was placed closest to the largest capacitor as shown in Fig. 3.19.

(c) *Interconnect meshing:* One way to reduce the interconnect inductance (and resistance) is to put them in parallel. In fact, the capacitor array can be interspersed by an interconnect mesh to reduce the overall inductance. Fig. 3.19 and Fig. 4.29 shows the implementation of such an interconnect mesh.

(d) *Interconnect interleaving:* The capacitors are connected using interconnects on the top and bottom plates. When ac current flows to the top plate, the bottom plate is being simultaneously discharged. As a result, the interconnects connecting the top and the bottom capacitor plates always carry currents in opposing directions. This can be exploited in the capacitor array to reduce the effect of interconnect inductance. For this, the top and bottom plate interconnects can be interleaved with each other such that any two adjoining interconnects carry opposing ac currents. This will cause a reduction in the inductance due to the coupling between these interconnects. In effect, the resultant mutual inductance due to the opposing current flow will significantly reduce the self-inductance of the interconnects. This effect is shown in Fig. 3.21

2. *Techniques to reduce the impact of interconnect inductance:*

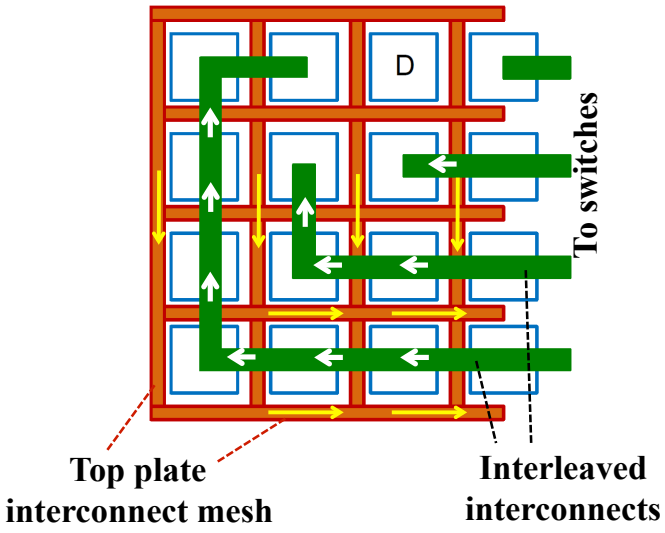

Figure 3.21: Interleaved interconnects to reduce the effect of interconnect inductance

(a) *G_m cell placement:* As discussed above, the placement of the G_m cell is critical to reduce the effect of interconnect inductance, as well as to avoid parasitic mode oscillations. As shown in Fig. 3.22, the G_m cell can be placed in multiple ways with respect to the tank. For option 1, shown in Fig. 3.22(a), the interconnect inductance will add to the tank inductor during normal operation, lowering the frequency of oscillation significantly. Additionally, parasitic modes of oscillation, as shown in the left of Fig. 3.20 can be triggered. For option 2, shown in Fig. 3.22(b), the interconnect inductance does not affect the tank inductance during normal operation (unless the capacitance is altered by the inductance as discussed in Section B.0.3). However, the possibility of parasitic oscillation modes remains. Option 3, shown in

Fig. 3.22(c), reduces the possibility of parasitic oscillation modes by reducing the distance between the G_m cell and any given capacitor (thus reducing the maximum interconnect inductance). However, a part of the interconnect inductance now adds to the tank inductor during normal operation. Option 3 was chosen for implementation in the Prototype II and the G_m cell was placed between the two capacitor arrays as shown in Fig. 3.19.

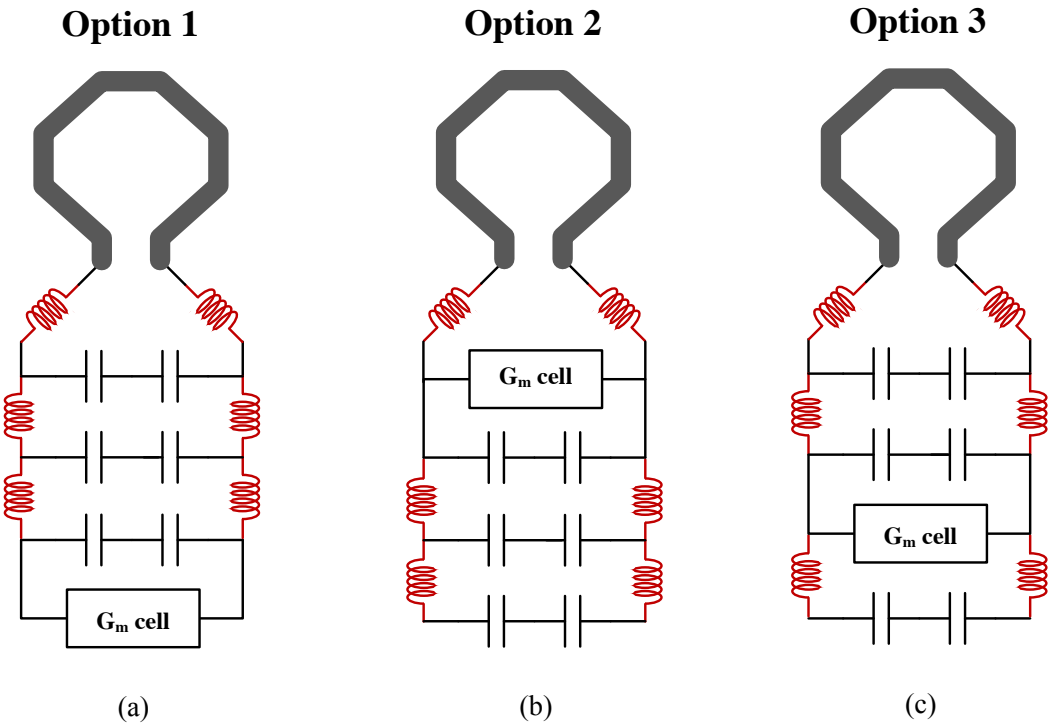

Figure 3.22: Different options to place the G_m cell with respect to the tank; interconnect parasitic inductance is shown in red

3. *Modeling and simulation of interconnect inductance:*

As evident from the discussion above, modeling the capacitor array interconnects

is critical to ensure desired operation of the VCO. Moreover, modeling the entire circuit in an electromagnetic simulator is impossible due to the required computations.

For modeling the interconnects, the schematic and layout views are partitioned into 2 parts: one that requires electromagnetic modeling, and another that does not. A layout view with only the interconnects is exported to an electromagnetic simulator. Results from the electromagnetic simulations (S-parameter data) is then imported into the circuit simulation environment (Cadence®). The resistive and capacitive parasitics of the rest of the circuit (without the interconnects) is extracted in the circuit simulator using extraction tools. For these extracted circuits, the interconnect inductance is minimal and does not affect the circuit functionality and performance. A few iterations might be required to arrive at the desired design. A flowchart of the design technique is shown in Fig. 3.23.

Switch Sizing

The MIM capacitors in the array are connected between the two differential VCO outputs 'Out+' and 'Out-' by NMOS transistor switches as diagrammatically shown in Fig. 3.24.

The switches contribute parasitic capacitance when **off** reducing the tuning range of the oscillator. Also, they contribute parasitic resistance when **on** degrading the phase noise of the oscillator.

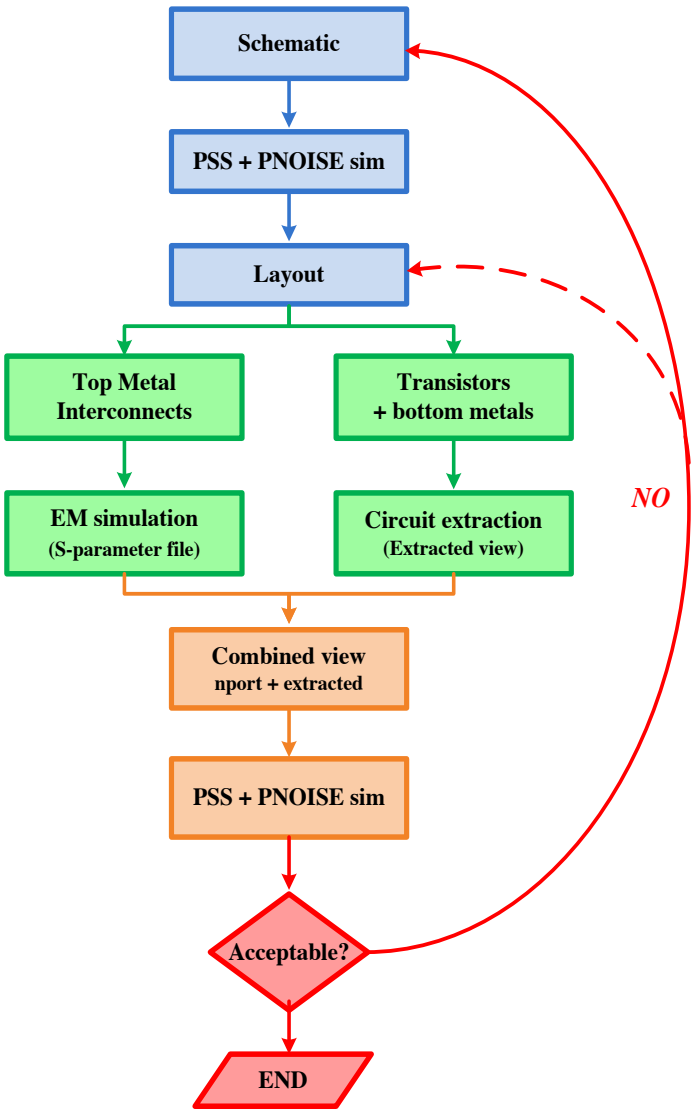

Figure 3.23: Electromagnetic simulation assisted VCO design flow

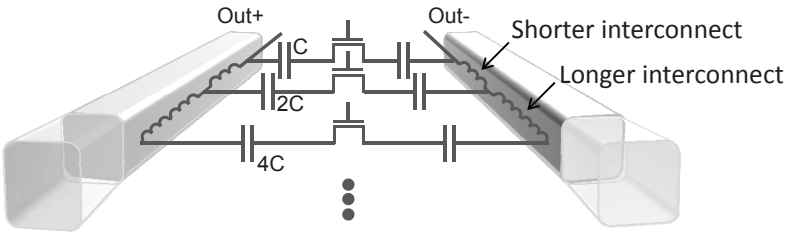

Figure 3.24: Diagram showing the connection of capacitors in a capacitor array

A naive way of sizing these NMOS switches would be as follows: the switch for the smallest capacitor bit is sized so as to minimally degrade the phase noise at the highest frequency of operation (corresponding to the worst phase noise as shown in Fig. 3.1). The other switches are sized in proportion to the increasing capacitance to ensure a constant RC product (constant Q at a particular frequency) and a smooth tuning curve. This sizing scheme for the lower frequency band is shown in column 4 of Table 3.2.

However, for large capacitor banks, these proportionally sized switches contribute considerable parasitic capacitance. Since the Q of a capacitor is given by $Q = 1/\omega RC$, the capacitor Q improves at low frequencies. Therefore, to maintain a constant capacitor Q, the size of the capacitor switches could afford to reduce for the larger capacitors which come into effect only at these lower frequencies. Such a sizing scheme is shown in column 5 of Table 3.2.

Note also, that since the quality factor of the inductor $Q = \omega L/R_s$ degrades at lower frequencies, the Q of the capacitor bank can actually be allowed to drop at the lower

Table 3.2: Capacitor array switch sizing (lower frequency bank)

Cap. units	Max. freq.	Ind. Q	Switch size in normalized units		
			Constant RC	Constant cap. Q	**Constant $L(\Delta\omega)$**
1	3.8	15.5	1	1	**1**
2	3.7	15.4	2	2	**2**
4	3.6	15.3	4	4	**4**
8	3.3	14.8	8	7	**7**
16	3	14.2	16	15	**11**
32	2.5	13	32	27	**16**
64	2	11.3	64	51	**11**
128	1.4	8.6	128	88	**15**
$\Sigma = 255$			$\Sigma = 255$	$\Sigma = 195$	$\mathbf{\Sigma = 67}$

frequencies without affecting the overall resonator quality factor as given by $1/Q = 1/Q_{ind} + 1/Q_{cap}$. Again, in this particular design, since the three largest capacitor bits are turned on only in the low frequency band, these switch sizes can be designed to be even smaller. The sizes can be fine-tuned through simulation so that the phase noise hits its worst case value every time a new capacitor is switched in. Such a scheme is shown in the last column of Table 3.2. The resultant parasitic capacitance contributed by the switch array is therefore reduced by 73.8%, i.e., by a $(255 - 67)/255$ ratio (Table 3.2). In effect, the Q of the capacitor bank is traded off with the parasitic capacitance of the switches, thereby improving the tuning range of the oscillator.

Other novel techniques for optimizing the capacitor array for improving the tuning

range of VCOs are discussed in Appendix B.

The measurement results from the design are presented below.

3.5.3 Measurement Results

Tuning range, phase noise, and power consumption measurements are presented. An unprecedented tuning range of 158% in measurement is obtained for LC VCOs. An excellent phase noise and power dissipation is maintained throughout the tuning range.

Tuning Range

The frequency tuning range (FTR) obtained from simulation spans 5.41GHz from 720MHz to 6.13GHz (158%) as shown in Fig. 3.25. The three inductances produce three frequency bands as shown. These bands are designed to overlap to ensure continuous frequency coverage in the face of process variations.

Phase Noise

The variation in phase noise measured at a 1MHZ offset over the tuning range is shown in Fig. 3.26. Due to the capacitor switch size optimization for the low frequencies, the phase noise is almost constant at these frequencies. The low frequencies were targeted for switch size optimization because these frequencies utilize the largest capacitors, and therefore the largest switches, contributing the maximum parasitic capacitance to the tank. Optimizing these switches for constant phase noise as discussed in Section 3.5.2 is instrumental in the achievement of the large tuning range discussed above.

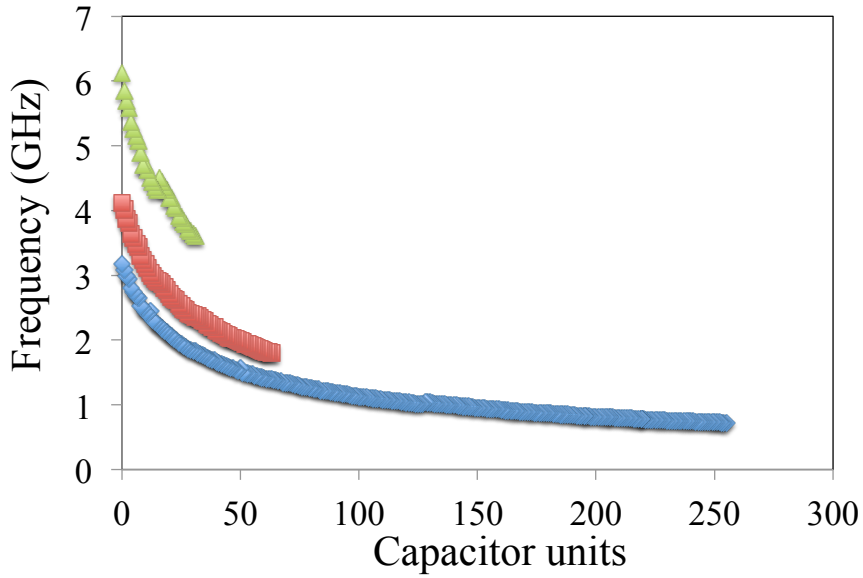

Figure 3.25: Frequency tuning range versus capacitance in terms of the 'x', where 'x' is the capacitance of a unit switched capacitor

The optimum phase noise at a 1MHz offset over the entire tuning range is seen to vary between -128.4 and -116.0 dBc/Hz. The degradation in phase noise as compared to the design described in Section 3.4 is due to the lower Q factor available in the IBM 65nm CMOS technology as compared to the IBM 130nm SiGe BiCMOS technology.

Power

The variation in power across the entire frequency range is shown in Fig. 3.27. The overall trends in power dissipation are as expected from the analysis in Section 3.2.

The figure of merit of the VCO is plotted across the entire tuning range in Fig. 3.28, and is seen to vary between 169 and 189 dBc/Hz with an average FOM of 176dBc/Hz

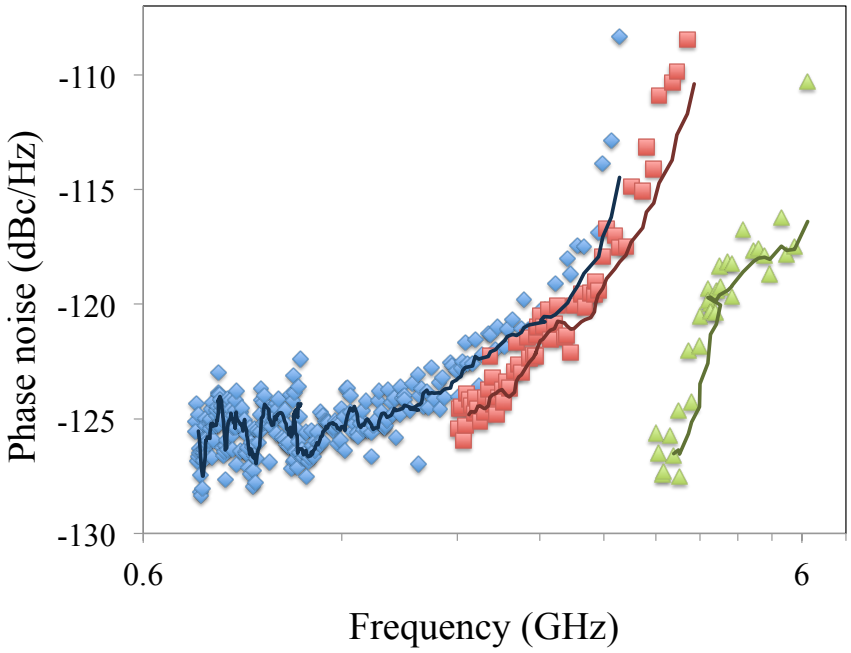

Figure 3.26: Variation of phase noise at 1MHz offset with frequency (log scale)

across the frequency points[7] . This represents an excellent performance suitable for SDR type applications.

A comparison of this work with previous wide-tuning range single LC tank designs is shown in Table 3.3. As seen from the comparison, this technique can provide much larger tuning ranges than are currently available in the state of the art designs. Also, the FOM_{PFTN} of these switched inductor oscillators is a significant improvement over other state of the art designs. Fig. 3.29 shows a number of previously published measured VCOs in IEEE conferences and journals between 1999 and 2010. The two measured

[7] Note that this average is extremely pessimistic since the density of points is much larger at the lower frequencies that represent a lower FOM as compared to the density of points at the higher frequencies represented by the higher FOMs in this case.

implementations discussed in this chapter are also plotted. As seen, a much larger tuning range, compared to other state of the art designs, is feasible using switched inductors.

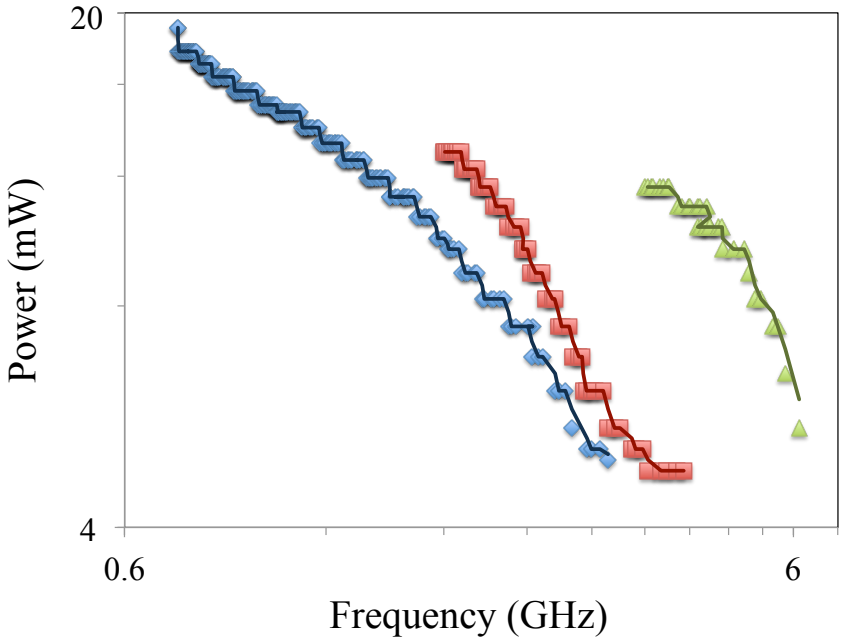

Figure 3.27: Variation of VCO core power dissipation with frequency (log log plot)

3.6 Conclusions

In this chapter, inductor switching was introduced as a viable solution for obtaining very wide tuning range oscillators with low phase noise. The advantages of inductor switching were identified over traditional pure capacitive tuning solutions. Significant advantages were seen to accrue from inductor switching, and a design methodology for switched-inductor oscillators was subsequently developed. Simple models were introduced to obtain further design insight into switched-inductor resonators. For a proof of concept,

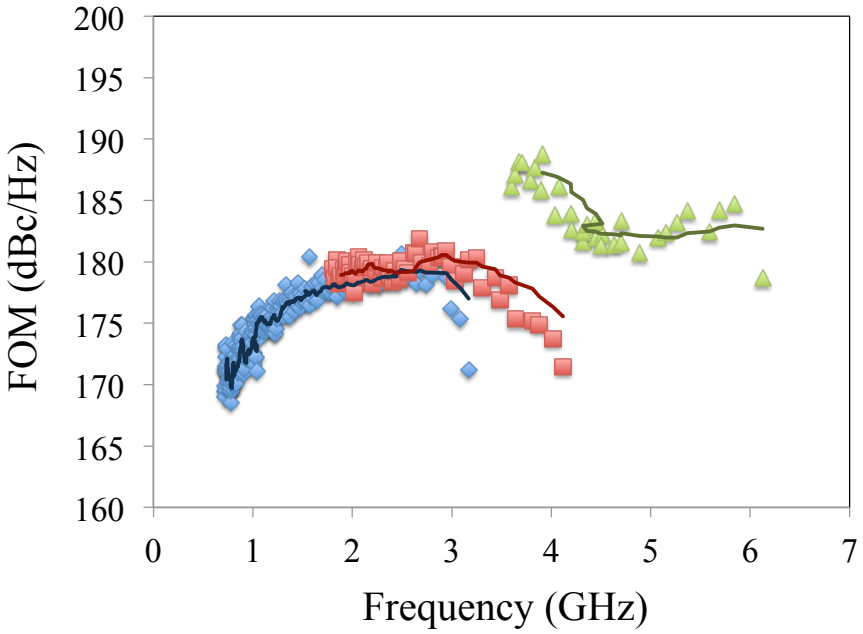

Figure 3.28: Variation of VCO FOM with frequency (log log plot)

a new inductor switching scheme was proposed. A first prototype design based on a single switch inductor was implemented in CMOS, and was seen to simultaneously achieve a phase noise between −117.2 and −122 dBc/Hz at 1MHz frequency offset, and an unprecedented tuning range of 87.2% (3.3−8.5 GHz) for single inductor LC VCOs in measurement. A second prototype was designed based on a three switch inductor for obtaining an increased tuning range. This prototype was seen to achieve an unprecedented 157% (0.7−6.1 GHz) tuning range for LC VCOs. The phase noise varies between −116.0 and −128.4 dBc/Hz at 1MHz frequency offset.

Table 3.3: VCO Performance Comparison

Ref	f_{min} tof_{max} (GHz)	FTR (%)	Phase Noise* (dBc/Hz)@1MHz	FOM_{PFTN} (dB) [82]	Implementation (μm)
[82]	2.0 to 2.6	26	−125.4 to −119.4	−3.1	0.35 BiCMOS
[88]	3.1 to 5.6	58.7	−120.8 to −114.6	5.9 to 10.3	0.13 SOI
[40]	1.1 to 2.5	73	−126.5 (f_0)	5.0 to 8.5	0.18 CMOS
[89]	3.6 to 8.4	74	−104 to −101.5	−4.6 to 4.0	0.13 CMOS
Prototype 1 (measured)	**3.3 to 8.4**	**87.2**	**−122 to −117**	**6.6 to 10.2**	**0.13 CMOS**
Prototype 2 (measured)	**0.72 to 6.13**	**157**	**−128.4 to −116.0**	**4.3 to 14.3**	**0.065 CMOS**

(*Assuming 20dB/decade drop with offset frequency)

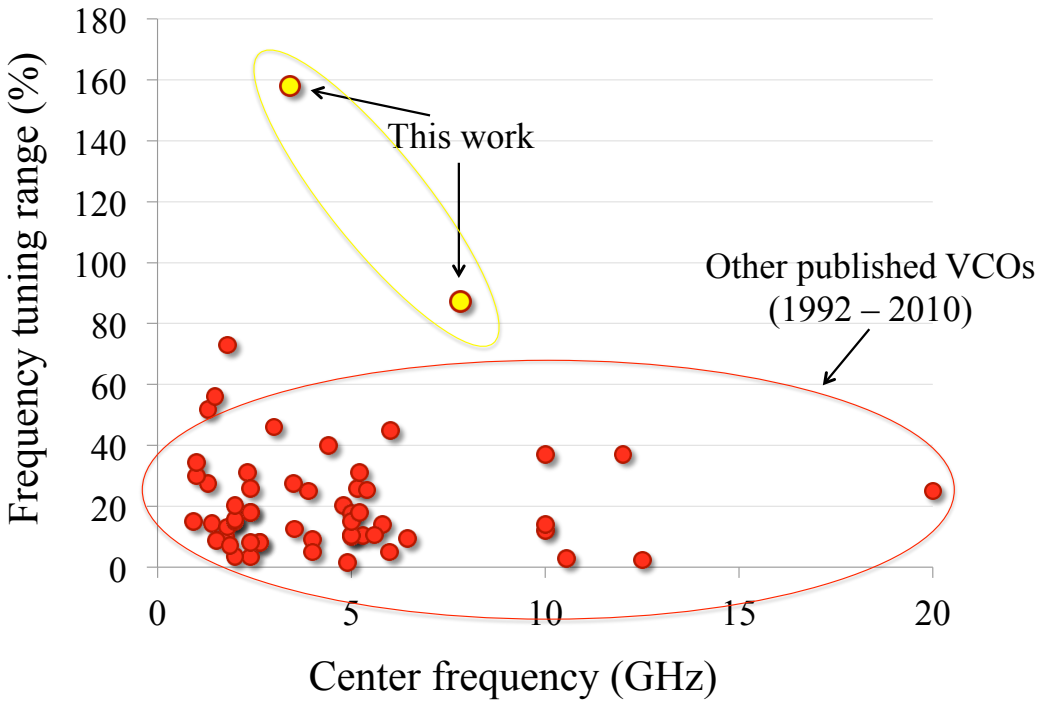

Figure 3.29: A comparison of the VCO implementations in this chapter with other measured implementations between 1999 and 2010

Chapter 4

Low Phase Noise VCO Topologies

4.1 Introduction

This chapter presents techniques to improve the phase noise of VCOs as compared to prior known topologies. These techniques are compatible with the switched inductor techniques described in Chapter 3.

Phase noise in CMOS LC VCOs is fundamentally limited by the oscillation amplitude, often determined by the breakdown voltage of the active devices, and the inherent device noise. In this chapter, we outline several new topologies based on capacitive feedback to improve the phase noise in LC VCOs. Using these topologies, we propose a way to operate the oscillator such that the amplitude of the oscillation is not limited by the breakdown voltages of the active device. The LC-tank is decoupled from the active devices and the coupling ratio is set so that the oscillation amplitude can reach

the maximum value determined by the breakdown of the passive components only.

The key idea is to expose the active devices in the oscillator to a small fraction of the full amplitude. This can be implemented with a capacitive divider, transformer or a similar passive decoupling device. An additional important benefit of the proposed approach is the reduced amount of noise that is injected from the active devices into the tank. The overall phase noise is then lowered by both the increased amplitude and relatively lower noise injection. The large amplitude is reduced for the following (active) stages by using capacitive dividers; note that the phase noise is not significantly affected by this amplitude division.

In this chapter, we discuss a number of VCO topologies [90], a few of which have been explored in prior literature, and others that are novel. We will discuss and classify the topologies by the degree to which they achieve the following objectives, all of which are advantageous to achieving low phase noise.

1. Increased oscillation amplitude in the oscillator LC tank
2. Reduced injection of active device noise into the LC tank
3. Improved loaded tank Q (or reduced degradation of tank Q due to active device and bias circuit loading)
4. Independent biasing of device drain and gate (or collector and base) for controlling oscillation amplitude and the region of device operation
5. Reduced waveform distortion

The benefits of 1, 3, and 4 above follow directly from Leesons formula [91]; the benefits of reduced noise and waveform distortion (2, and 5 above) are discussed in [92].

4.2 A Family of Linearized Transconductance (LiT) VCOs

4.2.1 Cross-coupled Topology

The first topology discussed is the widely used cross-coupled VCO topology shown in Fig. 4.1. For analysis, the circuit can be represented by an equivalent half circuit, shown in Fig. 4.2. From this, we can see that the oscillation amplitude will be determined by the G_m non-linearity; that is, the oscillation amplitude will increase until G_m decreases to the point that it is just high enough to overcome losses in the tank, represented by R_p in Fig. 4.2. This is shown graphically in Fig. 4.3, where G_m decreases with increasing swing until $G_{m,eq} = 1/R_p$. Also, the G_m varies significantly over the cycle as the MOSFET device moves from the saturated region into the triode region, distorting the waveform and periodically reducing the tank Q as the device impedances change with G_m. It can also be seen from Fig. 4.1 that all of the drain current noise in the devices flows through the tank. A constant current source for biasing is typically used as in Fig. 4.1, and all of its current noise flows into the tank as well.

The G_m non-linearity is critical in determining the phase noise in the oscillator. The small signal transconductance (g_m) is shown in Fig. 4.3 corresponding to G_m for an infinitesimally small amplitude of oscillation ($V_{g,sw} = 0$). This g_m determines both the

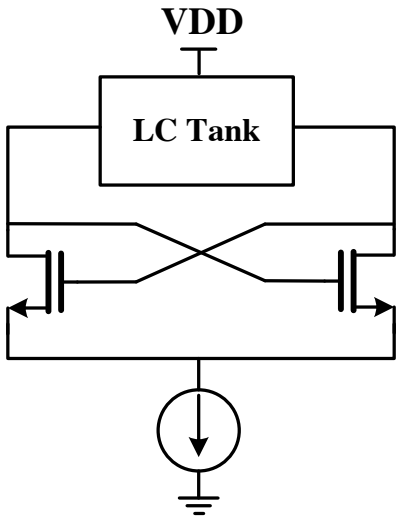

Figure 4.1: Cross-coupled VCO: Full circuit

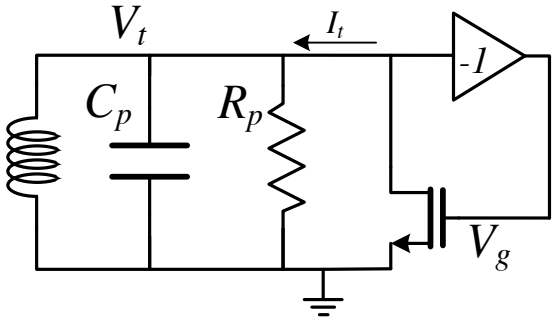

Figure 4.2: Cross-coupled VCO: Half circuit

start-up margin and the active device noise in the oscillator. This can be understood as follows: for a given amplitude of oscillation, the larger the G_m non-linearity, the greater the ratio of the small signal transconductance to the equilibrium large signal transconductance ($g_m/G_{m,eq} = g_m \cdot R_p$) which is the start-up margin provided in the oscillator. However, since g_m also determines the active device noise during zero crossings (when the oscillator is most vulnerable to phase noise [92], $g_m/G_{m,eq} = g_m \cdot R_p$, also determines the noise to signal ratio in the oscillator. Therefore, a lower start-up margin provides a better noise performance in the oscillator. Linearizing the transconductance curve so as to reduce $g_m \cdot R_p$ effectively reduces the start-up margin, and therefore the noise to signal ratio in the oscillator. Consequently, for a particular amplitude of oscillation, the active device noise is minimized.

The desired transconductance linearization may be achieved using several feedback techniques. These techniques can be broken up into two fundamental principles:

1. *Linearization using a lower swing at the node causing non-linearity:* This can be understood using the example of a cross-coupled oscillator half-circuit shown in Fig. 4.4. In MOSFETs, the node primarily responsible for non-linearity is the drain. The non-linearity is caused by the MOSFET entering the triode region as shown in Fig. 4.4(a) for a cross-coupled oscillator. However, a modified feedback technique can be used to reduce the swing on the drain, and reduce triode operation, as shown in Fig. 4.4(b). This reduction in triode operation would linearize the G_m curve as shown in Fig. 4.5, thereby increasing the amplitude of operation[1]

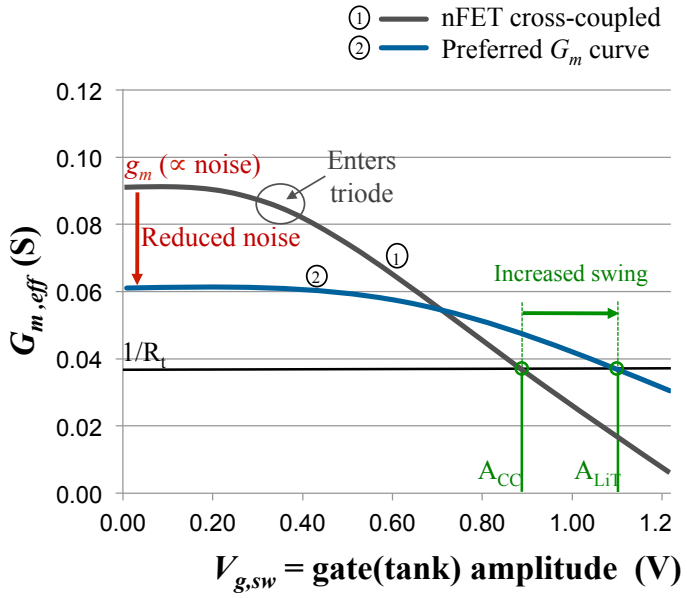

Figure 4.3: Transconductance vs. voltage swing showing a typical oscillator G_m curve, and a preferred curve for improved phase noise

.

2. *Effective linearization by using a lower swing on the active devices through feedback:* This can be understood using another example half-circuit model shown in Fig. 4.6. In this case, the transistor swing is reduced due to a reduction in the overall transconductance, and a lowering of the G_m curve as shown in Fig. 4.7. However, as shown in Fig. 4.7, the tank amplitude is increased by the same ratio as the lowering of the G_m curve. This causes an overall increase in the oscillation amplitude as shown in Fig. 4.7.

[1] Note that this also increases the small signal transconductance (g_m) increasing the active device noise. Later we describe a technique to achieve both larger amplitude as well as reduced noise as discussed in Section 4.2.1.

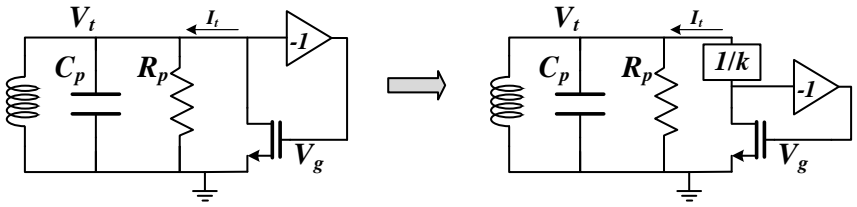

(a) Cross-coupled (b) With feedback

Figure 4.4: Feedback concept for G_m linearlization using lower voltage swing on the drain node

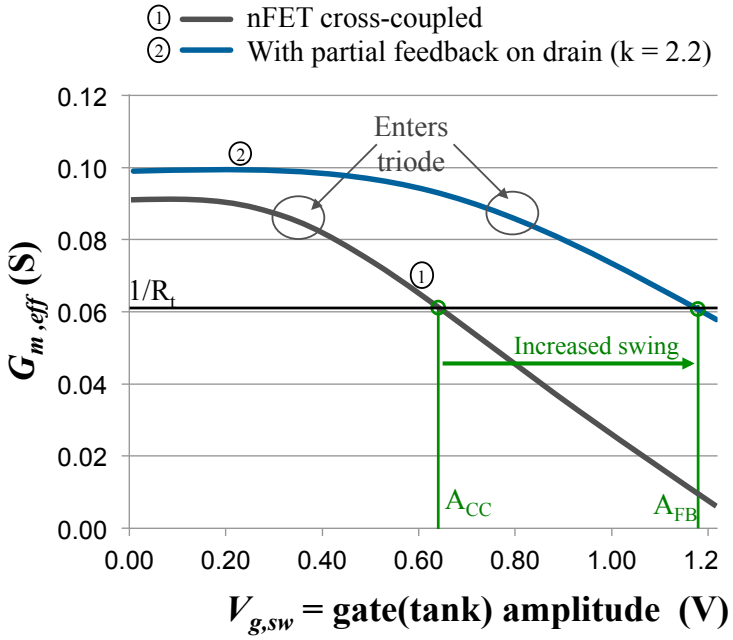

Figure 4.5: Transconductance vs. voltage swing showing a typical oscillator G_m curve, and one that uses lower drain voltage swing for linearization

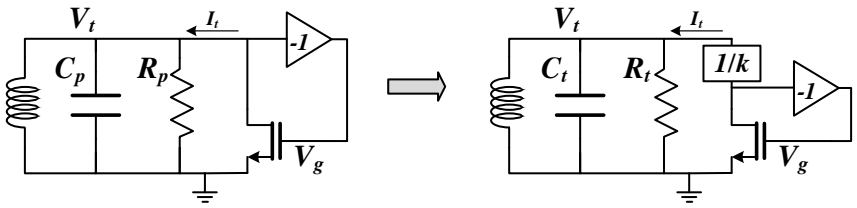

(a) Cross-coupled (b) With feedback

Figure 4.6: Feedback concept for G_m linearlization using lower voltage swing for the active device

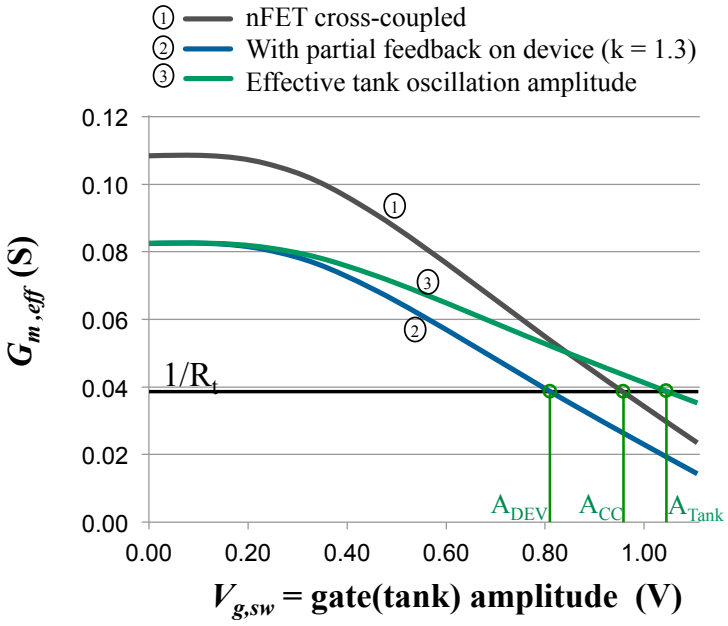

Figure 4.7: Transconductance vs. voltage swing showing a typical oscillator G_m curve, and one that uses lower active device swing for linearization

Next we consider a family of topologies that utilize decoupled tanks in order to achieve large oscillation amplitudes, and low phase noise. This is achieved by some form of transconductance linearization (LiT). We call this family Linearized Transconductance VCOs (LiTVCOs). The following LiTVCOs are considered in this chapter:

1. Drain-divider
2. Gate-divider/Drain-divider
3. Gate-divider Type B
4. Gate & Drain divider Type A - v1
5. Gate & Drain divider Type A - v2
6. Gate & Drain divider Type B - v1
7. Gate & Drain divider Type B - v2

4.2.2 Drain Divider

The first LiTVCO topology we examined is referred to here as the drain-divider topology [93,94], and shown in Fig. 4.8. In this topology, the LC tank is moved to a location between the gates of the two active devices, and additional biasing inductors are included between the drains of the active devices and V_{DD}. An equivalent half-circuit is shown in Fig. 4.9, and the simulated change in G_m with gate voltage over the cycle is shown in Fig. 4.10. The circuit builds up a higher oscillation amplitude across the tank

as in Fig. 4.10, but G_m varies much less across the cycle, and it is easier to arrange operating conditions to prevent the devices from entering triode region, where low drain impedance will lower the tank Q severely. We classify this circuit as a drain divider because the tank oscillation amplitude is divided onto the drains by a factor of $1/k$, where $k \approx \frac{C_c+Cd}{C_c}$. An additional benefit to this circuit is that the drain current noise does not all flow through the tank; some of it flows through the series combination of C_c and R_p, and some through C_d. Effectively, due to the linearization of the G_m curve, the noise from the active devices is reduced while maintaining the same oscillation amplitude.

Note that the output of the oscillator can be tapped at any node with almost identical close-in phase noise performance. In case the active devices (buffer transistors) require a low voltage swing input, the tank amplitude can be attenuated using capacitive dividers (without affecting the close-in phase noise adversely), or the outputs can be tapped from a low swing node (e.g. drain node). The output nodes for the other topologies described in this chapter can be selected based on a similar principle.

4.2.3 Gate-Divider or Base-Divider

The third topology is referred to here as the base-divider (or gate-divider) VCO [95], and shown in the simplified schematic of Fig. 4.11. In this topology, the LC tank remains unchanged from Fig. 4.1, but a capacitive divider reduces the signal swing to the bases of the cross-coupled transistors. The devices are shown here as BJTs rather than MOSFETs because this topology is primarily advantageous for bipolars. The highly

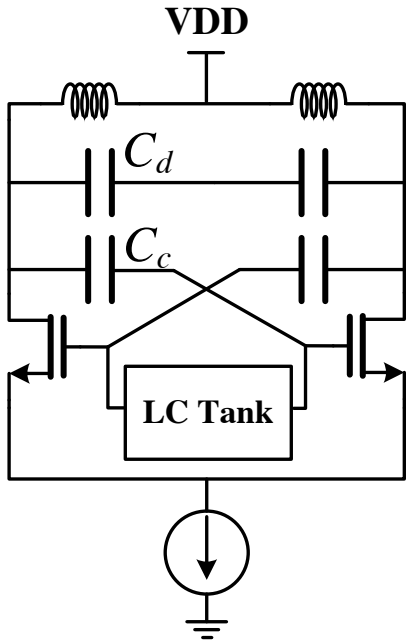

Figure 4.8: Drain divider VCO: Full circuit

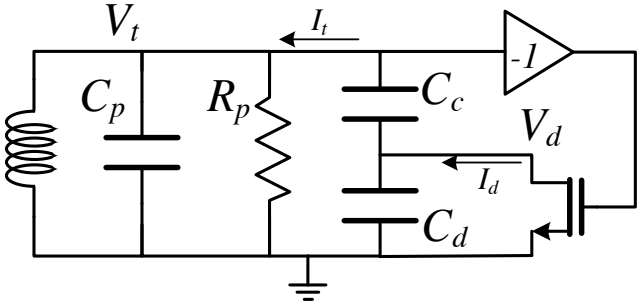

Figure 4.9: Drain divider VCO: Half circuit

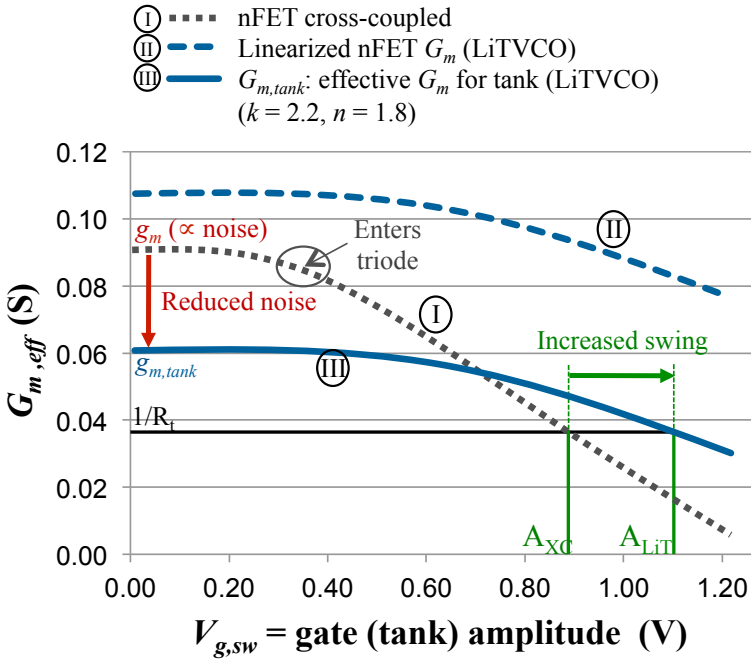

Figure 4.10: Transconductance linearization in LiT VCO

non-linear variation in the impedance seen looking into the base of a BJT reduces the Q of the tank and distorts the waveform; the capacitive base-divider permits a higher signal swing across the LC tank before this effect occurs. However, the maximum tank signal swing is still limited by the collector-base breakdown voltage of the BJTs. Also note that this topology is actually disadvantageous for MOSFETs because the triode region of operation is increased, resulting in higher G_m non-linearity, lower tank amplitude, increased degradation of the tank Q, and a more distorted waveform.

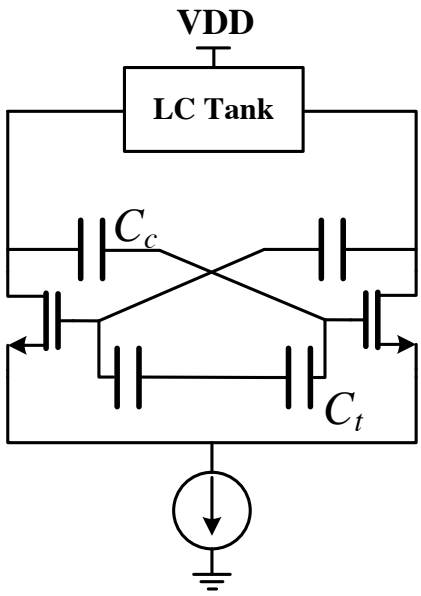

Figure 4.11: Gate divider / base divider VCO: Full circuit

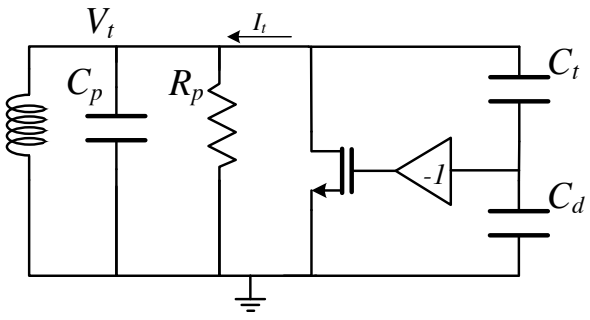

Figure 4.12: Gate divider / base divider VCO: Half circuit

4.2.4 Gate-Divider Type B

A new topology which we refer to as Gate-divider Type B resolves many of the problems with the base-divider topology, and is shown in Fig. 4.13. In this topology the LC tank is moved to a position between the gates of the devices, isolated by capacitors C_t. The equivalent half-circuit is shown in Fig. 4.14, where it can be seen that the tank oscillation amplitude is divided onto the gates of the active devices by ratio

$$\frac{1}{k} \approx \frac{C_t}{C_d + C_t}$$

The tank is decoupled from the active devices, and the capacitive divider allows higher oscillation amplitude in the tank without chances of breakdown at the device gates. In the case of MOSFETs, this circuit is more advantageous than Fig. 4.11, because it does not increase the region of triode operation (and corresponding G_m non-linearity) compared to Fig. 4.1, unlike the base-divider topology of Fig. 4.11. Like the drain divider, an additional benefit to this circuit is that the drain current noise does not all flow through the tank; some of it flows through the series combination of C_t and R_p, and some through C_d. If we define $\frac{1}{n}$ to be the fraction of the drain current noise that flows into the tank, then

$$n \approx \frac{Z_{C_d} + Z_{C_t} + R_p}{Z_{C_d}} = \frac{C_t + C_d + R_p \cdot sC_d \cdot C_t}{C_t}$$

The capacitive divider in Fig. 4.13 reduces active device loading on the tank, therefore tank Q is degraded less than in the case of the cross-coupled oscillator (Fig. 4.1),

and there is less waveform distortion. However, although the Gate-divider Type B is an improvement over the base-divider of Fig. 4.11, further improvements are possible.

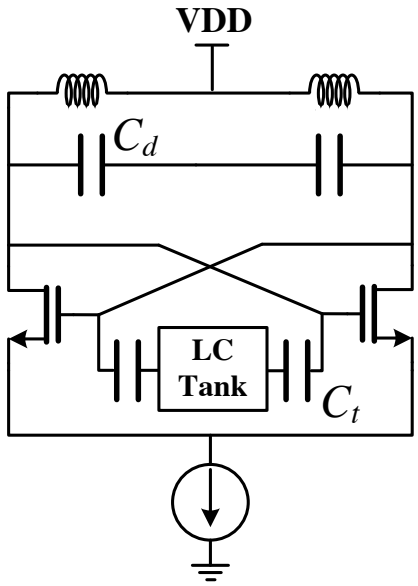

Figure 4.13: Gate divider / base divider VCO Type B: Full circuit

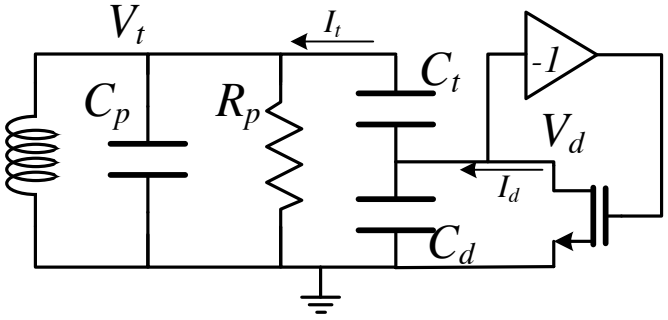

Figure 4.14: Gate divider / base divider VCO Type B: Half circuit

4.2.5 Gate & Drain Divider Type A

An advantageous new topology which we refer to as the Gate & Drain divider Type A - v1 is shown in the simplified schematic of Fig. 4.15. It removes the constraints of the topologies of both Fig. 4.8 and Fig. 4.11 by more completely decoupling the LC tank from the active device terminals, providing a way to increase oscillation amplitude above the device breakdown limits at both gate (or base) and drain (or collector), while also reducing the effects of device non-linearity at both gate and drain. The equivalent half-circuit is shown in Fig. 4.16. Once again, $\frac{1}{n}$ represents the fraction of drain current noise that flows into the tank, and $\frac{1}{k}$ represents the voltage division ratio from the tank to the device gates, where

$$n \approx \frac{Z_{C_c} + Z_{C_t} + R_p + Z_{C_d}}{Z_{C_d}} = \frac{C_t C_d + C_d C_c + s R_p C_d C_c C_t + C_t C_c}{C_t C_c}$$

and

$$k \approx \frac{C_t + C_k}{C_t}$$

where

$$C_k = \frac{C_c C_d}{C_c + C_d}$$

This circuit provides an increase in the oscillation amplitude using two different techniques. The capacitor feedback ensures that even though the inherent device transconductance non-linearity limits the oscillation amplitude on the gate/base, the amplitude on the tank is larger than this limitation by a factor of k. Additionally, the device non-linearity is reduced as has been shown in Fig. 4.7 which increases the amplitude

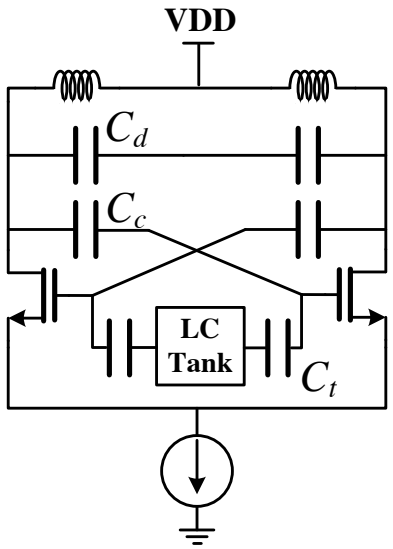

Figure 4.15: Gate & drain divider Type A - v1 VCO: Full circuit

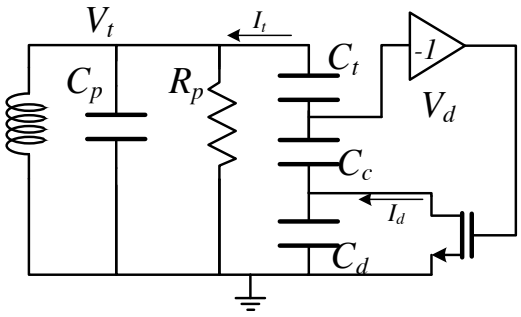

Figure 4.16: Gate & drain divider Type A - v1 VCO: Half circuit

of oscillation even on the gate/base for a given amount of injected active device noise. Moreover, by completely decoupling the tank from the active devices, the oscillation amplitude is not limited by the breakdown of these devices. To summarize, the advantages of the Gate & Drain divider Type A - v1 topology are enumerated below:

1. Amplitude of oscillation is increased by tank to gate (base) feedback-ratio
2. The amplitude is also increased G_m linearization from a lower drain voltage swing
3. The tank is completely decoupled from the active devices, and can therefore sustain much larger voltage swings
4. Only a fraction of the active device noise is injected into the tank
5. Only a fraction of the noise from the biasing circuitry flows into the tank
6. The tank Q is not degraded by the MOSFETs entering the triode region of operation, causing lower distortion, and improving the phase noise

Advantages 2 and 6 are less dominant in full BJT versions of this topology. The relative advantages between MOSFET and BJT versions of this design is discussed further below.

Fig. 4.17 shows a variant of the Gate & Drain divider Type A - v1 topology, referred to as Gate & Drain divider Type A - v2. This topological rearrangement moves the LC tank to a location between the device drains, isolated by capacitors C_t. The corresponding half-circuit is shown in Fig. 4.18. By comparison with Fig. 4.16, the half-circuit of

Fig. 4.18 shows clearly that the voltage division ratio from the tank to the device gates is increased. That is, the fraction of the tank voltage on the device gates is lower (by $\frac{1}{k}$), where

$$k \approx \frac{C_l + C_g}{C_l}$$

where

$$C_l = \frac{C_c C_t}{C_c + C_t}$$

Conversely, the fraction of the tank voltage on the device drains is higher than in Fig. 4.16 , and is given as

$$\frac{C_t + C'_k}{C_t}$$

where

$$C'_k = \frac{C_c C_d}{C_c + C_g}$$

Thus, the Gate & Drain divider Type A - v1 topology provides more G_m linearization than the Gate & Drain divider Type A - v2 topology since the swing on the drain is relatively lower causing lower non-linearity. However, the feedback ratio to the gate is relatively larger for the Gate & Drain divider Type A - v2 topology as compared to the Gate & Drain divider Type A - v1 topology, therefore increasing the amplitude on the tank as a result of this feedback ratio. Therefore, these two topologies provide a way to trade off between the two techniques for amplitude increase depending on the device characteristics, and the requirements for a particular design. For example, in practical implementations, since the MOSFET non-linearity is highly dependent on

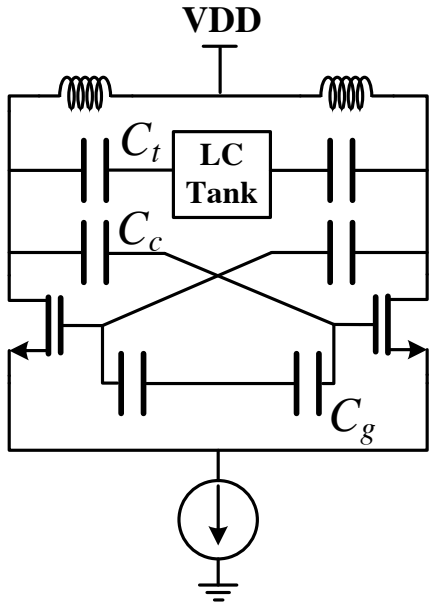

Figure 4.17: Gate and drain divider Type A - v2 VCO: Full circuit

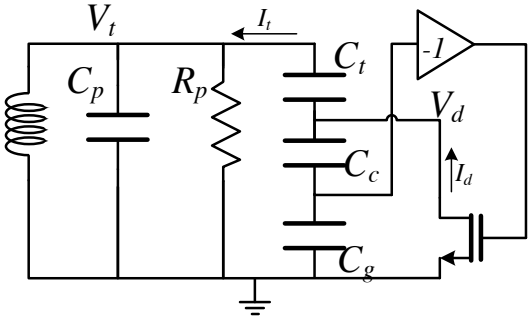

Figure 4.18: Gate and drain divider Type A - v2 VCO: Half circuit

the drain voltage swing, Gate & Drain divider Type A - v1, with its lower voltage swing, is suitable for MOSFET implementations. On the other hand, considering BJT versions of the same topology, the collector voltage swing does not affect the BJT non-linearity as significantly as the swing on the highly non-linear base voltage node. Therefore, the amplitude increase from the large tank-to-gate feedback (k) ratio in Gate & Drain divider Type A - v2 makes this topology particularly suitable for the BJT version of the design.

4.2.6 Gate & Drain Divider Type B

Another advantageous new topology referred to as Gate & Drain divider Type B - v1 is shown in the simplified schematic of Fig. 4.19. The equivalent half-circuit is shown in Fig. 4.20. The Gate & Drain divider Type B - v1 topology shares the same advantages as the Gate & Drain divider Type A - v1 topology, namely:

1. Amplitude of oscillation is increased by tank to base feedback-ratio
2. The tank is completely decoupled from the active devices, and can therefore sustain much larger voltage swings
3. Only a fraction of the active device noise is injected into the tank
4. Only a fraction of the noise from the biasing circuitry flows into the tank
5. For MOSFET versions, the amplitude is also increased G_m linearization from a lower drain voltage swing

6. Also, for MOSFET versions, the tank Q is not degraded by the MOSFETs entering the triode region of operation, causing lower distortion, and improving the phase noise

7. Additionally, this topology offers an extra degree of freedom for the choice of variables, and therefore provides more control over the performance specifications for this design.

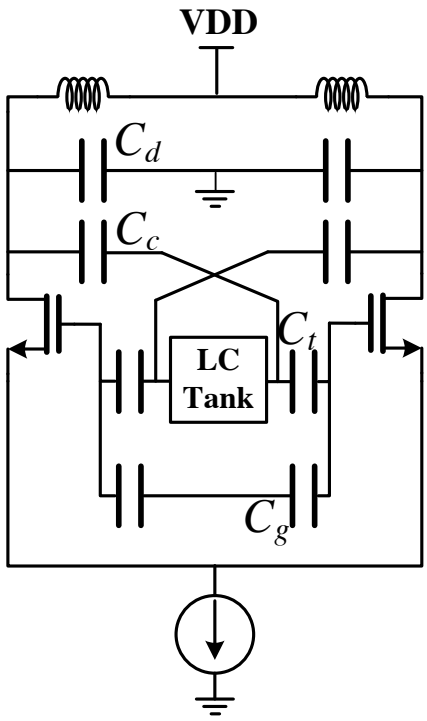

Figure 4.19: Gate & drain divider Type B - v1 VCO: Full circuit

As discussed above, compared to the Gate & Drain divider Type A topologies, the Gate & Drain divider Type B - v1 topology adds an additional capacitor (compare Fig. 4.20 with Fig. 4.16), which now gives freedom choosing total capacitive load,

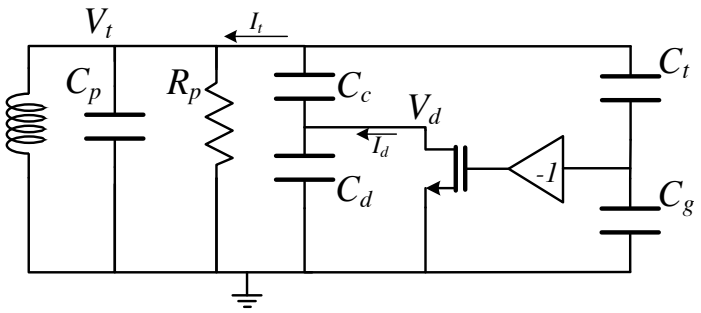

Figure 4.20: Gate & drain divider Type B - v1 VCO: Half circuit

bias tuning, and feedback ratios (that is, there are 4 capacitors and 4 specifications). This is particularly useful in bipolar implementations where smaller base voltage swing is desired, due to higher device transconductance and greater non-linearity in the base-emitter circuit.

A closely related topology is shown in Fig. 4.21, referred to as the Gate & Drain divider Type B - v2 VCO. Referring to the equivalent circuit in Fig. 4.22, it can be seen that this topology also has 4 capacitors and thus gives complete flexibility in choosing total capacitive load, bias tuning, and feedback ratios. In fact, the Gate & Drain divider Type B - v1 and the Gate & Drain divider Type B - v2 are equivalent in the sense that one circuit can be transformed into the other, with only a change in the value of the capacitors C_c, C_d, C_t, and C_g. Which circuit would be preferred would depend on which one gave more convenient and easily realizable capacitor values for a given application.

Table 4.1 shows a summary of topologies and corresponding advantages, in accordance with an embodiment of the present principles. That is, Table 4.1 summarizes

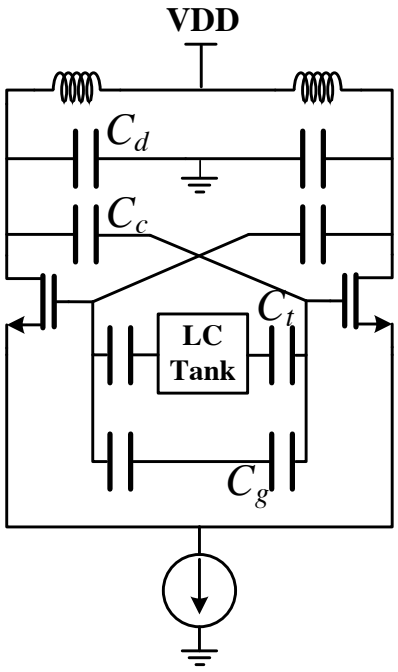

Figure 4.21: Gate & drain divider Type B - v2 VCO: Full circuit

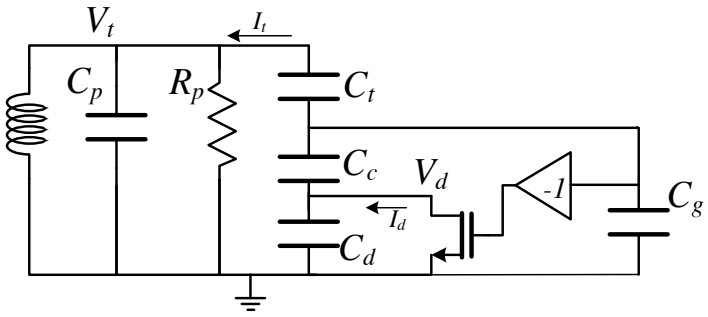

Figure 4.22: Gate & drain divider Type B - v2 VCO: Half circuit

the different voltage controlled oscillator (VCO) topologies that were discussed and compares their advantages (and disadvantages). The four aspects compared include the amplitude of oscillation (A), the noise injected into the tank (N), the loaded quality factor of the tank due to the impact of the transconductor (Q), and the waveform distortion due to varying resistance in the loaded tank (D). These aspects have a direct impact on the phase noise performance of the VCO. The mechanisms for improvement/alteration in these parameters for the different VCOs are mentioned in parentheses. Additional comments regarding the topologies have also been emphasized in the Comment column. As can be seen from Table 4.1, our proposed architectures improve these aspects (A, N, Q and D) of the VCO as compared to prior art cross-coupled topologies. As a consequence, the new topologies provide different mechanisms to achieve a low phase noise performance. Table 4.1 also provides an intuition of the trade-offs between these topologies. The choice of a particular topology can be determined by the specific requirements for that design.

4.3 High swing capacitor switch design

As discussed above, these topologies exploit high oscillation amplitudes in order to obtain better phase noise performance. For this, the active devices are decoupled from the tank so that the amplitude of oscillation is not limited by the active device breakdown limits. However, in order to realize large frequency tuning using these topologies, switched capacitors are necessary. The switches are realized using active devices, and

their possible breakdown, as well as turn-off, turn-on characteristics may limit the realizable voltage swing. Here, we discuss a technique for switched capacitor design capable of handling large voltage swings.

Since the total voltage swing is comparable/larger than VDD, there is a possibility for the switches in the switched capacitor array to turn on even when disabled. To solve this problem, a series combination of two switches (2X wider than if one was used) can be utilized as shown in Fig. 4.23. This allows the gate nodes (node B) to be small signal opens and swing partially as shown. Consequently, the gate source voltage (V_{gs}) swing is reduced allowing the switches to remain off when so desired through the entire oscillation cycle as shown in Fig. 4.23. Also, by allowing node B to be a small signal open, the effective parasitic capacitance resulting from the C_{db} of the two 2X switches is equivalent to that of one 1X switch. Therefore, there is minimal additional loading on the tank due to this technique. pFET switches typically have a thicker oxide (the on resistance is identical to an nFET's) in the lower technology nodes, and can be utilized to enable them to tolerate the large tank voltage swing.

4.4 Biasing

The LiTVCO topologies typically feature a tank that is decoupled from the active devices. Moreover, in a number of topologies, the individual nodes in the active device are decoupled from each other. This is unlike the traditional cross-coupled topology shown in Fig. 4.1. As a result, it is possible to bias the active devices independently.

Moreover, a fixed current based biasing need not be used. This voltage based biasing can provide robustness to process, supply voltage, and temperature variations.

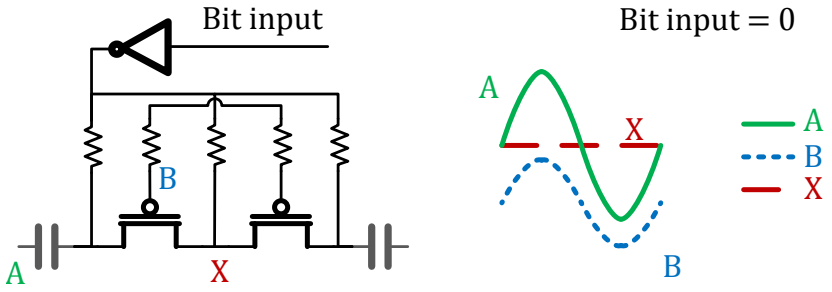

Figure 4.23: A proposed switch scheme for the high amplitude LiTVCO tank switched capacitors

Table 4.1: VCO summary

Topology	Advantages	Figures	Comments
Cross-coupled	A → Normal N → All active + all bias Q → Degraded Q D → Distorted	4.1, 4.2	This is the popular topology to which we are comparing
Drain-divider	A → Increased (feedback), no breakdown N → Fraction of active, fraction from bias Q → Degraded (less than cross-coupled) D → Lower distortion than cross-coupled	4.8, 4.9	The biasing network can be tuned close to the oscillation frequency to improve the phase noise
Gate-divider/ base-divider	A → Increased (feedback) & reduced (g_m non-lin.) N → All active, no bias Q → Degraded Q (more than cross-coupled)	4.11, 4.12	Not suitable for MOSFET designs because g_m non-linearity increases reducing the swing

Continued on next page

Table 4.1 – continued from previous page

Topology	Advantages	Figures	Comments
	D → Potentially more distorted than cross-coupled		
Gate-divider/ base-divider – Type B	A → Increased (feedback), no breakdown N → Fraction of active, fraction from bias Q → Degraded (less than cross-coupled) D → Lower distortion than cross-coupled	4.13, 4.14	More suitable for BJT designs compared to MOSFET designs
Gate-Drain-divider Type A - v1	A → Increased (feedback, g_m lin.), no breakdown N → Fraction of active, fraction from bias Q → Not degraded D → Minimal distortion	4.15, 4.16	Emphasis on g_m linearization
Gate-Drain-divider Type A - v2	A → Increased (feedback, g_m lin.), no breakdown N → Fraction of active, fraction from bias Q → Not degraded	4.17, 4.18	Emphasis on amplitude increase from feedback

Continued on next page

Table 4.1 – continued from previous page

Topology	Advantages	Figures	Comments
	D → Minimal distortion		
Gate-Drain-divider Type B - v1	A → Increased (feedback, g_m lin.), no breakdown N → Fraction of active, fraction from bias Q → Not degraded D → Minimal distortion	4.19, 4.20	Complete freedom in choosing all parameters
Gate-Drain-divider Type B - v2	A → Increased (feedback, g_m lin.), no breakdown N → Fraction of active, fraction from bias Q → Not degraded D → Minimal distortion	4.21, 4.22	Different cap values from Type B - v1 (design adv.); complete freedom in choosing all parameters

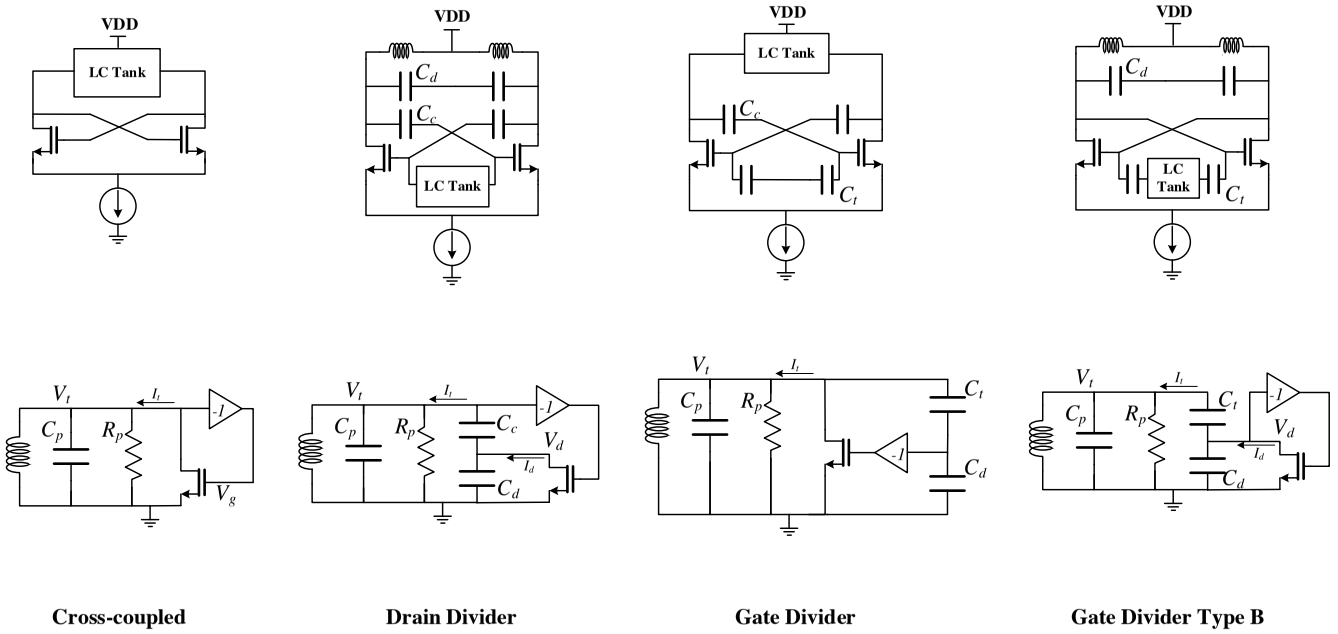

Figure 4.24: A (diagrammatic) summary of the different VCO topologies discussed (I)

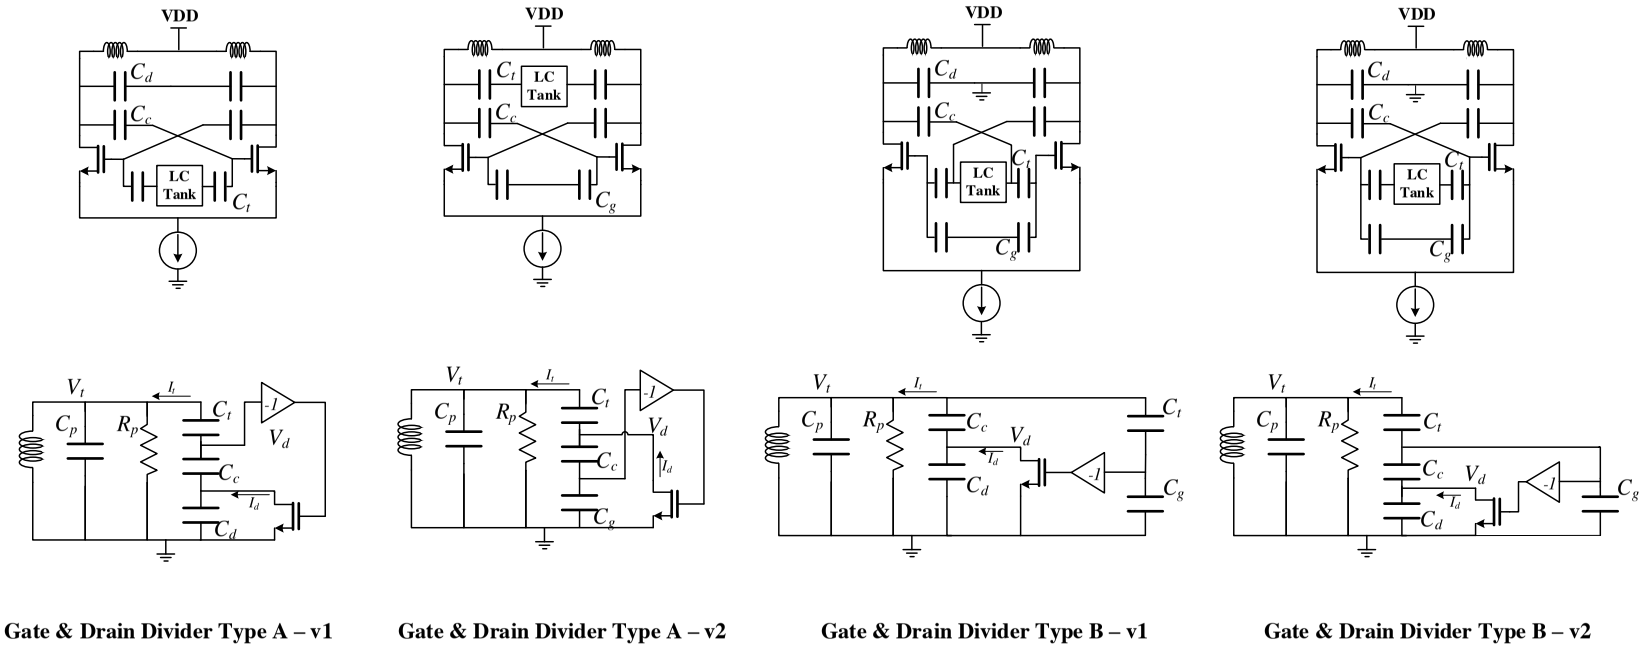

Figure 4.25: A (diagrammatic) summary of the different VCO topologies discussed (II)

4.5 Drain-Divider LiTVCO

In this section, we introduce the new approach based on *transconductance linearization* of the active devices that increases the signal swing while reducing the active device noise contribution in LC VCOs [94]. Consequently, this provides a significant improvement in the VCO phase noise performance. We implement this approach using a signal feedback technique similar to that presented in [93]. Additionally, for the capacitor array, we utilize a new layout approach based on interconnect inductance mitigation. Using these novel approaches, we are able to globally optimize the VCO for phase noise, power, tuning range and PVT robustness. Consequently, the implemented PLL incorporating the resulting linear transconductance VCO (LiT VCO) achieves both an excellent measured phase noise of -130dBc/Hz at 22GHz (10MHz offset), and, in contrast to the 6.7% tuning range of the VCO-only result reported in [93], a large frequency tuning range (FTR) of 23% (21.8$-$27.5GHz). This range covers, with margin, operation over 22.8$-$26.4GHz, supporting 60GHz superheterodyne radios using a frequency doubler [96].

4.5.1 Transconductance Linearization

The oscillation amplitude in an LC VCO is limited by the non-linearity of the large signal device transconductance (G_m)[2] . The G_m vs. gate ac swing ($V_{g,sw}$) for a FET based cross-coupled VCO is shown in Curve I in Fig. 4.10. As shown, the G_m drops with increasing oscillation amplitude ($= V_{g,sw}$). The equilibrium amplitude, A_{XC}, is

[2] Here, G_m is defined as $G_m = \frac{I_\omega}{V_\omega}$, where I_ω and V_ω are the ac drain current and ac gate voltage at the frequency of oscillation, ω.

reached when $G_m.R_p = 1$. The main source of non-linearity is the transistor entering the triode region.

In contrast, if the VCO is designed to avoid the triode region, G_m is linearized. To achieve this in the LiT VCO, the LC tank is placed between the gates of the FET devices as shown in Fig. 4.26. The half circuit model for the VCO is shown in Fig. 4.27. The capacitive divider ensures that the drain swing is a fraction of the swing across the LC tank ($V_{g,sw}$), eliminating FET triode-operation. From Fig. 4.27,

$$V_d = \frac{V_t}{k}, \quad \text{where } k \approx \frac{C_c + C_d}{C_c} \tag{4.1}$$

Moreover, the capacitive dc isolation allows the FET gate ($V_{g,bias}$) to be biased lower than its drain (V_{DD}), further eliminating triode operation. Consequently, FET non-linearity is reduced, resulting in a more linear G_m as shown in Curve II in Fig. 4.10 for $k = 2.2$.

Also note that due to the capacitive division, only a part of the FET ac current flows into the tank as shown in Fig. 4.10. From Fig. 4.27, where I_d is the FET drain current, and I_t is the fraction of I_d flowing into the tank, at the resonant frequency,

$$\begin{aligned} I_t = \frac{I_d}{n}, \quad \text{where} \quad n &= \frac{Z_{C_d} + Z_{C_c} + R_p}{Z_{C_d}} \\ &= \frac{C_c + C_d + R_p.sC_d.C_c}{C_c} \end{aligned} \tag{4.2}$$

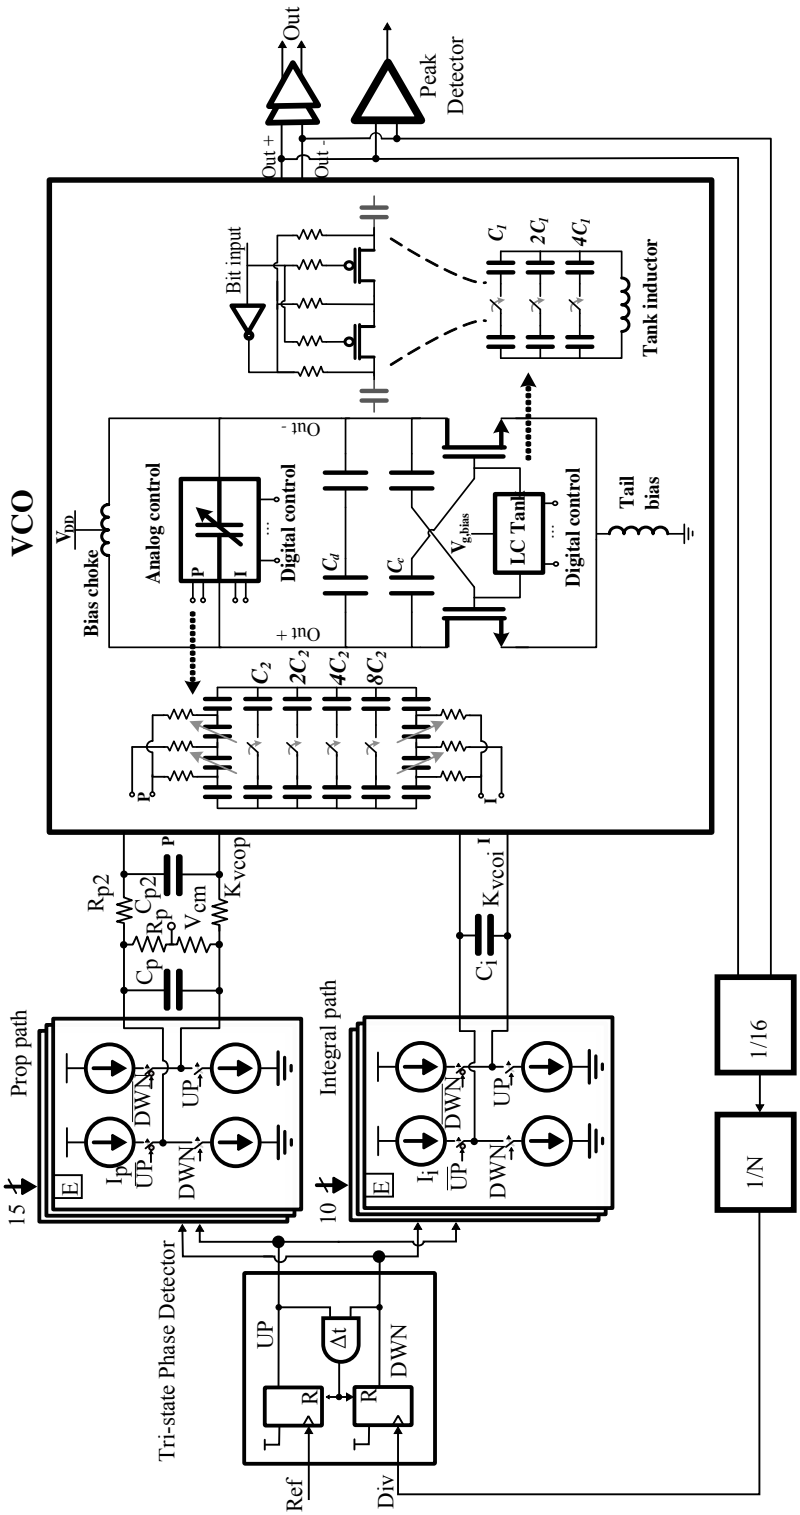

Figure 4.26: Block diagram of the PLL showing the detailed VCO schematic including frequency tuning schemes

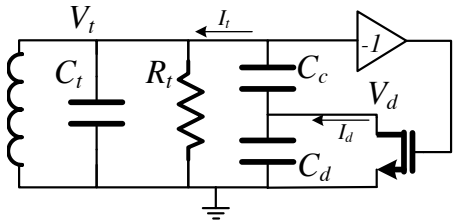

Figure 4.27: An equivalent half circuit model of the LiT VCO showing the capacitive feedback technique employed (biasing details not shown)

This reduces the effective G_m (to $G_{m,tank}$) as plotted in Curve III in Fig. 4.10 for $k = 2.2$ and $n = 1.8$. The resulting oscillation amplitude, A_{LiT}, using curve III for the LiT VCO, is larger than A_{XC}, and results in improved phase noise performance in the LiT VCO.

As shown in Fig. 4.28, the larger oscillation amplitude can be viewed as a voltage limit extension in the LiT VCO. This enables a power vs. phase noise trade-off beyond the voltage limit imposed in a cross-coupled VCO.

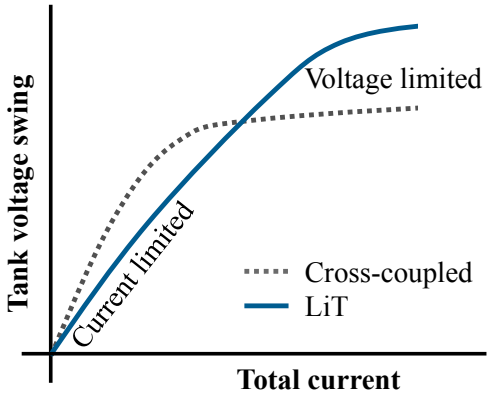

Figure 4.28: Voltage limited regime extension in LiT VCO

Additionally, phase noise contribution from the FET is reduced: active device noise

power at the oscillation zero crossings (when the oscillator's phase noise sensitivity is at its highest) is proportional to $g_m (= G_m(V_{g,sw} = 0))$. Note that for the LiT VCO, the effective transconductance across the tank ($g_{m,tank}$) is lower than the device g_m by n (i.e. $g_{m,tank} = g_m/n$). Therefore, at the zero crossings, the noise contribution of the active devices is reduced by the same amount as shown in Fig. 4.10. Additionally, by eliminating triode operation, the LiT VCO reduces degradation of the loaded tank Q. Consequently, the LiT VCO enables a larger tank amplitude and yet injects less noise into the tank as compared to the cross-coupled VCO leading to significant improvements in phase noise.

The foregoing theory of LiT VCO operation enables design optimization by utilizing additional power to enable a larger tank amplitude, lower noise injection and improved loaded tank Q as compared to traditional cross-coupled VCOs. Based on G_m curves as shown in Fig. 4.10, any cross-coupled VCO optimized using [97] can be converted into a LiT VCO design. The target amplitude of oscillation, A_{LiT}, can be maximized to the tolerable limits imposed by the technology. G_m curves can then be constructed for the active device for different combinations of k and n and a global optimum for power, phase noise, tuning range and robustness can be determined[3] .

[3] For example, increasing k linearizes the G_m at the expense of power. Similarly, increasing n lowers the active device noise at the expense of power and start-up margin. Also, k and n are related through equations 4.1 and 4.2.

4.5.2 Implementation

The LiT VCO PLL shown in Fig. 4.26 has been implemented in IBM's 32nm SOI CMOS process. The PLL architecture incorporating the LiT VCO includes two separate fully differential control paths: an integer path consisting of a charge pump and capacitor, and a proportional path consisting of charge pump and resistor (as well as a ripple capacitor). The divider consists of a static prescalar of division ratio of 8 or 16, followed by a fully synchronous divider which can divide from 2 to 128. The PFD consists of a classic tri-state phase detector. The charge pump currents, proportional path resistance and integral path capacitance value are all made programmable in order to support a wide range of bandwidths and reference frequencies. More details of the PLL can be found in [98, 99]. The LiT VCO design is optimized at $k = 2.2$ and $n = 1.8$ (Fig. 4.10). C_c, C_d and C_t are selected so as to obtain this k and n combination while adding minimum capacitive load on the tank.

Tuning range For the targeted 60GHz application, a large frequency tuning range (FTR) is a critical requirement [96]. We utilize switched capacitors on the drain and tank simultaneously (Fig. 4.26) to manipulate the feedback ratio and extend G_m linearization (and low phase noise performance) over a large FTR. At these high frequencies, a large FTR (in conjunction with a low K_{VCO}) also involves implementing a large switched capacitor array in the tank. For optimal phase noise, a small inductor is utilized [97]. Consequently, the capacitor array and the inductor occupy comparable areas (Fig. 4.29)

and the interconnect inductance significantly impacts VCO performance.

To avoid this problem in the LiT VCO, the interconnect parasitics are minimized. As shown in Fig. 4.29, the unit capacitors are arranged in a square, reducing the perimeter to area ratio of the structure. The top plates of the MIM-capacitors are connected using a mesh to reduce parasitics by parallelization. The bottom-plate interconnects are interleaved into the top plate mesh and carry an opposing current so that the resulting mutual inductance subtracts from the self-inductance of the interconnections. Consequently, compared to the 100pH tank inductor, the interconnects connecting the 70x80μm^2 capacitor array contribute only 12pH of inductance. By limiting the inductive parasitics, parasitic oscillation modes are eliminated, and a large tuning range of 23% is obtained.

Discrete frequency tuning is achieved using a 3-bit coarse and a 4-bit fine switched capacitor array (Fig. 4.26). Varactors are used for continuous tuning and provide around 100MHz analog tuning range achieving the desired low K_{VCO}.

Biasing As shown in Fig. 4.26, an inductor is used to resonate out, at $2f_0$, the parasitic capacitance at the tail node. The resulting source degeneration at $2f_0$ reduces the G_m cell noise contribution when one transistor is in cut-off. The tail inductor is implemented using a transmission line to retain the symmetry in the LiT VCO layout (see Fig. 4.29). A DAC digitally controls the VCO gate bias for performance optimization. This bias control also eliminates the tail current source which can be a significant source of noise.

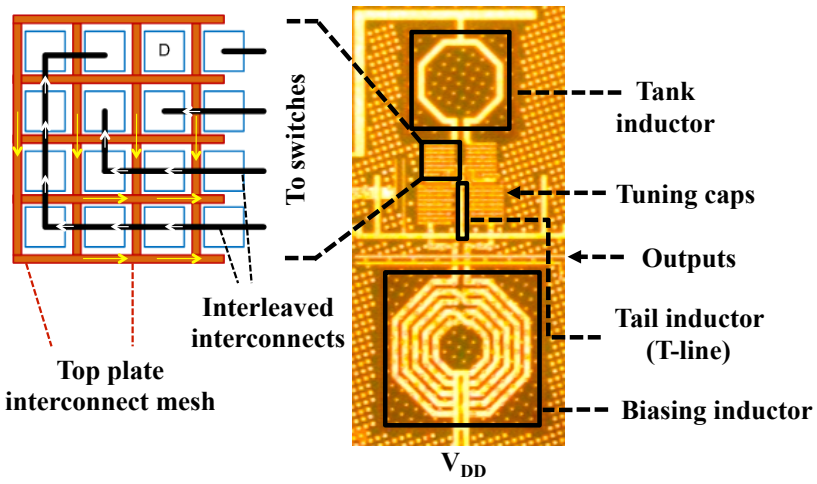

Tank inductor inductance	100pH
Tank inductor area	80 x 80 = 6400μm^2
Cap array interconnect inductance	12pH
Capacitor array area	70 x 80 = 5600μm^2
Biasing inductor	6nH (300μm dia.)
T-line tail inductor	70pH

Figure 4.29: Die photo of the LiT VCO showing the capacitor array design details for minimizing interconnect inductance, and other implementation details

Also, only a fraction of the noise from the biasing choke (6nH) (Fig. 4) flows into the tank reducing the total noise from the biasing circuitry in the LiT VCO as compared to the cross-coupled topology.

4.5.3 Measurement Results

The LiTVCO is implemented in the IBM 32nm SOI CMOS process. Measurement results from the PLL are plotted in Figs. 4.30−4.38. Fig. 4.30 shows the phase noise vs. frequency offset from a 22.6GHz carrier with −4.4dBm output power (single-ended). The differential output power is greater than 0dBm over the entire FTR.

Fig. 4.31 shows the measured phase noise @10MHz offset over the FTR for 3 dies at different radii on the wafer. The phase noise of the LiT VCO-based PLL varies from −130dBc/Hz at 22GHz to −126dBc/Hz at 27GHz. This performance is notably superior to other PLL's for 60GHz applications (Table 4.2).

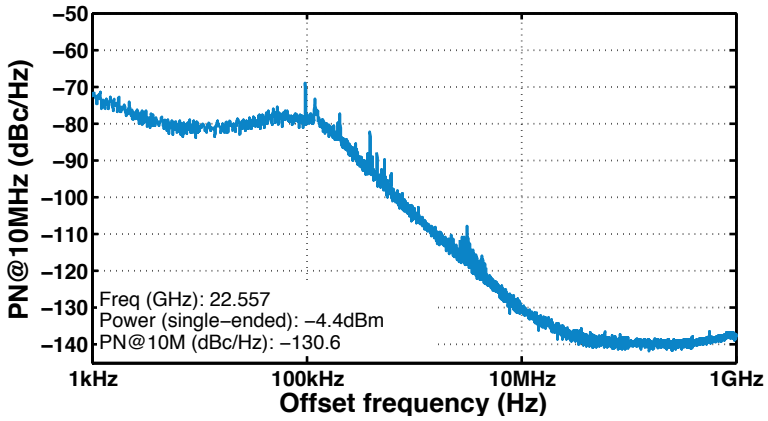

Figure 4.30: Phase noise vs. frequency offset from a 22GHz carrier

The measured tuning range of 21.8−27.5 GHz covers the required range (22.8−26.4

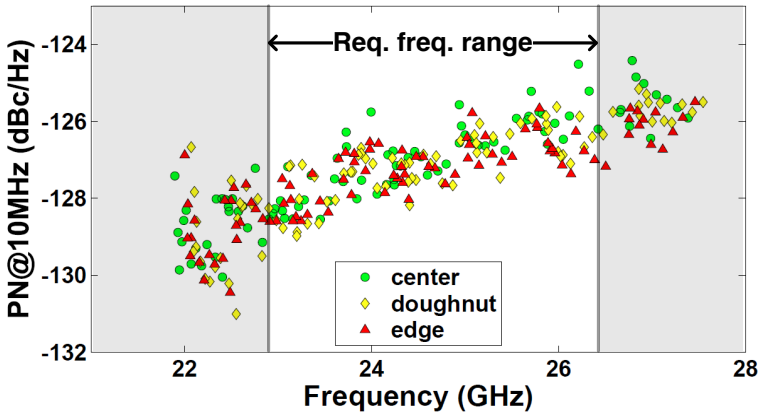

Figure 4.31: Phase noise over the tuning range for different die locations on a wafer

GHz) and is shown in Fig. 4.32 over all 3-bit coarse tuning combinations. The resulting 23% FTR makes this a robust solution that is manufacturable in volume.

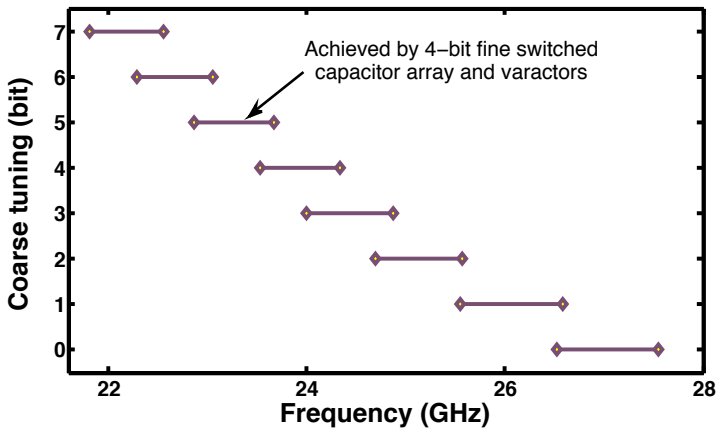

Figure 4.32: Coarse frequency tuning showing frequency overlap

Die to die phase noise variation for 46 dies on a wafer is shown in Fig. 4.33. Measurements using 3 (out of 128 possible combinations) equally spaced switched capacitor

settings on each die are denoted by 3 different markers. All other settings are kept identical. The measured phase noise variation is $\sigma < 0.6$dB across all dies. The PLL's performance robustness across temperature at a fixed capacitor setting is shown in Fig. 4.34. The measured phase noise at 10MHz offset degrades by 2dB from −130dBc/Hz at 5°C to −128dBc/Hz at 85°C. The center-band frequency drops by 30MHz.

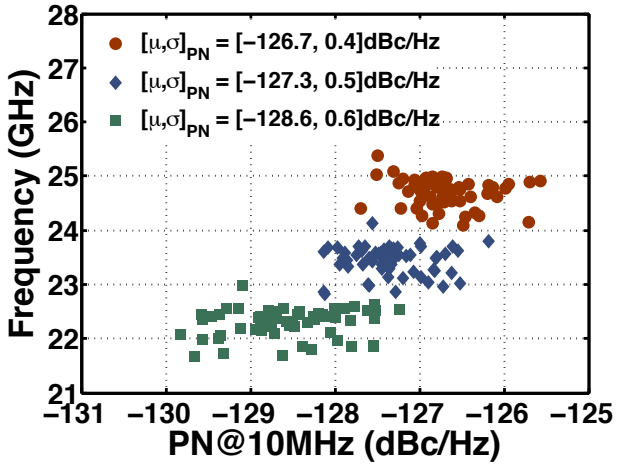

Figure 4.33: Robustness of the LiT VCO design to process variations (die to die)

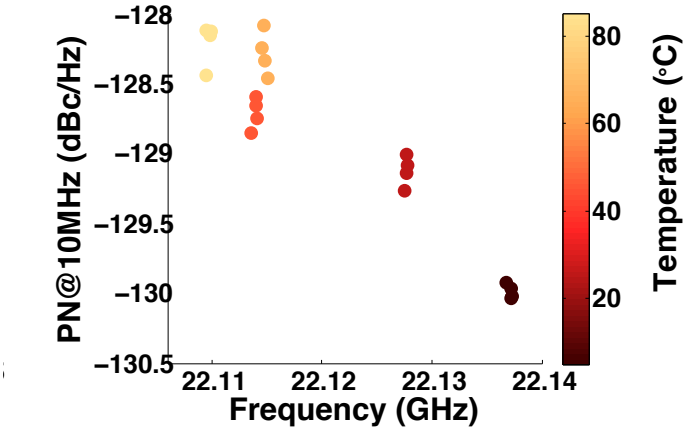

Figure 4.34: Robustness of the LiT VCO design to temperature variations

The effect of varying V_{DD} and $V_{g,bias}$ on the phase noise is shown in Fig. 4.35. For low $V_{g,bias}$, the VCO is in the current limited regime. Phase noise initially improves with increasing $V_{g,bias}$ until the FET becomes non-linear due to triode operation (particularly for lower V_{DD}), after which the phase noise degrades. This demonstrates the effect of triode operation on LiT VCO phase noise. Finally, the PLL's measured phase noise varies by only 2dB across 800mV supply variation (from 0.7 to 1.5 V).

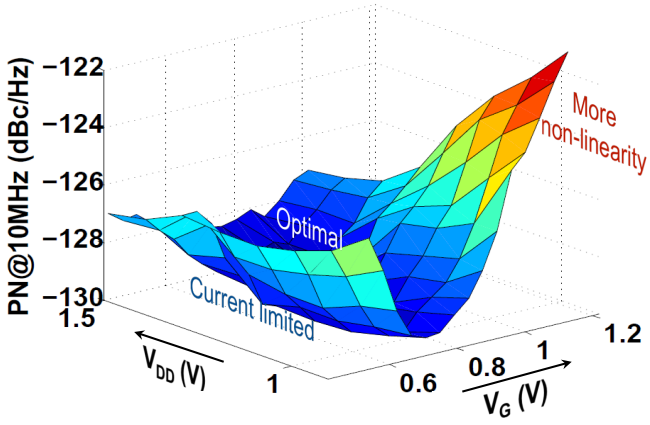

Figure 4.35: Robustness of the LiT VCO design to supply variations

Fig. 4.36 shows the phase noise versus varying gate bias ($V_{g,bias}$) for a particular die (top left subfigure) and for all dies at three different frequency bands (top right subfigure, bottom left subfigure, and bottom right subfigure). The top left subfigure shows the bias optimization for best phase noise. In the other three subfigures, each curve represents a unique die on the wafer. 46 dies were measured. The following observations can be made:

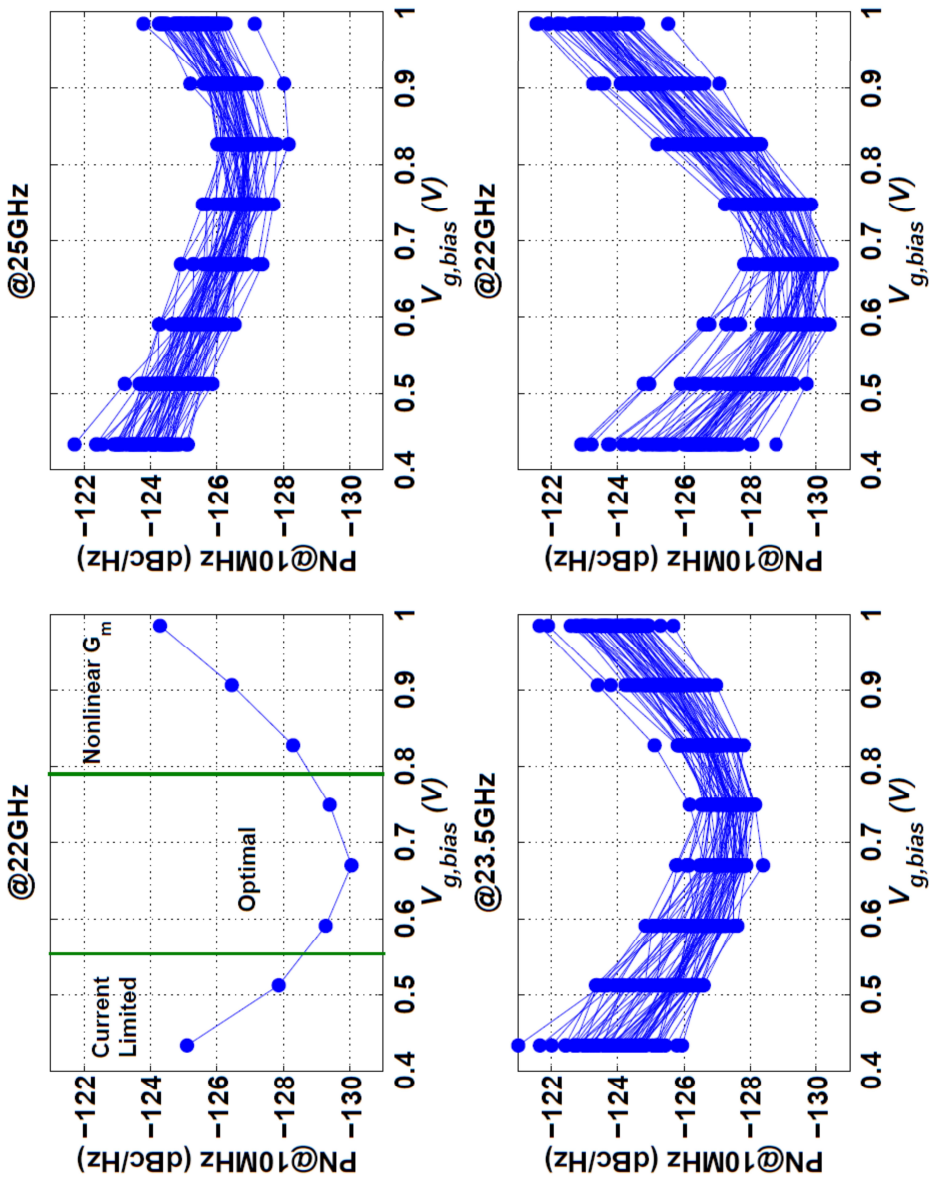

Figure 4.36: Lack of variation of optimal bias point with process and frequency

- There is very little die to die variation in the optimal bias point at each frequency.
- The variation in optimal bias point across the frequencies is 200mV. Consequently, the optimal bias point has low process dependence.

Fig. 4.37 shows the circuit schematic of a peak detector circuit used to measure the output amplitude. V_b denotes the bias voltage for the tail MOSFET of the peak detector. Results from randomly selected dies at 25GHz are shown. Phase noise is plotted versus peak detector output (zeroed with its replica) for varying gate bias ($V_{g,bias}$) values. Each curve represents a die. The following observations can be made:

- The optimal bias for the different dies is approximately equal.
- Due to process variation, there is a weak correlation between phase noise and peak detector output.

Fig. 4.38 shows the phase noise versus peak detector output zeroed with its dummy replica for a particular die across three frequency bands and three temperature values as is shown. Each curve shows the sensitivity of the phase noise to the gate bias voltage ($V_{g,bias}$). The following observations can be made:

- The optimal bias point is almost constant with variation in temperature. This also shows the robustness of the LiTVCO to temperature variation.
- The phase noise as well as the output amplitude degrade at higher temperatures, as expected.

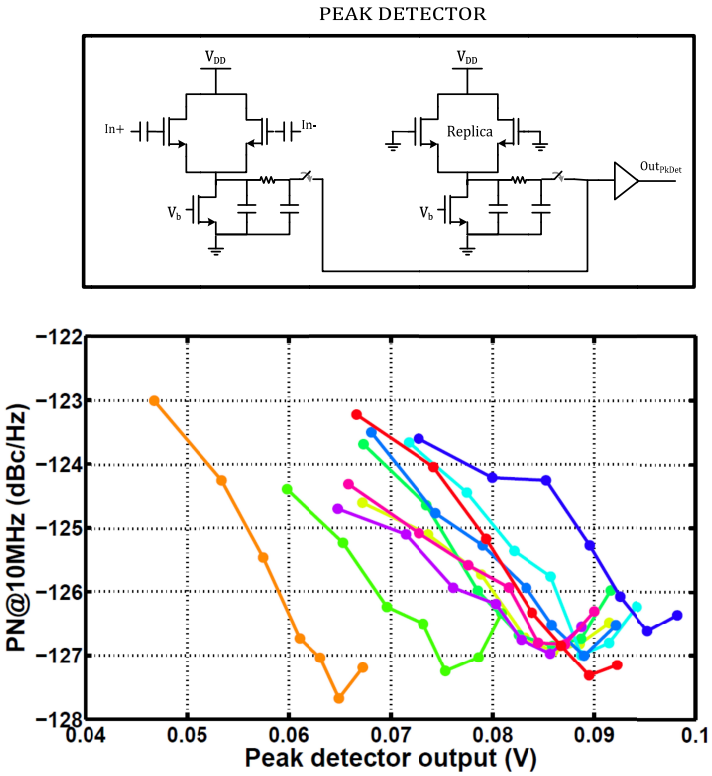

Figure 4.37: Peak detector circuit (above) and measured phase noise vs. peak detector output (below)

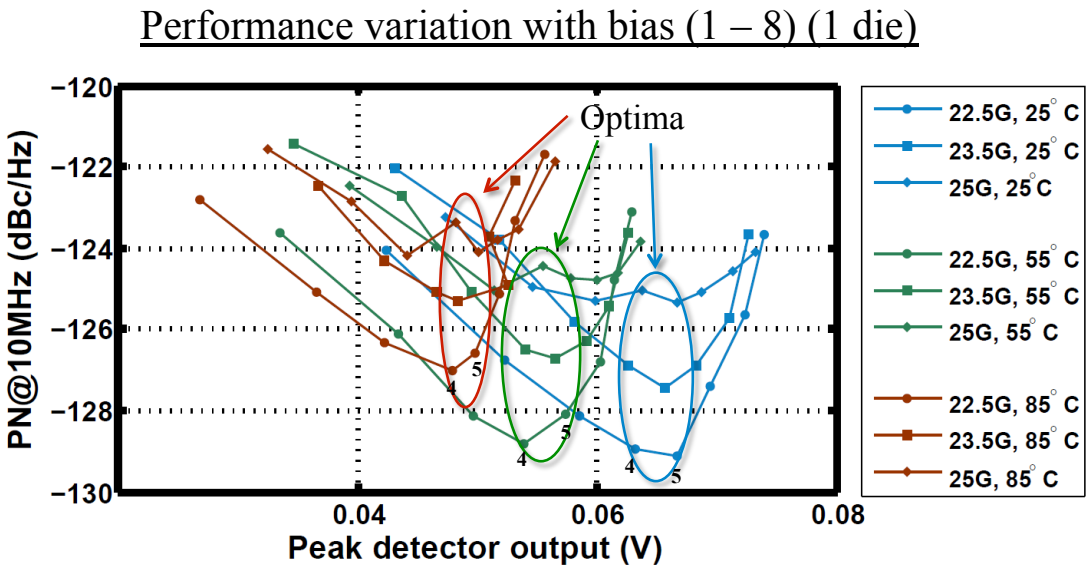

Figure 4.38: Phase noise vs. peak detector output across frequency and temperature

Table 4.2 and Fig. 4.39 compare the LiT VCO PLL performance with other PLLs designed for 60GHz applications. The LiT VCO PLL demonstrates the lowest phase noise of −127.3dBc/Hz at 10MHz from its center frequency, 24.7GHz, a large FTR of 23%, the lowest power consumption of 36mW and the best FOM_T of 188.6dBc/Hz compared to the other designs. Finally, measured results show the design is robust to process, temperature and supply variations and maintains excellent performance across all variations.

Acknowledgements: This work was partially supported by DARPA under AFRL contract # FA8650-09-C-7924. The views, opinions, and/or findings contained in this article/presentation are those of the author/presenter and should not be interpreted as representing the official views

Table 4.2: Table for comparison with other PLLs

Ref	Floyd, JSSC 08 [92]	Osorio, et al. ISSCC 11 [95]	Richard, et al. ISSCC 10 [96]	Murphy, et al. JSSC 11 [97]	Scheir, et al. ISSCC 09 [98]	Pellerano, et al. ISSCC 08 [99]	**This work**
FTR (GHz) (f_{osc}) (GHz)	16.0-18.8 (17.4)	21.7-27.9 (24.8)	17.5-20.9 (19.2)	42.1-53.1 (47.6)	57.0-66.0 (61.5)	39.1-41.6 (40.4)	**21.8-27.5 (24.7)**
PN @10MHz from f_{osc} (dBc/Hz)*	-123.9	-121.0	-126.0 (@21GHz)	-117.5	-95.0	-112.0	**-127.3**
Technology (nm)	130 BiCMOS	45 CMOS	65 CMOS	65 CMOS	45 CMOS	90 CMOS	**32 SOI**
PLL power (mW)	144	40	80	72	78	64	**36‡**
PN @10MHz from 24.7GHz (dBc/Hz)†	-120.9	-121.0	-124.6	-123.2	-102.9	-116.3	**-127.3**
VCO tuning range (%)	16.5	25	17.7	23.1	14.6	6.2	**22.9**
VCO FOM_T (dBc/Hz)§	182.8	185.9	___	186.7	___	___	**188.6**

* Calculated assuming 20dB/decade degradation with offset frequency

† Calculated assuming 20dB/decade degradation with oscillation frequency

‡ Includes the consumption of the micro-controller and sensors for digital calibration and optimization; the VCO consumes 24mW

§ $FOM_T = L(\Delta\omega) - 20 \cdot \log\left(\frac{\omega_0}{\Delta\omega} \cdot \frac{FTR}{10}\right) + 10 \cdot \log\left(\frac{P_{diss}}{1mW}\right)$ where P_{diss} is the VCO power dissipation

or policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the Department of Defense.

4.6 Conclusions

In this chapter, a number of VCO topologies that provide superior phase noise performance compared the LC cross-coupled VCO were presented. These topologies were discussed, and their benefits as well as disadvantages compared against each other. The improvement in phase noise is obtained by increasing the amplitude of oscillation in conjunction with reducing the active device noise injected into the tank. This is achieved by different capacitive feedback techniques that achieve transconductance linearization.

As a proof of concept for one of the topologies, a prototype low phase noise VCO

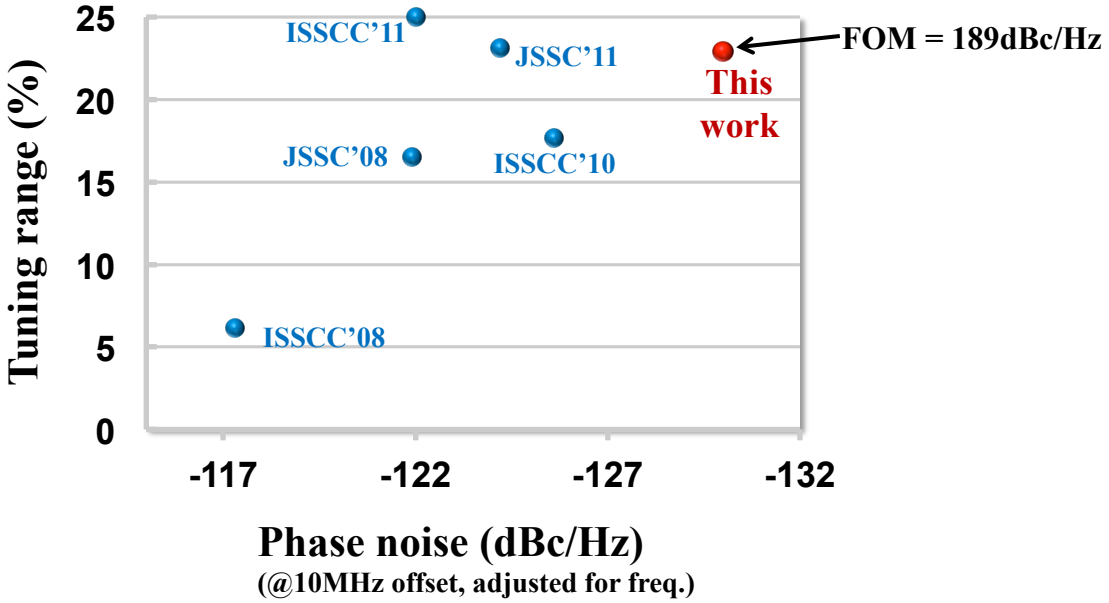

Figure 4.39: Phase noise and tuning range comparison with other PLLs for 60GHz applications

integrated in a dual loop PLL for 60GHz applications was presented. A novel G_m linearization technique was used to achieve a larger oscillation amplitude that, along with lower active device noise, resulted in significantly improved phase noise performance. Additionally, a new capacitor array layout approach was used to reduce interconnect inductance, enabling a large 23% FTR.

Chapter 5

RF Sampling and Signal Processing

5.1 Introduction

As seen in Fig. 6.2 in Chapter 2, this spectrum sensing architecture uses an RF sampler followed by discrete time signal processing in the analog domain. Specifically, passive charge domain computations are utilized for signal processing followed by digitization. For RF sampled processors, the RF sampler remains a significant bottle-neck. However, recently, it has been shown that the sampling requirements for narrowband systems employing RF samplers are significantly more relaxed [105]. Moreover, it is possible to use charge domain sampling due to its superior robustness to jitter [106].

A general architecture for an RF sampling receiver is shown in Fig. 5.1. This use of

RF samplers and subsequent discrete-time processing provide a number of advantages in deep sub-micron CMOS processes [107]. Recently, other discrete time radio receivers using RF sampling have been demonstrated using CMOS technology for Bluetooth [48], GSM/GPRS [47], WLAN [108], and SDR-type applications [109, 110]. The advantages and motivation for the use of sampled charge processing is briefly reviewed here.

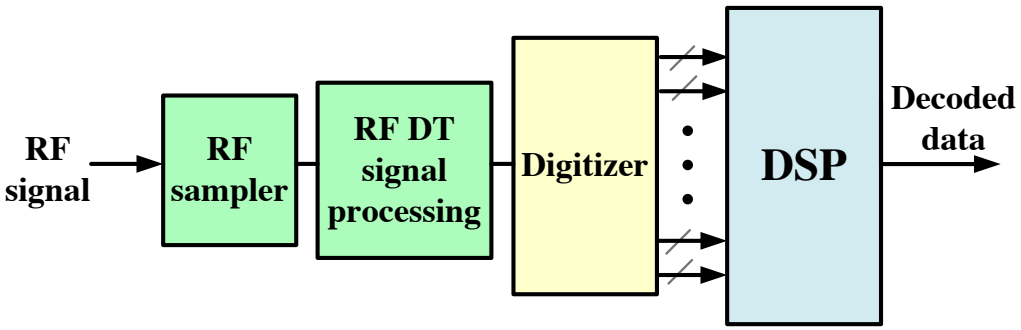

Figure 5.1: A spectrum sensing architecture based on RF sampling and signal processing

5.2 Sampled Charge Processing

Signal sampling and variable-rate analog signal processing is performed in the charge domain due to the inherent benefits of including a built-in anti-alias filter into the sampler [111], robustness to jitter [106], and the ability to vary the resulting filter notches by simply varying the integration period [46, 106, 112–114].

Many of the benefits of the discrete time FFT architecture are based on the use of passive discrete-time charge based computations. This is best illustrated with the help of an example design. The passive switched-capacitor shown in Fig. 5.2 is able to

operate at RF sampling speeds [46, 112–114]. In this circuit the input signal is sampled progressively in time ($\phi_1 - \phi_n$). After N clock periods the averaged output is sampled onto the capacitor C_s, which has previously been discharged. The complete circuit implements an N-tap FIR filter that is decimated by N. Interestingly, if the capacitor C_s is not discharged between each rotation then the circuit implements an N-tap FIR filter combined with a first-order IIR filter that is decimated by N. Note there is no active element (i.e., amplifier) in this circuit. The circuit consists only of switches and capacitors, so the maximum sampling rate is only dependent on the RC settling times of the switches. Additionally, the only power dissipation, other than that required for sampling the signal from the input, is due to the charging and discharging of transistor gate capacitors in a very digital-like way. As a result, a variety of functions on the sampled signal can be computed very fast and using minimal power.

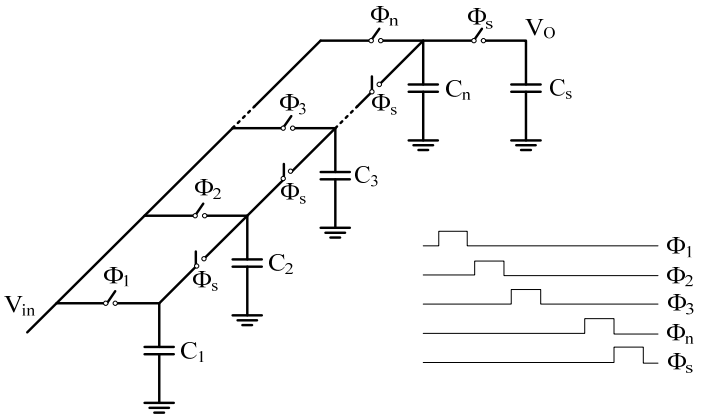

Figure 5.2: Switched capacitor implementation of a passive N-tap FIR with a decimation by N

5.3 Passive Computations

For performing any linear function, addition and multiplication operations need to be performed. Note that all passive switched capacitor operations are destructive in nature. Therefore, once an operation is performed, the input values are lost. For performing multiple operations on a single input, multiple copies of the input need to be maintained. We have explored different techniques to perform these operations using passive switched capacitor circuits. We have compared these techniques based on their robustness to non-idealities, ease of implementation, power consumption, speed, etc. and selected suitable architectures for these computations.

5.3.1 Addition

1. Parallel connection: Using passive switched capacitors, addition may be performed by sharing the charges on two participating capacitors by connecting them in parallel as shown in Fig. 5.3. The result of this operation is the average value $(V_1 + V_2)/2$ of the input voltages V_1 and V_2, which is a scaled version of their sum operation. Also note that two copies of the output are obtained and these can be used for two independent operations later. However, the operation inherently attenuates the output by half. From an implementation perspective, use of parallel capacitors allows the sharing of one plate (ground plate) for all the capacitors. This can greatly reduce the parasitic capacitance and resistance of the capacitor, and the area of the overall implementation.

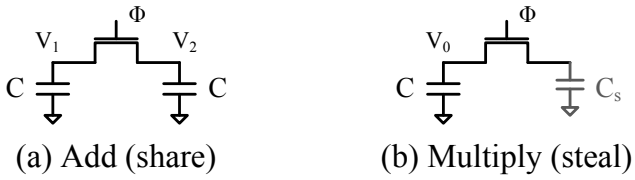

Figure 5.3: Techniques for charge domain addition and multiplication operations

2. Series connection: An alternative technique is to connect the capacitors in series. The result of this operation is the sum $(V_1 + V_2)$ of the input voltages V_1 and V_2. In this scheme, it is possible to use slightly delayed clock phases for the top and bottom plate switches in order to make the charge injection independent of the input voltage [115]. However, in this latter technique, switches are required both on the top and bottom plate, thereby increasing the power consumption in this circuit. Also, the two switches placed in series halves the speed of this circuit for identical switch sizes. Moreover, only one output (which can be used for exactly one subsequent operation) is obtained.

5.3.2 Multiplication

1. Charge stealing: Multiplication in the charge domain can be performed by scaling the voltage on a capacitor using a share operation with another known capacitor (stealing capacitor). The charge on the stealing capacitor is not utilized later. The overall operation causes a sub-unity scaling on the original value. The scaling factor for a capacitor of value C and a stealing capacitor of value C_s is given by $m = C/(C + C_s)$.

Fig. 5.3(b) shows a scaling operation using a stealing capacitor of size C_s with no initial voltage on it. After the sharing operation, the final value on the capacitor with initial value V_0 becomes $V_0 \cdot C/(C + C_s)$. C_s can be chosen appropriately to obtain a particular scaling factor. Note that though this technique is capable of performing both sub-unity scaling and multiplication with a known attenuation, at least one of the operands needs to be known in advance for this implementation. In case variable capacitors are utilized, dynamic operands can also be used.

2. Pulse-width modulation (PWM): Another technique to perform multiplication using passive switched capacitors is to modulate the turn on time of the switch and perform an incomplete share operation with a fixed stealing capacitor. The duration of the operation determines the multiplication factor. It is possible to multiply two unknown operands using this technique. However, considering the non-linearity in the resistance and the share operation, the errors caused by this technique make it unusable. However, the concept can be used to devise another PWM scheme which allows complete settling thereby making it more reliable. In this modified technique, the switch can be turned on using a sequence of randomly placed pulses and sharing the capacitor charge using a small stealing capacitor for each clock cycle. The stealing capacitor is discharged at the end of each cycle. Complete settling is allowed in each cycle. The total number of on-pulses determines the amount of scaling. Maximum scaling is obtained when all the clock cycles have on pulses, while no scaling is obtained when all the clock cycles

have off pulses. Although this technique is relatively accurate, and is able to handle dynamic operands, it is slow and consumes more power than the charge stealing technique. Also, depending on the accuracy required, the attenuation is considerable.

3. Current domain: If the charge is converted to the current domain, a single, variable-duration pulse PWM scheme can be used to perform multiplication. Also, multiplication would not entail an inherent attenuation. However, the technique is very power hungry, and the accuracy of the transconductance amplifier that translates from charge to current domain needs to be very high. We have focused on the charge stealing concept for performing multiplications due to their low power characteristics. For most linear algebra problems, multiplication using fixed coefficients is sufficient, and this technique lends itself easily to such applications.

5.3.3 Switching Schemes

For reducing the attenuation, different sharing techniques can be used. Two capacitors are shared as shown in Fig. 5.3(a). We can combine a share followed by scaling into a single operation by connecting 3 capacitors (2 with input samples and 1 empty) and sharing their charges. This can be performed in different ways using 2 switches or 3 switches as shown in Fig. 5.4(a,b,c). Additionally, it is possible to reduce the total settling error by using 3 appropriately sized switches in the scheme of Fig. 5.4(c). Similarly, different

schemes may be used for sharing 4 and 5 capacitors for scaling by complex factors of the form '$c + c \cdot j$ as shown in Fig. 5.4(d,e,f,g,h,i). While some schemes (Fig. 5.4(b,d,g)) ensure symmetry, others (Fig. 5.4(a,e,h)) offer faster computations for the same/lower number of switches as compared to Fig. 5.4(b,d,h). Some schemes (Fig. 5.4(c,f,i)) provide both speed and symmetry at the cost of a larger power dissipation. Also, different schemes, with their appropriate switch sizes, provide different trade-offs with regard to noise contribution, charge injection error, clock feed-through error, etc.

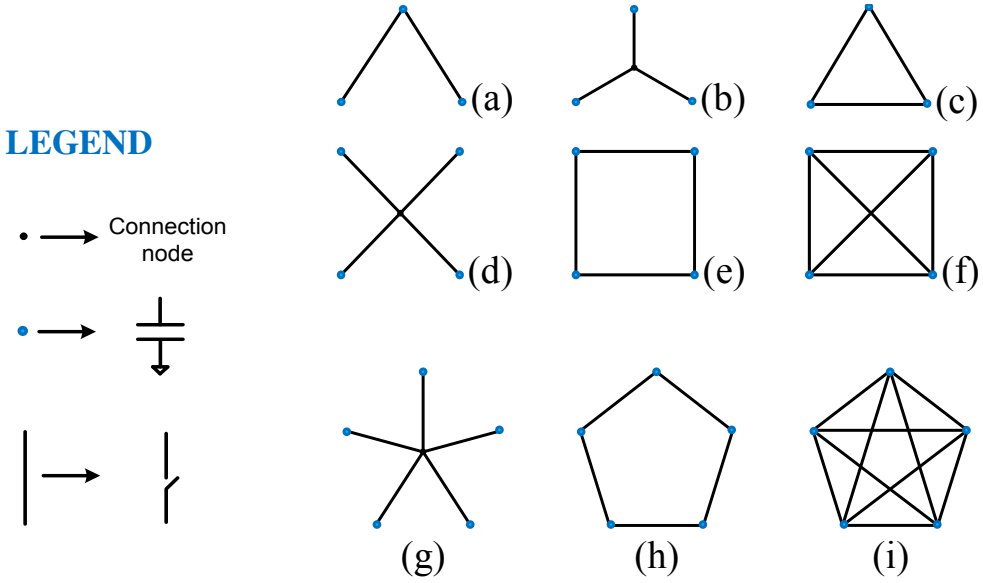

Figure 5.4: Different switching topologies for charge domain operations

5.4 Non-idealities

Several non-idealities haunt passive switched capacitor circuits. The problem of non-idealities is aggravated by the absence of a virtual ground node unlike in op-amp based

switched capacitor circuits. The effect of sampling clock jitter in passive switched capacitor circuits has been analyzed. We use current mode sampling to make the circuits tolerant to sampling clock jitter [106]. Two important non-idealities: clock feed-through and charge injection, become a nuisance in the absence of a virtual ground node. Traditional circuit techniques such as bottom plate sampling are consequently difficult to implement. Also, poor matching between nMOS and pMOS switches, and the reducing difference between V_{dd} to V_{th} in scaled technologies makes the use of transmission gate switches less effective for mitigating these non-idealities. The noise in the system is dominated by the $k\mathrm{T}/C$ noise of the RC filter formed by the switch-capacitor combination. Moreover, for a multi-stage switched capacitor operation, the sampled noise voltages from one stage recombine in the later stages. These combining noise samples in a particular stage are correlated, and therefore, the final noise becomes a complicated function of the noise sampled at each stage of the switched capacitor operation. The switch resistance (along with the capacitors capacitance) determines the settling time constant. However, the switch resistance is inherently non-linear and input signal dependent. Consequently, in the case of high speeds of operation, incomplete settling can cause significant signal dependent errors in computations.

5.5 Example Comparison

In order to appreciate the kind of power gains available in an analog signal processing architecture implemented using passive switched capacitor circuits against a digital

signal processing architecture, we will consider two example architectures, and compare their power performance. Let us consider the simple and almost ubiquitous architecture shown in Fig. 5.5 (Arch A).

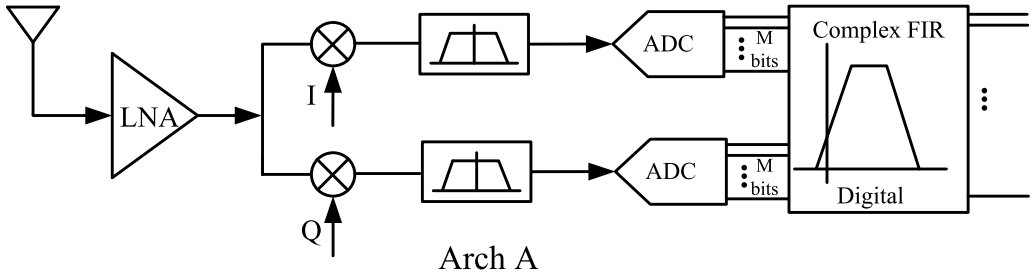

Figure 5.5: Receiver architecture incorporating digital domain FIR filtering

As shown in the figure, the incoming RF signal is first amplified, and then down-converted into real and imaginary signal paths. Each of these signal paths is low-pass filtered and digitized. The digital signal is then filtered using a complex FIR filter in the digital domain to sift out the desired signal.

The same function can be implemented using an alternative architecture where the complex FIR filtering is performed in the analog domain as shown in Fig. 5.6 (Arch B).

Let us now perform a comparative analysis of the power dissipation at a particular technology node. Later we will also compare the performance across technology and develop architectural insights based on the results. Also, since the architectures differ only in the last two elements, namely the complex FIR filter and the ADC, we will consider only these for our comparison. A complex N-tap FIR filtering operation can be mathematically represented through the following equation:

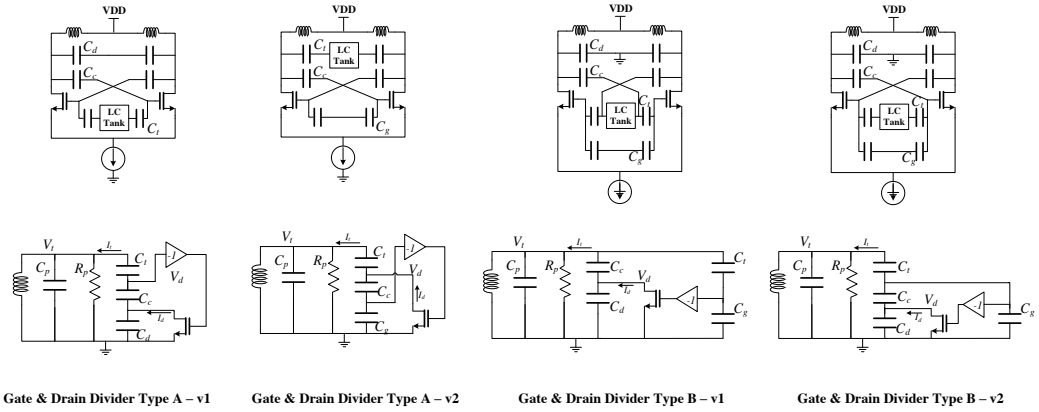

Figure 5.6: Receiver architecture incorporating analog domain FIR filtering

$$y[n] = \sum_{i=0}^{N} b_i x[n-i]$$

where the coefficients b_i and the samples x_{n-i} are complex.

Power Analysis of Arch A:

1. N-tap digital FIR filter: For the power analysis, we will assume an M-bit ADC output (representing each $x[n]$) and C-bit coefficients (representing each b_i). The power dissipation of a complex FIR filter implemented in digital can be analyzed by breaking down the total power in terms of the power used for multiplications (P_{mul}), accumulation operations (P_{acc}) and memory operations (P_{mem}): $P_{FIR} = P_{mul} + P_{acc} + P_{mem}$. These operations can be written down in terms of the more fundamental operations of addition, shifting and memory access.

$$P_{mul} = N(M+1)(C-1)(P_{add} + P_{shift})$$

$$P_{acc} = (N-1)(M + C + \log_2 N)P_{add}$$

$$P_{mem} = (N)(M+C)P_{sram}$$

Letting E_{gate} be the energy dissipation for a minimum sized nMOS transistor gate in 65nm technology. From an analysis of a ring oscillator in IBMs 65nm technology, E_{gate} is found to be 1.39nW/MHz. Using this value and assuming a static CMOS implementation, P_{tot} is calculated for $M = 12$ and $C = 6$ and plotted in Fig. 5.7 for different tap numbers (N).

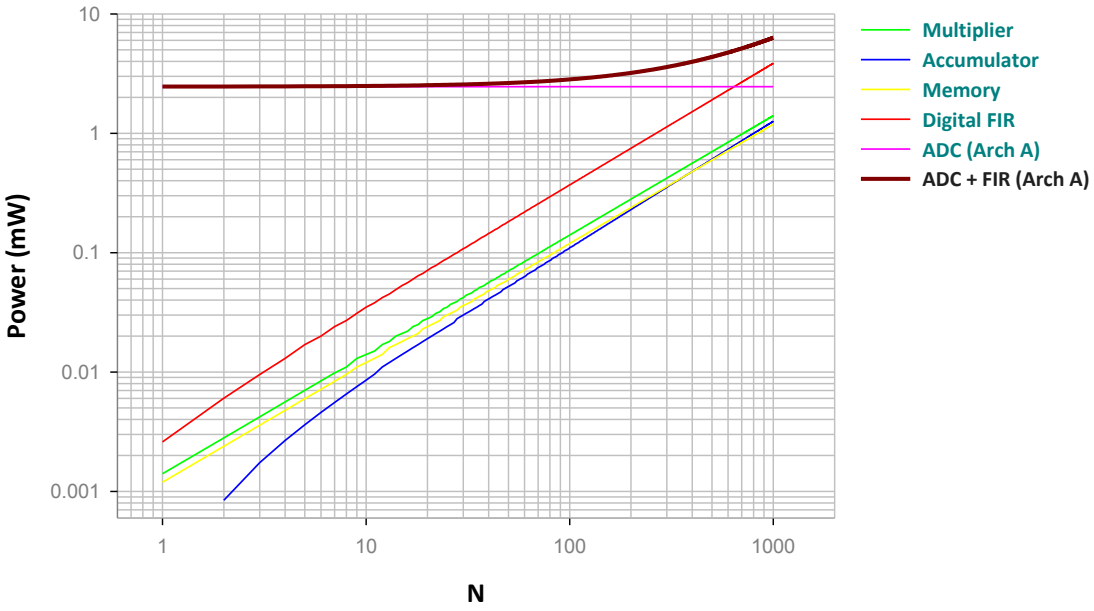

Figure 5.7: Power dissipation for different functions in Arch A vs. number of taps in the FIR filter

2. M-bit ADC: The power consumption for the ADC can be roughly calculated from the following figure of merit expression:

$$FOM = \frac{P_{diss}}{f_s 2^{ENOB}}$$

where P_{diss} is the power consumption of the ADC, f_s is the sampling frequency and ENOB is the effective number of bits. Using an empirical minimum value of 0.15pJ/conversion for the FOM based on previous literature, the power dissipation for a 12-bit ADC with an f_s of 1GHz is found to be 2.46W. The total power for the FIR and ADC in Arch A with individual breakdowns for the different functionalities involved is shown in Fig. 5.7.

Power Analysis of Arch B:

1. N-tap analog FIR filter: For implementing an N-tap analog filter, we consider the 2-tap example architecture shown in Fig. 5.8.

   The filtering operation is performed through three steps:

   (a) Sampling charge onto a capacitor and scaling the sample with a coefficient b_i: This is performed using a G_m cell and series capacitors such that the voltage at the intermediate node is a scaled version of the sampled voltage. Since each sample corresponding to a time instant is used in the computation of N outputs, each time with different coefficients, in an N-tap FIR filter the charge at each time instant is sampled onto N capacitor-pairs, each with

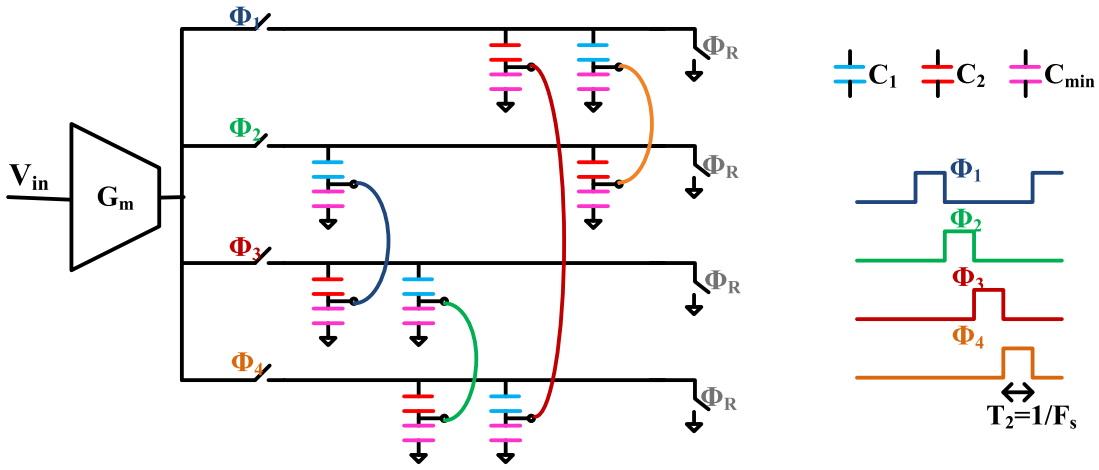

Figure 5.8: Architecture for analog 2-tap FIR using passive switched capacitors

appropriate capacitive ratios. Therefore, if a capacitor combination of C_1 and C_{min} is used for charge sampling, $b_i = \frac{C_1}{C_1+C_{min}}$.

(b) Accumulating the voltages from different samples: This can be performed by shorting together one copy of the charge sampled and scaled from each time instant. The accumulated voltage then becomes $V_{acc} = \frac{1}{N}\sum_{i=1}^{N} b_i x[n-i]$ which is a scaled version of the required result.

(c) Resetting the capacitors: The capacitors need to be reset so as to be ready for the next iteration.

As a result of the two additional steps of accumulation and reset, two extra rows of capacitors are required (over and above the N rows) to accommodate a pipelined operation as shown in Fig. 5.8. The curved lines denote the accumulation operation in the appropriate clock cycle as depicted through its color. The reset switches

for the intermediate nodes have not been explicitly shown to reduce clutter. For a calculation of the power dissipation of the system, a constant frequency of operation (1GHz) has been assumed (with varying N). For this, the switches had to be sized in order to maintain a constant RC time constant for all operations. The designed switch sizes allow a maximum frequency of operation of 10GHz for the 65nm technology used for calculations, and therefore, a 1GHz operating frequency should be easily achievable. The architecture in Fig. 5.8 shows only real number operations. Complex number operations can be easily performed by using a few additional multiplication and scaling operations.

The power dissipation of this system can be divided into two parts:

(a) Power dissipated for charging the capacitors during the sampling operation

(b) Power dissipated in charging the transistor gates for all the switching operations

These values are calculated for the architecture shown, and the results are plotted in Fig. 5.9. The total power dissipated in the digital FIR block is shown in the same graph for comparison. As seen from this analysis, the power dissipation in the analog implementation of the N-tap FIR filter is more than two orders of magnitude lower than in the digital implementation. In fact, this result is hardly surprising. While, in analog, each voltage is represented and computed using a few capacitors connected in series or parallel, the same operation in digital involves

an M-bit sample and C-bit coefficient and a corresponding complex algorithm to compute the result of an addition or scaling operation.

2. K-bit ADC: The ADC requirements in Arch B are different from those in Arch A. Since the input to the ADC is already filtered, the dynamic range requirements for the ADC (K bits, $K < M$) are reduced considerably. The blocker attenuation is assumed to be proportional to the number of taps used. Also, the bandwidth seen by the ADCs is now reduced due to this filtering operation and a reduction factor of $\frac{1}{2}$ has been used for calculations.

The graph below in Fig. 5.9 shows the power requirements for the FIR filter and ADC for Arch B. The total requirement for these two components of Arch A is also plotted for comparison purposes. As seen in the plot, for lower number of taps, the ADC power requirements remain high (comparable to the digital power) and therefore, does not provide a substantial benefit. However, the power requirement for Arch B reduces with N and is optimal at about 55 taps with a total power dissipation of 1mW. Comparing this power with that in Arch A ($> 2.5W$), a $600X$ improvement in power dissipation is achievable.

5.6 Conclusions

In this chapter, architectures based on RF sampling and high-speed charge domain processing were considered. RF sampling can be a viable technique (with feasible jitter

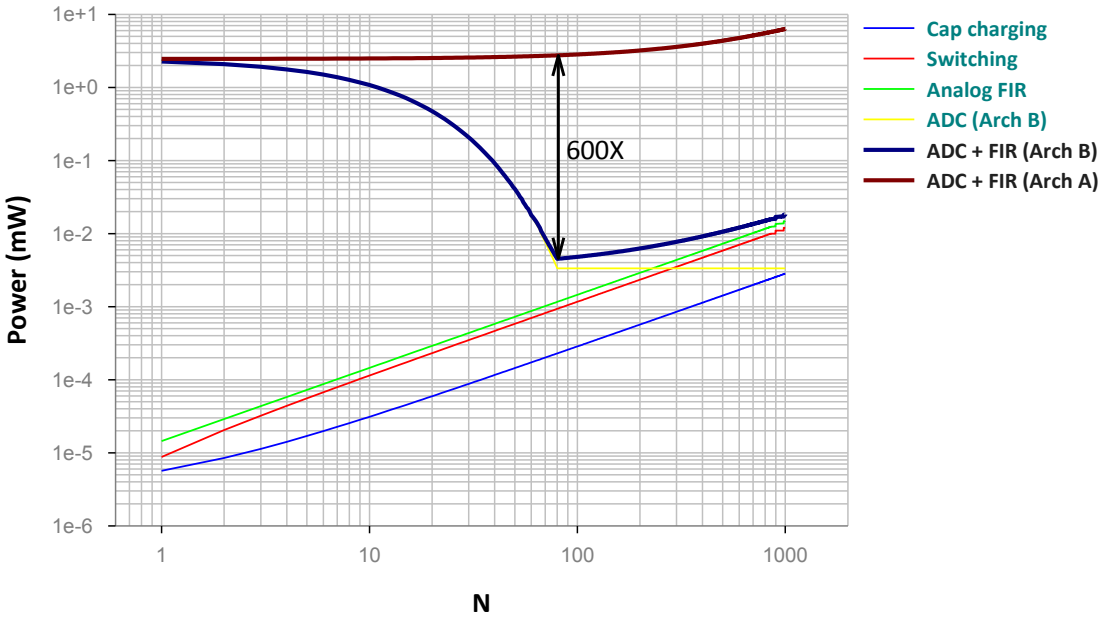

Figure 5.9: Power dissipation for different functions in Arch B vs. number of taps in the FIR filter; also shown is the total (ADC and FIR) Arch A power dissipation for comparison purposes

requirements) for narrowband systems, or if charge domain samplers are employed. The fundamental concepts of sampled charge processing, and different charge domain computation techniques were introduced. An example comparison between a system relying solely on digital signal processing was compared to one that employed some filtering in the sampled charge domain. A large amount of power savings for wideband systems is expected.

Chapter 6

CRAFT: Charge Re-use Analog Fourier Transform

6.1 Introduction

In this chapter, we discuss the design of a wideband digitizer introduced in Fig. 2.7 in Chapter 2. The digitizer is based on a frequency domain divide and conquer approach that could be followed by multiple ADCs that digitize the input in the frequency domain. For this, we propose to utilize an RF sampler followed by a discrete-time Fourier transform engine to perform channelization of the wideband RF input. The use of RF samplers and subsequent discrete-time processing, prior to digitization, provide a number of advantages in deep sub-micron CMOS processes [105, 107, 116]. Recently, discrete time radio receivers using RF sampling have been demonstrated using CMOS

technology for Bluetooth [48], GSM/GPRS [47], WLAN [108, 117], and SDR-type applications [109, 110].

In this work, we use a discrete time DFT as a functionally equivalent linear phase N-path filter (see Fig. 6.1) to perform channelization [118]. This scheme reduces both the speed and dynamic range of the ADCs, and, by virtue of being minimal phase, allows for simple reconstruction in the digital domain. A few current based analog DFT filters have been designed recently [119, 120]. However, these designs are speed-limited, and consume significant power (Table 6.3) minimizing the overall gains. Additionally, these use active devices in the signal processing unit, and are therefore expected to be have a higher non-linearity at high operating speeds. In this work, we perform a charge domain DFT to significantly reduce the power overhead, making a DFT-based wideband digitizing front-end feasible.

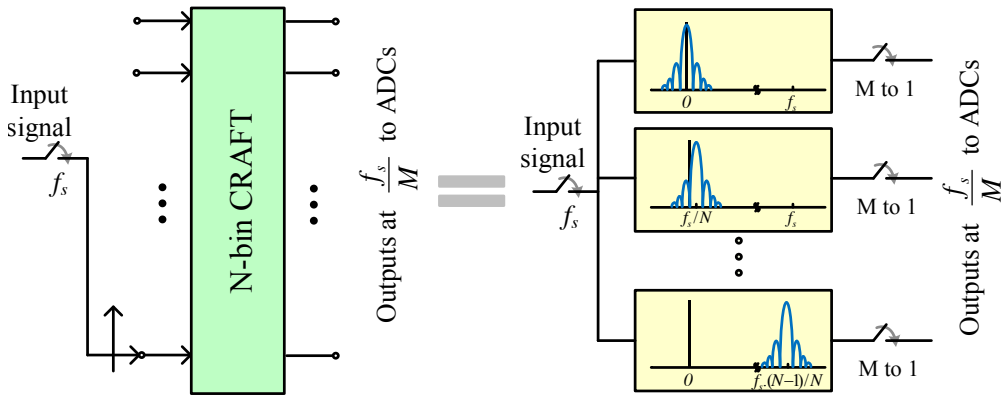

Figure 6.1: The frequency domain functional equivalent of a DFT comprising of N-path bandpass filters

This paper describes the design of such a DFT filter, CRAFT (Charge Re-use Analog Fourier Transform), based on passive switched capacitors. It performs an analog domain 16 point DFT running at input rates as high as 5GS/s and uses only 12.2pJ of energy per conversion. The design was first presented in [121, 122]; this paper further includes the details of the implementation, modeling and mitigation of non-idealities, design methodology and optimization, and additional measurement results.

As shown in Fig. 6.2, an architecture with a CRAFT RF front-end reduces the ADCs' speeds by N, at a negligible power overhead. The ADC dynamic range is also reduced due to the removal of out-of-band signals per ADC. The impact of the CRAFT front-end on the ADC input bandwidth and dynamic range is shown in Fig. 6.3. The expected ADC requirements are plotted amongst measured ADC implementations [60].

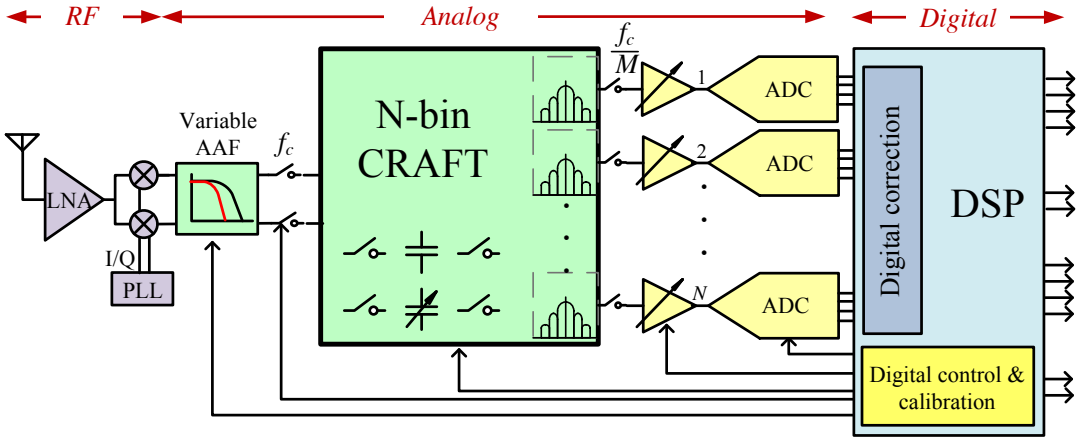

Figure 6.2: An envisioned SDR architecture enabled by CRAFT

For dynamic range calculations, signals are assumed to be distributed in frequency as shown in Fig. 6.4(a). The peak to average power ratio (PAPR) in this case is 10.

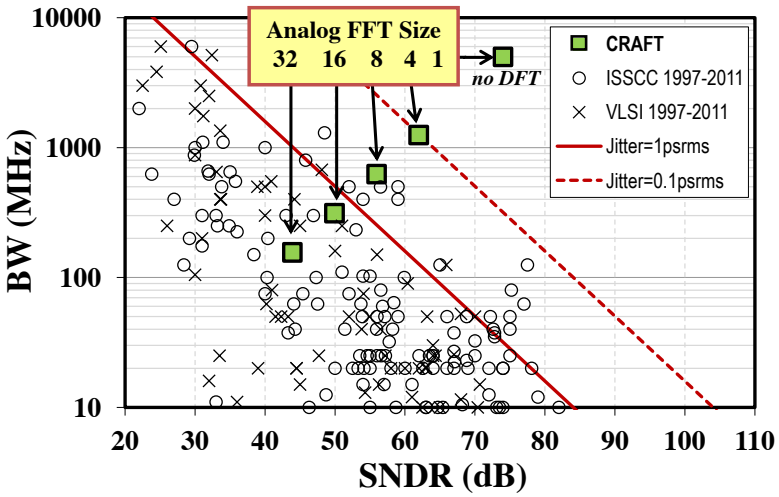

Figure 6.3: Feasibility of ADCs vs. bin size of the CRAFT front-end

In comparison, breaking it up into 5 equal frequency channels, one of which is shown in Fig. 6.4(b), reduces the PAPR to 2 (factor of 5 reduction). In general, an N path channelization of spread signals reduces the dynamic range by N times[1] , causing an N times dynamic range reduction for the ADCs.

As seen, the CRAFT front-end brings the required ADC specifications from being infeasible (top right corner) toward being achievable (bottom left half). By solving this critical broad-band digitizing problem, CRAFT is expected to enable the realization of wideband SDRs in the future.

The paper is organized as follows. First, the conceptual charge-based passive design utilized in CRAFT is described. The CRAFT implementation is then discussed in detail, including a brief description of the on-chip circuitry that interfaces the design

[1] Additionally, unlike the perfect channelization example shown, even in the case the signals are congested, but not directly on-bin, the DFT provides sufficient signal spreading for this approximation to roughly hold.

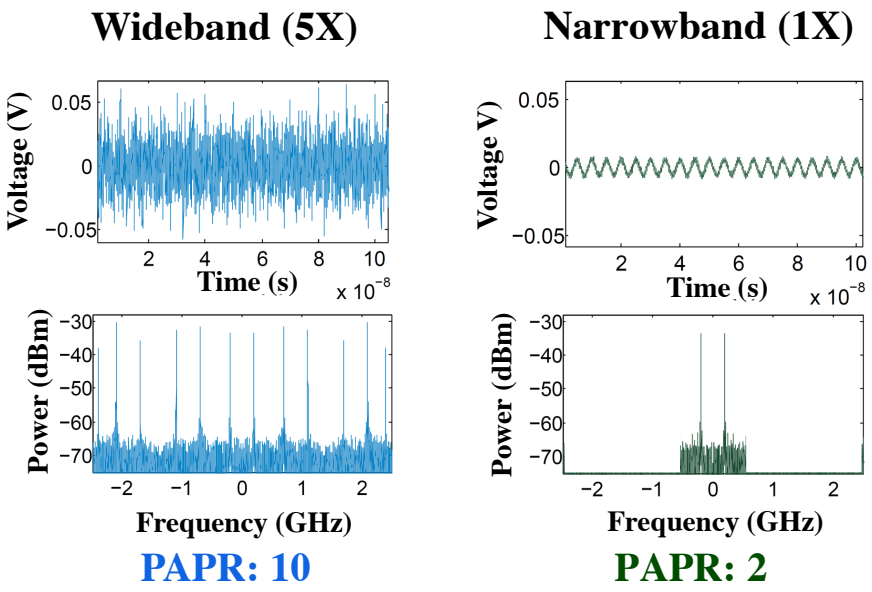

Figure 6.4: Wideband and channelized signals showing dynamic range reduction

for performance measurement. The circuit performance is ultimately limited by several non-idealities. Modeling of these non-idealities, and design techniques and circuit optimization to mitigate them are discussed. Measurement results from a prototype 16 point 5GS/s design are then presented, followed by conclusions.

6.2 CRAFT Design Concept

CRAFT operations are based on charge re-use. Once sampled, the charge on a capacitor is shared and re-shared with other charge samples such that the resulting mathematical manipulation is an in-place DFT. By basing the design only on toggling switches (transistor gates), low power and high speeds are made feasible. Additionally, the power scales with frequency, supply, and technology in a digital-like fashion.

Fig. 6.5(a) shows the diagrammatic representation of the radix-2 FFT algorithm used in CRAFT.

As seen, the FFT uses only two types of operations: addition, and multiplication by twiddle factors. These twiddle factors are shown as powers of W in the figure, where $W = e^{-\frac{2\pi j}{16}}$, and are equi-spaced points on an unit circle in the complex plane as shown in Fig. 6.5(b). As a result, for these scaling factors, W^k, $\Re\{W^k\} \leq 1, \forall k$ and $\Im\{W^k\} \leq 1, \forall k$. Since passive computations inherently attenuate the signal, these operations are particularly suited for sub-unity scaling.

For performing the CRAFT operations, the following charge domain computations are used. Addition is performed by sharing the charges on two capacitors as shown in Fig. 6.6(a). Multiplication is performed by stealing charge away from a capacitor using a suitably-sized stealing capacitor (C_s) as shown in Fig. 6.6(b). These 2 simple operations form the basis of all operations performed in CRAFT.

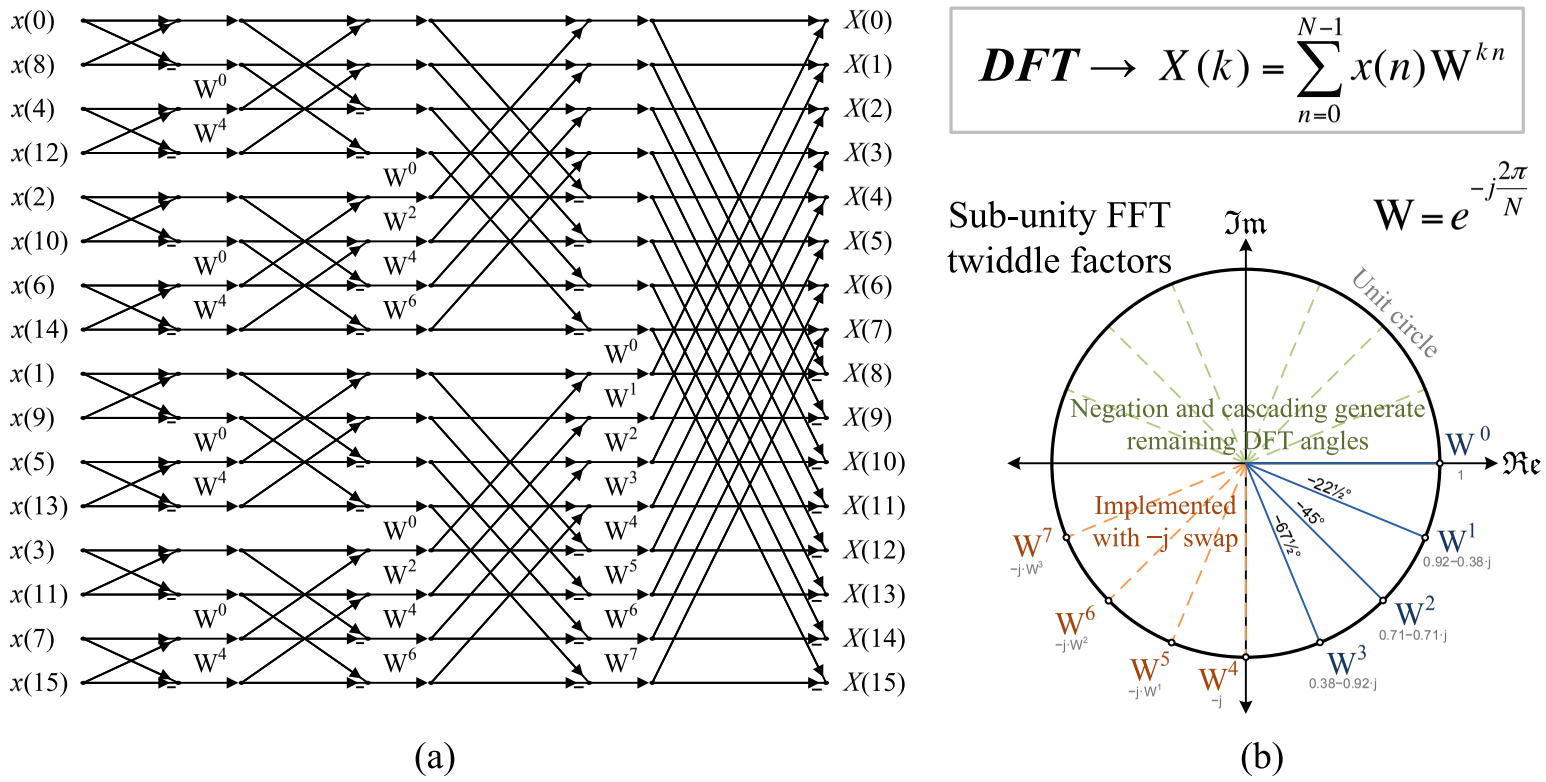

Figure 6.5: (a) A 16 point radix-2 FFT signal flow diagram using the decimation-in-time algorithm, and (b) DFT sub-unity twiddle factors on the unit circle

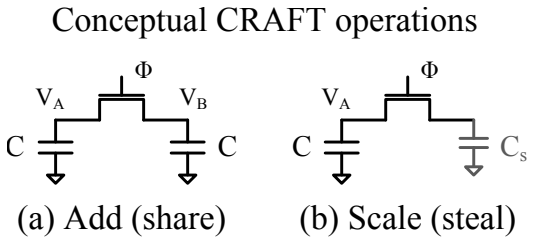

Figure 6.6: Conceptual implementation of addition and sub-unity scaling operations on which CRAFT is based

The 16 point, radix-2 CRAFT operation can be represented as a linear matrix transform:

$$\mathbf{X} = \frac{1}{k}\ \mathbf{F}\ \mathbf{x} \tag{6.1}$$

where k is the scaling factor due to attenuation inherent in the CRAFT operation. Also, note that any linear transform with a fixed matrix can be performed using the addition and multiplication techniques outlined above. The result is a scaled version of the ideal transformed output.

These operations are utilized in the CRAFT processor as shown in Fig. 6.7. Row 1 shows the 'Share' operation. Subsequent sub-unity twiddle-factor multiplication, if required, is performed by charge stealing as shown in row 2: 'Share and multiply' operation in Fig. 6.7. Negation is performed by swapping the positive and negative wires as shown in row 3, while multiplication by j $(= \sqrt{-1})$ is performed by swapping the appropriate wires of the real and imaginary components as shown in row 4 of Fig. 6.7. These techniques are extended to perform complex multiplication, as shown in row 5. An example butterfly operation using multiple complex multiplications is also shown in row

6. Note that each butterfly operation requires 2 clock phases. However, in the CRAFT engine, we optimize the operations such that only 5, instead of 8, clock phases are utilized for the 4 butterfly stages. This optimization is discussed further in Section 6.3. The optimized butterfly blocks are then used to construct the 16 point CRAFT engine shown in Fig. 6.8. As shown, the following different types of butterflies are utilized: 1 stage share (Type 1), 1 stage scalar multiply (Type 2A), 1 stage complex multiply (Type 2B), 1 stage scalar multiply (Type 3A), 1 stage complex multiply (Type 3B), and 2 stage complex multiply (Type 3C). The circuit schematics for these butterflies is shown in Fig. 6.10.

The in-place CRAFT computations are inherently destructive; therefore, multiple copies of each data value are required for multiple operations. Since a radix-2 FFT algorithm performs the DFT with minimum number of operations per operand (less copies required), it is selected for CRAFT.

Conceptually, the FFT is computed as follows: The input signal is sampled onto capacitors. Since each input is operated on twice in an FFT butterfly, and twice for complex operations, 4 copies of each sample are required. Also, considering I, Q (= 2) and differential (= 2) inputs, the 16 point FFT requires 16 x 2 (complex math) x 2 (butterfly branches) x 2 (I, Q) x 2 (differential) = 256 sampling capacitors.

Function name	Symbol used	Implementation
Share	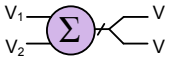	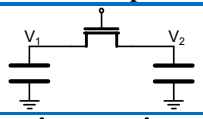 $V = \frac{V_1 + V_2}{2}$
Share & multiply	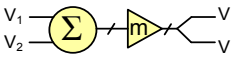	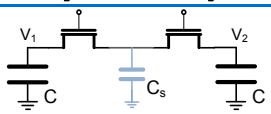 $V = (V_1 + V_2).m$ $m = \frac{C}{2 \cdot C + C_s}$
Negate	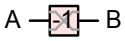	$A \times (-1) = B$ 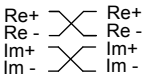
Multiply by '*j*' $j = \sqrt{-1}$	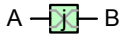	$A \times j = B$ 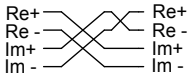
Share & complex multiply	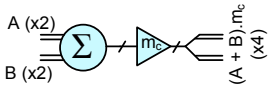 $m_c = m_r + j \cdot m_i$	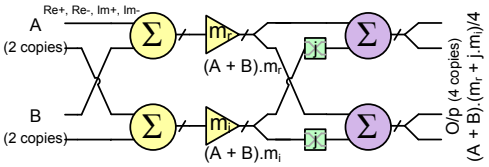
Butterfly	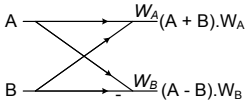	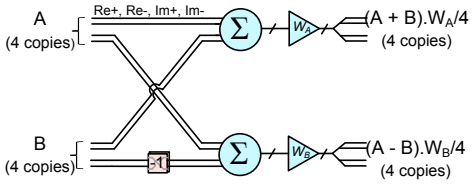

Figure 6.7: Mathematical operations in the CRAFT processor using charge-sharing

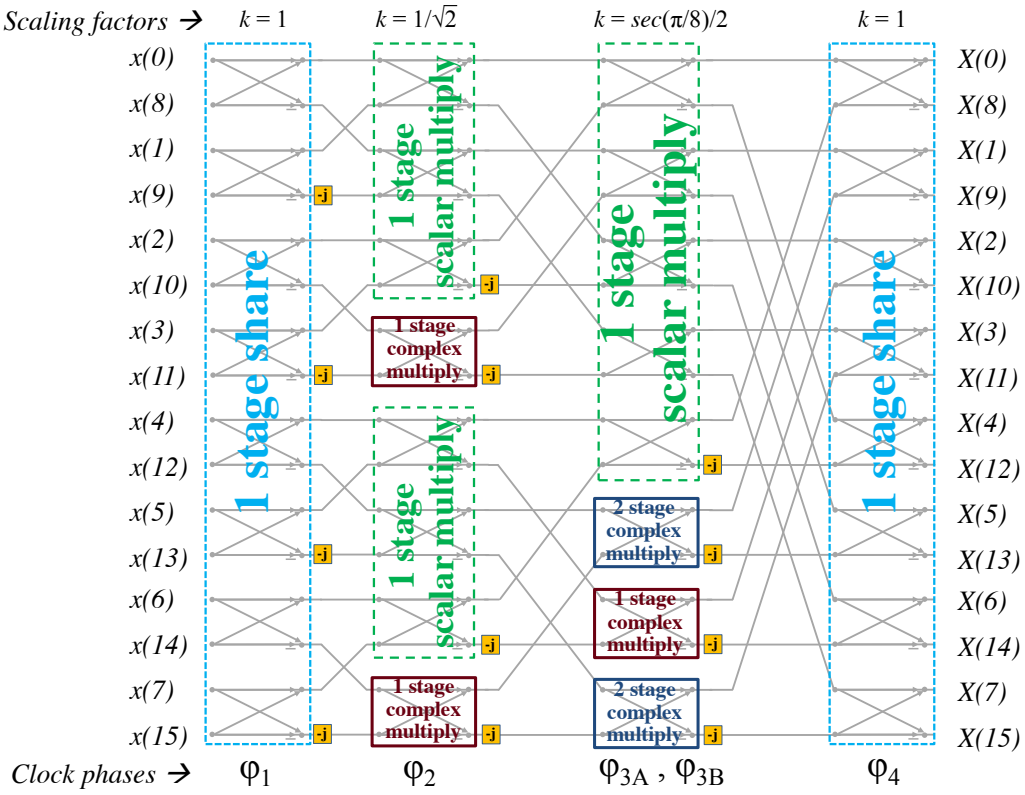

Figure 6.8: The implemented CRAFT algorithm showing the different kinds of butterflies used and the scaling in each stage

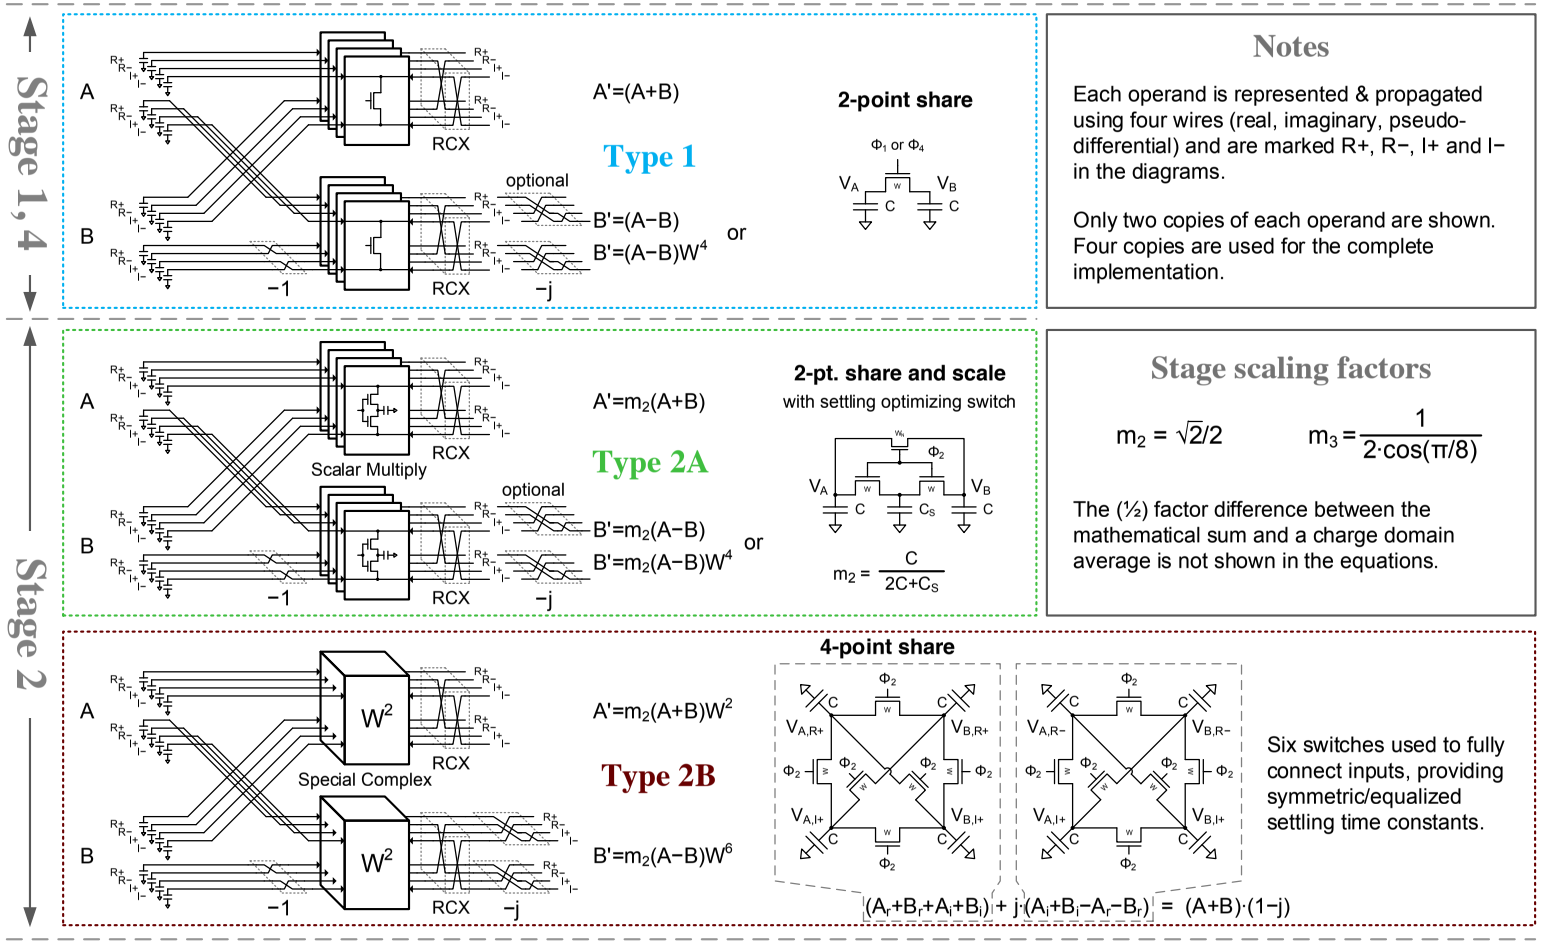

Figure 6.9: The four FFT stages and their component types of butterflies are shown. Note that the capacitors holding the operands for each buttefly (at the left of the diagram) are the sampling capacitors.

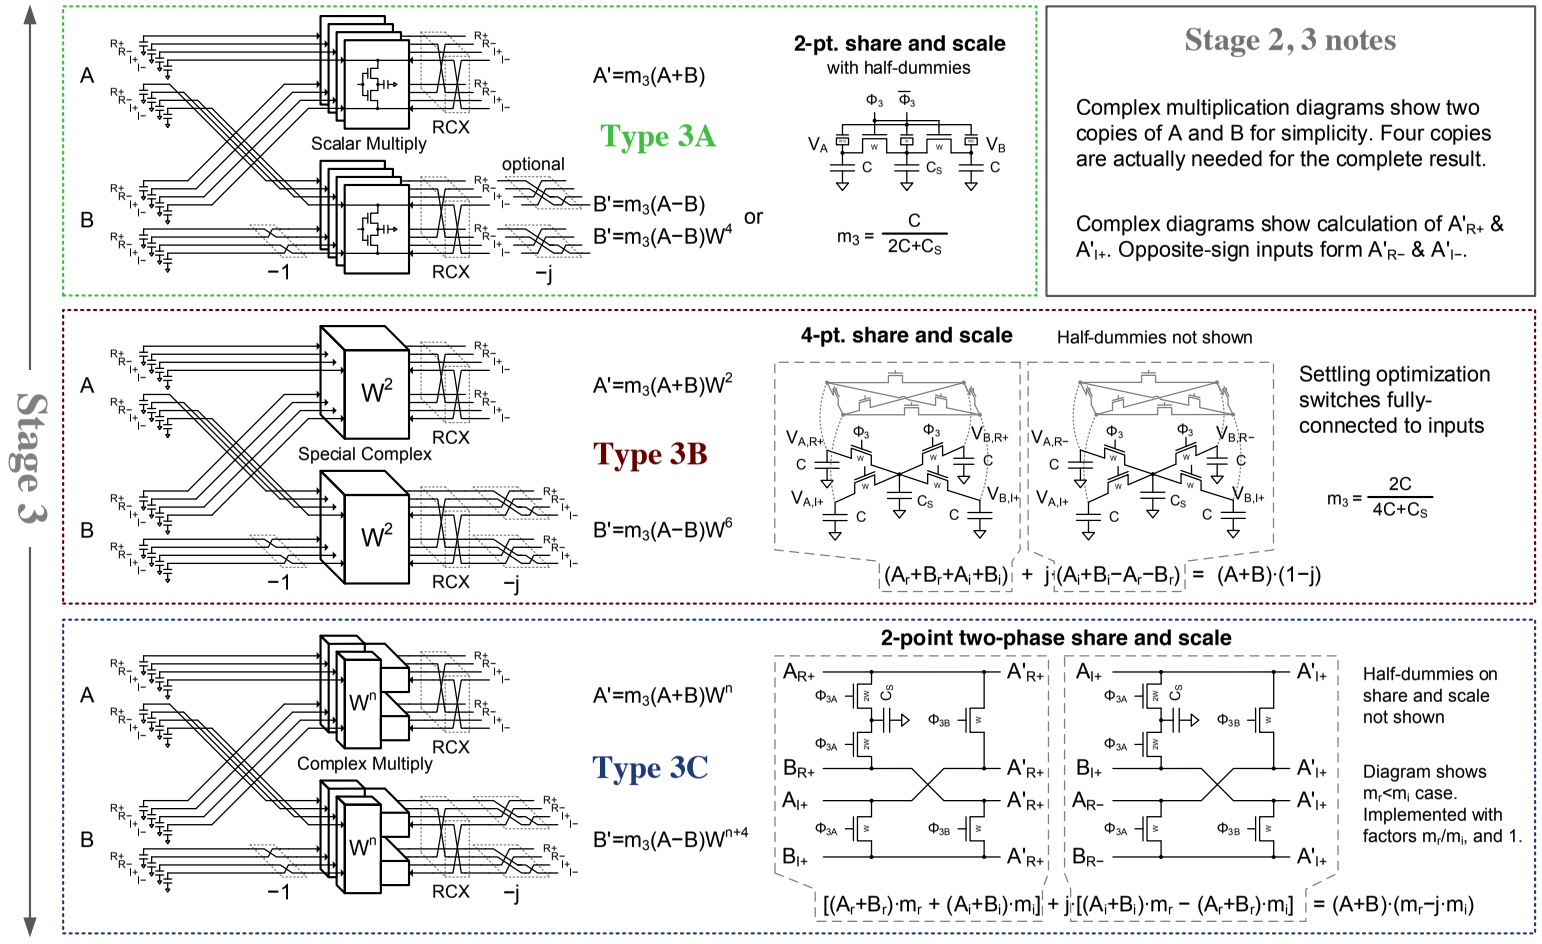

Figure 6.10: The four FFT stages and their component types of butterflies are shown. Note that the capacitors holding the operands for each buttefly (at the left of the diagram) are the sampling capacitors.

6.3 CRAFT Implementation

The CRAFT engine is implemented as shown in Fig. 6.11. The figure closely emulates the layout implementation. A layout screenshot of the on-chip implementation, as well as one of only the core are also shown in Fig. 6.12 and Fig. 6.13 respectively. The design core, shown in Fig. 6.11, includes $4(= \log_2(16))$ stages of FFT butterflies similar to the signal flow graph representation in Fig.6.8. Also, in Fig. 6.11, note the different types of butterflies corresponding to those in Fig.6.8 and in Fig. 6.10. As shown, the core design is supplemented with input voltage samplers for data input, and output latches and multiplexers for data read-out. A block-level detailed description of the figure follows later in this section.

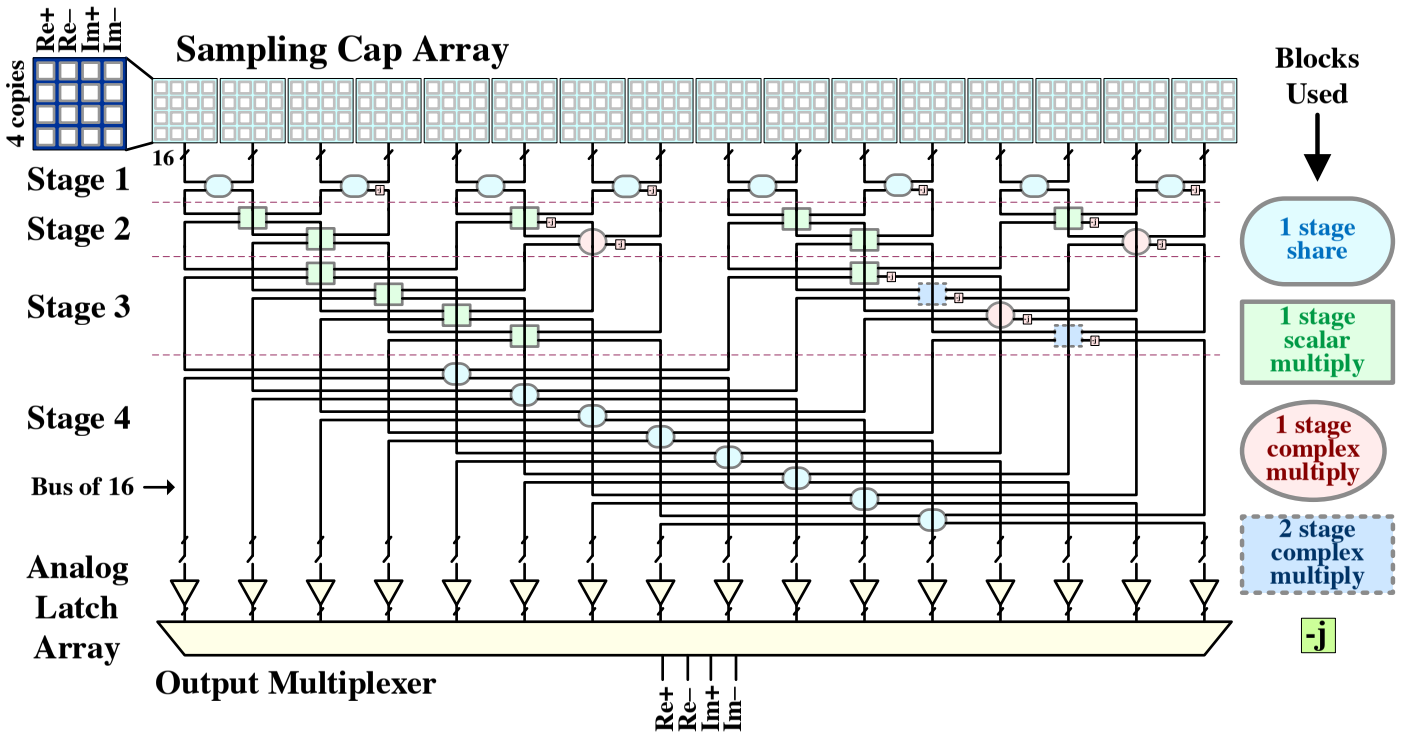

Figure 6.11: A block-level schematic imitation of the CRAFT layout floorplan showing the different blocks used for each stage

All the blocks shown are clocked and controlled by state machines. The timing diagram for the operations is shown in Fig. 6.14. The sampling operation (16 clock phases) and processing operation (4 clock phases) are allowed equal amounts of time in anticipation of an interleave-by-two implementation. The slower processing clock phases allow for more settling time, lowering both the error and the power consumption in the operations. In the current implementation, the reset phase prevents contiguous time window conversions in an interleave by 2 design. However, this is only a trivial limitation of this implementation, and the processing clock durations can be easily reduced for only a marginal increase in power[2] . Similarly, the processing phase can be further reduced to obtain a lower latency design at the expense of higher power. Moreover, dynamic optimization is possible through digital control of the clock phases. For the energy optimized design described in this work, the optimization of the exact duration of each phase is discussed in Section 6.4.

6.3.1 Sampler

The I and Q inputs are probed into the chip and sampled on the 16 sample phases shown in Fig. 6.14. Fig. 6.15(c) shows the pseudo-differential sampler used in CRAFT. Compared to other traditional approaches shown in Fig. 6.15(a) and (b), this topology provides excellent signal feedthrough cancellation for differential signals.

During the sampling hold mode, the following mechanisms cause signal feed-through

[2] Reducing the clock duration would require having larger switches for the same settling accuracy. This would directly translate to a larger switching power dissipation.

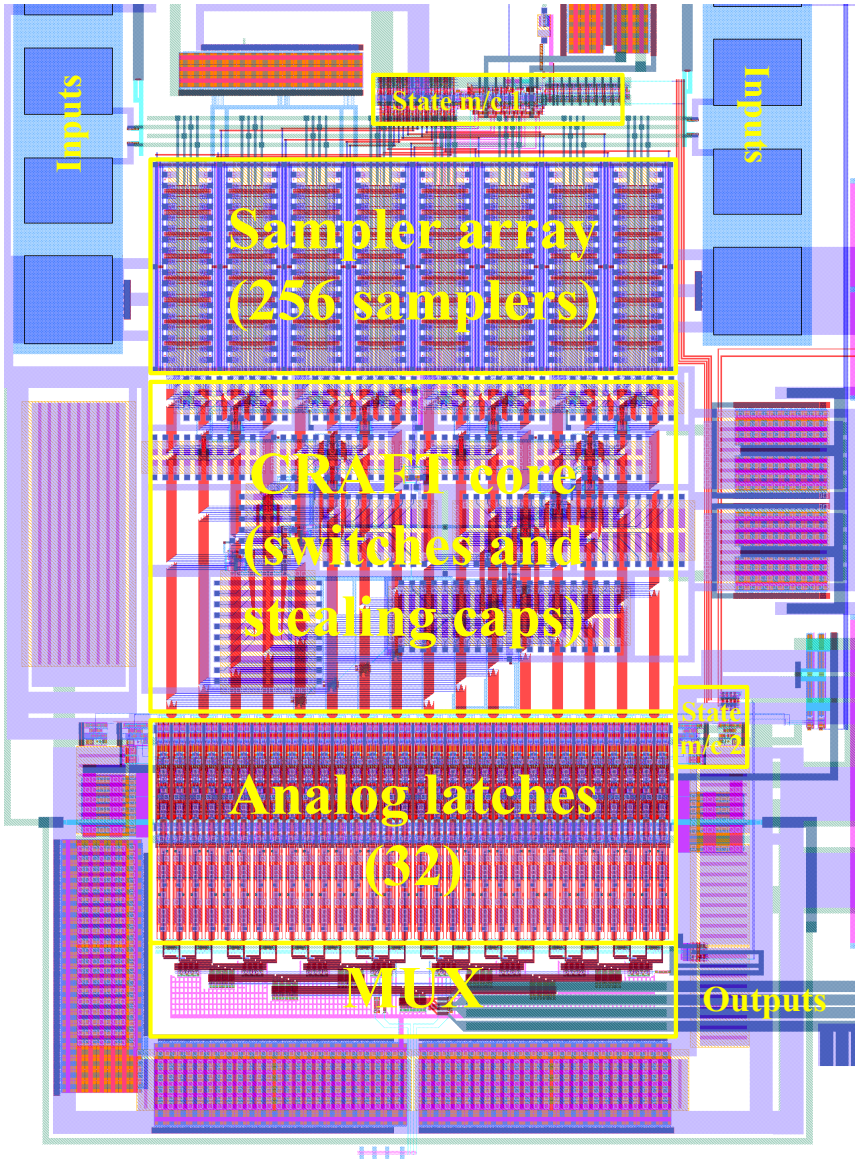

Figure 6.12: A screenshot of the CRAFT implementation layout

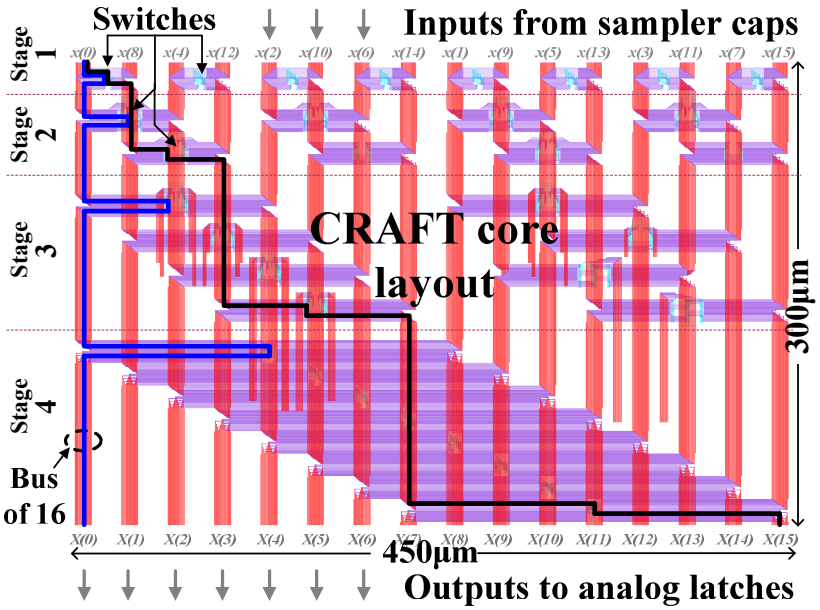

Figure 6.13: A screenshot of the actual CRAFT core layout showing the nominal matching between the parasitic wirings on two sampling capacitors

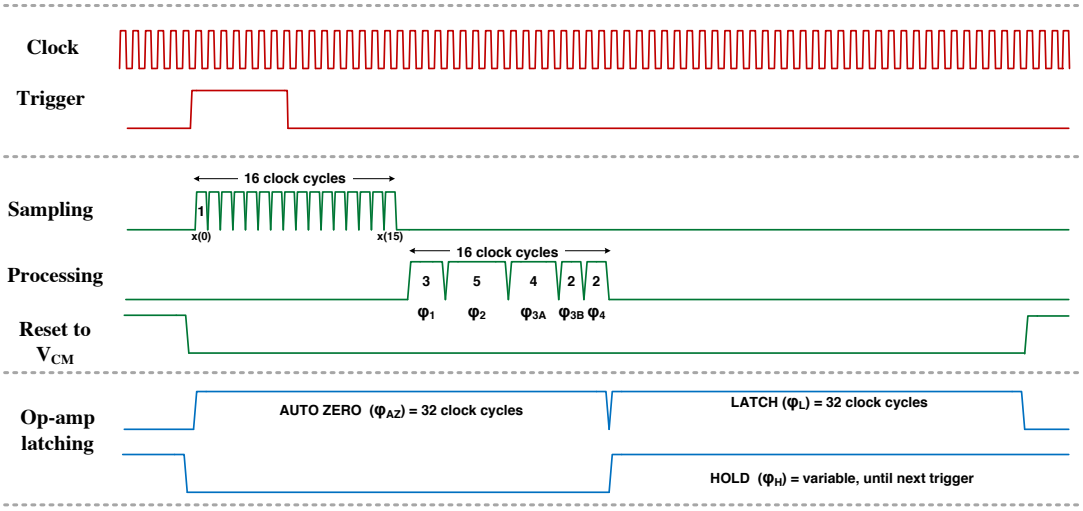

Figure 6.14: Timing diagram showing the clock, trigger, sampling, processing, reset, and output latch clocks

and leakage: (a) capacitive feed-through due to the appreciable C_{ov}/C ratio, and (b) sub-threshold leakage dependent on the input as well as the sampled and held voltages (see Fig. 6.15(a)).

A traditional way to improve the isolation is by grounding the intermediate node during the hold-mode as shown in Fig. 6.15(b) [123]. This reduces the capacitive feed-through as well as the input depence of sub-threshold leakage. However, to maintain the same on-resistance, power consumption increases by more than 4X, the total parasitic loading of the input driver increases, and matching degrades. For example, for the requirements in CRAFT, a ∼50dB isolation improvement compared to (a) can be realized for a 4.1X power increase.

In comparison, the proposed sampler configuration (Fig. 6.15(c)) uses matched cross-coupled feedthrough paths from the differential inputs to cancel the capacitive signal feed-through at the output node. Note that, in this configuration, no additional switching or active power is necessary. Moreover, sub-threshold input leakage effects are differentially-suppressed for small hold voltages. This is useful as in-place processing operations may create small results which should not be disturbed. Simulations suggest that the proposed design in Fig. 6.15(c) provides perfect isolation (>100dB) compared to Fig. 6.15(a) and is only limited by device and layout mismatch.

As shown in Fig. 6.11, each input sample is stored using a set of 16 capacitors (4 copies of pseudo-differential, complex inputs). For a 16 point FFT, 16 such sets (total of 256 capacitors) are used. The 256 sampling switches are designed to run at 5GS/s

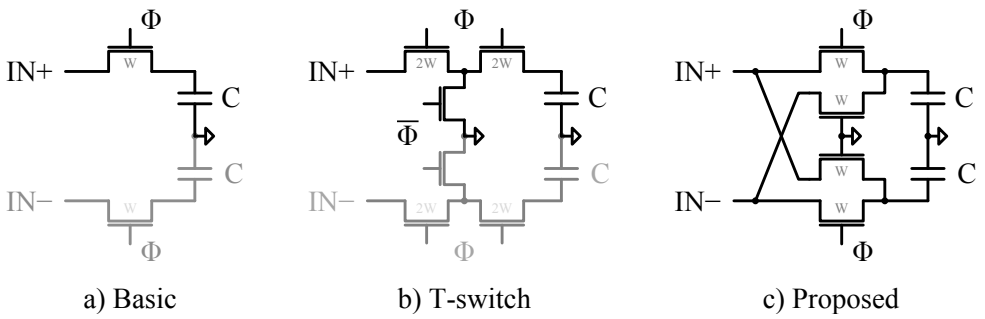

Figure 6.15: The proposed sampler, compared with 2 traditional samplers

with a total differential input swing of 1.2V and a 60dB SFDR in simulation.

The design samples and processes signals using a pseudo-differential representation; two capacitors, each with voltages ranging from 0 to $2V_{cm} = V_{cm} \pm V_p$, give a signal range of $V_{pp,\text{diff}} = 4V_{cm}$. Advantages of this arrangement include: maximized swing while using nMOS-only switches; issues inherent to below-ground swing are avoided; wire swapping can be used for performing negation without bottom-plate switches; capacitors can use a shared bottom-pate for considerable area savings; and signal-independent non-idealities, like clock-feedthrough, appear as common-mode. Although not used in this work, it is also possible to use bottom plate sampling to alleviate charge injection errors in the passive sampler [115].

The choice of the size of each sampling capacitor is a trade-off between the speed and the integrated noise in the circuit. For capacitor size selection, the non-linearity in the circuit is estimated at approximately 60dB and the integrated noise level is placed lower than this so that the circuit is not noise limited. For this, the sampling capacitors are sized at $200fF$ each such that the total output referred $k\text{T}/C$ noise contribution by

the CRAFT engine due to the sampling and processing operations is 64.4dB below the output full-scale signal. Details of the noise contribution of the processing operations is discussed in Section 6.4.

As shown in Fig. 6.11, 16 buses, each 16 wide, are connected to the sampling capacitors, and run through the CRAFT core. These wires are always connected to the sampling capacitors; consequently, their parasitic capacitance forms a part of the sampler and needs to be accurately matched. Additionally, since the operations are performed in-place, the outputs appear on the sampling capacitors at the end of the processing phase.

6.3.2 CRAFT Core

The CRAFT core performs the FFT operation in 4 ($= \log_2(16)$) stages as represented in Fig. 6.8. The processing clock phases were shown in Fig. 6.14. The CRAFT operation, represented in matrix form as shown in equation (6.1), can be further broken down into 4 share and 4 multiply operations in the 4 constituent stages as shown: $\mathbf{F} = \frac{\mathbf{1}}{\mathbf{16}} \cdot \mathbf{S_4 S_3 S_2 S_1 I}_{\text{bitrev}}$, where each of the 4 stages are denoted by $\mathbf{S_i}$, and i is the stage number. The $2^4 = 16$ division factor is because of the inherent scaling by 1/2 in each charge averaging operation. The $\mathbf{I}_{\text{bitrev}}$ and $\mathbf{S_i}$ matrices are expanded in Appendix C. The CRAFT matrix equation (6.1) can be written as $\mathbf{X} = \frac{\mathbf{1}}{\mathbf{16}} \cdot \mathbf{S_4 S_3 S_2 S_1 I}_{\text{bitrev}}\mathbf{x}$ representing the signal flow graph in Fig. 6.5 and Fig. 6.8.

As described in Fig. 6.7, each butterfly is conceptually implemented using switches

and scaling capacitors that perform a set of share and multiply operations on the input samples. Moreover, the *a priori* knowledge of the exact FFT operations are exploited to hard-wire a number of modifications that reduce the number of clock phases required per stage, and limit the signal attenuation inherent to passive computation. The implementation details of these improvised butterflies in each of the 4 stages is shown in Fig. 6.10 and described below in detail.

Stage 1

The first stage comprises only share operations as shown in Fig. 6.5, and by the share matrix, $\mathbf{S_1}$ in Appendix C. All the butterflies in this stage are almost identical, as shown in Fig. 6.8, and Fig. 6.11. The implementation of the butterflies is shown in the top of Fig. 6.10 (Type 1).

Note that the magnitude of the multiplicand in this stage is unity. This is a trivial operation; consequently, the multiplication stage can be completely eliminated as shown in Fig. 6.10. As shown in the Type 1 butterfly, two operands A and B perform a share operation. Each operand is represented using 2 copies of $\Re\pm$ and $\Im\pm$ pseudo-differential voltages stored on 8 capacitors[3] . Corresponding ($\Re\pm$ & $\Im\pm$) charges on these capacitors are shared when the stage 1 processing clocks are enabled. RC error cancellation, as discussed in Section 6.4.2, is performed using wire-swapping at the outputs. For multiplication by '$\pm j$', as required in some butterflies as shown in Fig. 6.8,

[3] Two copies of this butterfly are implemented in the actual circuit to provide 4 copies for the complex butterfly stages utlized in stages 2 and 3.

wire-swapping is performed as shown. This stage consumes a single processing-clock cycle as shown in Fig. 6.14.

Stage 2

The second stage comprises share and multiply operations as shown in Fig. 6.5, and represented by the $\mathbf{S_2}$ matrix in Appendix C. In this stage two types of operations are performed: share and multiply by unity, and share and multiply by $W^2 = \frac{1}{\sqrt{2}} - \frac{1}{\sqrt{2}}i$. We note that a share followed by multiplication by a complex twiddle factor $(m_r - jm_i)$ with equal real and imaginary parts $(m_r = m_i = m)$ can be simplified as shown below:

$$[(A_r + jA_i) + (B_r + jB_i)] \cdot m(1 - j)$$

$$= (A_r + B_r + A_i + B_i) \cdot m + j(-A_r - B_r + A_i + B_i) \cdot m$$

As a result, the 2-stage complex multiplication can be replaced using a single-stage share and multiply operation with 4 operands. Moreover, the share and multiply operation can be replaced by only a share operation to reduce the attenuation, as shown in Fig. 6.10 (Type 2B). For perfect symmetry and equalized settling, the switching configuration as shown in Type 2B is used. The mathematical scaling due to a 4-share operation is $\frac{1}{4}$. Scaling this to the intended attenuation of $\frac{1}{\sqrt{2}}$, we note that a unity gain needs to be replaced by a share (scaling by $\frac{1}{2}$), and a scaling factor of $\frac{1}{\sqrt{2}}$. Consequently, the other butterflies in this stage which originally required a scaling factor of unity should instead be scaled by $\frac{1}{\sqrt{2}}$. This scaling is obtained using a share and multiply operation shown in Fig. 6.10 (Type 2A). The third switch, in Type 2A, is used

for improving the differential settling as discussed in Section 6.4.2. The improvisation using a 4-share operation eliminates one clock phase, thereby providing for more settling time, and minimizes the mathematical attenuation inherent to charge-computations by $\sqrt{2}$.

Stage 3

The third stage comprises share and multiply operations as shown in Fig. 6.5 and represented by matrix $\mathbf{S_3}$ in Appendix C. As seen in the figure, the twiddle factors in this stage include $W^1 = j \cdot W^5 = \cos(\frac{\pi}{8}) - j \cdot \sin(\frac{\pi}{8})$, and $W^3 = j \cdot W^7 = \sin(\frac{\pi}{8}) - j \cdot \cos(\frac{\pi}{8})$. These butterflies require the complete complex multiplication stage shown in Fig. 6.7 and consume 2 clock cycles: ϕ_{3A} and ϕ_{3B} in Fig. 6.14. Also, for the 4 twiddle factors, only two types of stealing capacitors, in conjunction with 'j' wire-swaps, are used in order to reduce mismatch. The implementation of these butterflies is shown in Fig. 6.10 (Type 3C). As shown, the complex multiplication is implemented in 2 stages. The entire operation is scaled by $1/m_i$ in order to reduce the total attenuation by $sec(\pi/8)$. Consequently, only one scaling capacitor is necessary for this butterfly type as shown.

The share and multiply by unity is performed using a share operation as before. Also, similar to stage 2, these butterflies incorporate a constant scaling factor using a stealing capacitor to equalize the total stage scaling to the scaling incurred in the complex multiplication butterflies. These butterflies, shown in Fig. 6.10 (Type 3A), utilize both ϕ_{3A} and ϕ_{3B} for their operations. Half-dummy switches are used as shown to cancel

clock feed-through and reduce charge injection (details in Sections 6.4.4 and 6.4.5).

Multiplication by W^2 is performed using a technique similar to the 4-point share switch used in stage 2. For stage 3, however, further attenuation is required to adjust for the inherent scaling in the Type 3C butterflies in this stage. Therefore a stealing capacitor is added to the Type 3B butterflies. The optimized switching configuration is determined and shown in Fig. 6.10 (Type 3B). Extra switches added for improving the differential settling as shown. Half-dummy switches added for clock feed-through cancelation are omitted to avoid clutter in the diagram.

Stage 4

The fourth stage comprises only share operations as shown in Fig. 6.5 and in the matrix below. These butterflies are identical to the butterflies in stage 1 and are not redrawn in Fig. 6.10.

Using the improvised share and multiply techniques, only 5 clock phases are used for FFT computation. Moreover, the total attenuation in the CRAFT core is limited to 0.38 ($= 1 \times \frac{1}{\sqrt{2}} \times \frac{1}{2} \cdot sec(\frac{\pi}{8}) \times 1$) as shown in Fig. 6.8. This attenuation is 6.5X (16.3dB) superior compared to an unmodified implementation[4] .

Note that the butterfly implementation depicted in Fig. 6.11 imitates the signal flow shown in the last row of Fig. 6.7; as a result, after each operation, half the wires return

[4] An unmodified implementation is one which uses generic complex scalar butterflies (Fig. 6.7, row 6) everywhere.

to their bus while the rest continue on the other bus. To equalize the sampler wiring parasitics, the switches are always placed midway between two operand buses. Two example wires, one always returning to its own bus while the other always shifting on to the other operand bus, are highlighted in the layout screenshot in Fig. 6.11. As seen, the two wire lengths (and their associated parasitics) are nominally matched.

6.3.3 Output Latch

On the far end of the core, the wires connect through switches to operational transconductance amplifier (OTA) based analog latches that save the outputs prior to being read out (Fig. 6.11). To match the CRAFT performance, a two-stage, folded cascode, differential OTA with 70dB gain and 900MHz UGB is designed. The OTA is used in a differential switched capacitor analog latch, shown in Fig. 6.16. The corresponding clock phases used are shown in Fig. 6.14. The latch performs offset cancelation during the CRAFT sampling and processing phases (ϕ_{AZ}), and latches the CRAFT output with a 10τ settling accuracy during the next 32 clock phases (ϕ_L). The output is then held (ϕ_H) for output read until the external measurement system reads the outputs. Thirty-two (16×2 for real and imaginary) of these latches capture the FFT output. Switched-capacitor based common-mode feedback is performed during the ϕ_{AZ} phase.

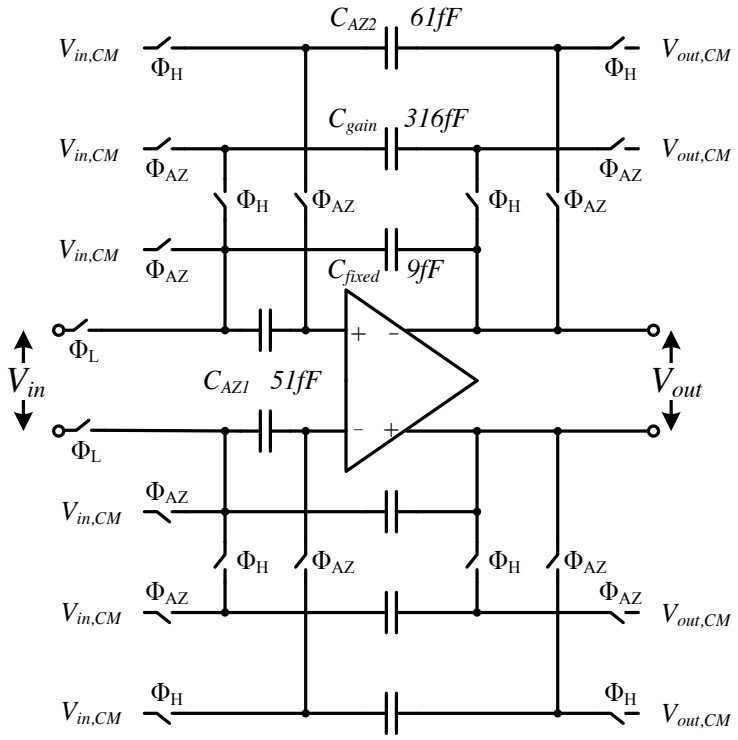

Figure 6.16: One of the 32 output latches used to latch the CRAFT results

6.3.4 State Machines

As shown in Fig. 6.17, the entire sampling, FFT computation, and testing interface require multiple clocks to be generated and interfaced with the testing equipment. The design uses an input clock to generate all internal signals. State machine 1 (SM1) is externally triggered and generates 16 sampling clock phases as shown in Fig. 6.14. It then generates the processing clocks that operate the CRAFT core switches. Another state machine (SM2) uses handshaking with SM1 and with an external FPGA (NI-7811R) to determine if the outputs and the FPGA are ready and subsequently generates the analog latch clocks. The latched outputs are observed sequentially using an FPGA controlled output MUX.

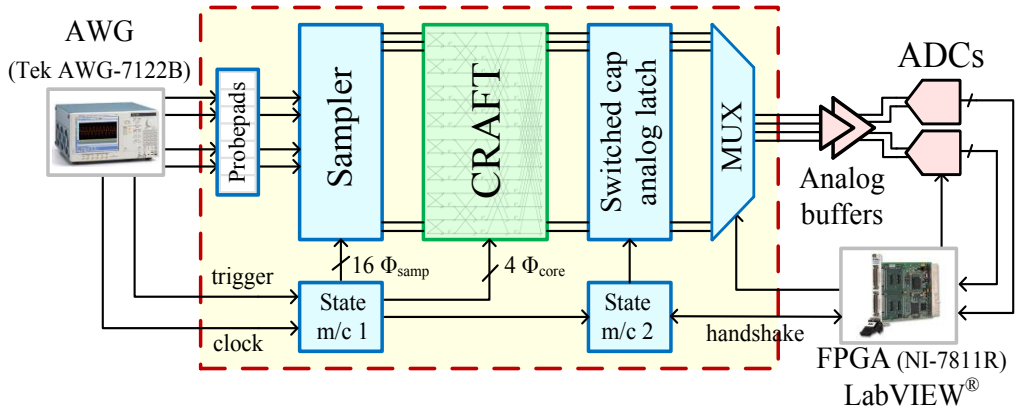

Figure 6.17: Test setup for the CRAFT processor (on and off chip)

6.4 Circuit Non-idealities and Optimization

As expected, the design relies heavily on digital circuits. Also the design's regularity and complexity makes it comparable to digital designs. However, despite these similarities, the CRAFT engine is an analog circuit, and remains vulnerable to circuit non-idealities including noise, matching, and non-linearity. Consequently, accurate modeling of non-idealities is critical. Moreover, switched capacitor circuit noise, charge-injection, and charge-accumulation are not adequately modeled in circuit-level simulators. For this design, CRAFT-specific models of dominant non-idealities (noise, settling, clock-feedthrough, charge-injection, etc.) are developed in MATLAB®. Each non-ideality is modeled as an independent error source. These error sources can be enabled and disabled independently in simulation for each stage of the CRAFT implementation to enable isolation of the impact of each error source. Also, *a priori* knowledge of the operations is utilized to devise novel circuit techniques such as correlated noise reduction, differential settling error cancellation, and high-linearity switches. Performance is optimized using system simulations that model both the non-idealities and the new circuit techniques. This enables the high dynamic range even at 5GS/s. The different sources of non-ideality, their modeling, and mitigation techniques used for CRAFT are outlined below.

The different operations in CRAFT can be distinguished into the initial sampling phase, and the subsequent processing phases. In this section, the effect of noise in the sampler as well as the core is discussed. This is followed by a discussion on other

non-linearities in the CRAFT core processing[5] switches, namely: operand-dependent processing-switch resistance, clock-feedthrough, charge injection and absorption, and incomplete settling, and outline some techniques for their mitigation.

6.4.1 Noise

Sampling Noise

Motivation During a voltage sampling operation, noise presents as a final-value disturbance with power $k\mathrm{T}/C$. Noise during the reset operation does not matter[6] . In this design, 4×2 capacitors, $200f\mathrm{F}$ each, are used to sample 4 copies of each input pseudo-differentially. The full-scale input is $V_{pp,diff} = 1.2V$. This sampler yields a sampled-noise voltage of $\sqrt{\overline{V_{n,rms}^2}} = V_{n,rms} = 144\mu V$ on each of the 4 single-ended copies (–63.4dBFS). As the noise is uncorrelated, averaging these copies at the output and forming a differential output gives a total noise of $V_{n,rms} = 144 \cdot \frac{1}{\sqrt{4}} \cdot \sqrt{2} = 102\mu V$ (–72.4dBc) for a full-scale single-tone input.

Also, for the sampling bandwidth used, the noise of the test equipment's 50Ω source resistance is approximately 6dB below the sampler noise per single-ended copy. However, it is correlated among the output copies and does not benefit from copy averaging, effectively doubling the total sampled noise power ($N_{samp} = 2 \times N_{input} = 2 \times N_{switch}$) during testing.

[5] The sampling process and its related non-idealities have been well-studied in literature; we focus on the processing non-idealities in this paper.

[6] This is not true for a current-domain sampler: noise during the reset phase disturbs the initial value, and this noise combines with the noise added when the sampling switch opens at the end of the sampling period.

Modeling Sampling noise effects are included into our MATLAB®system simulation model. After sampling or resetting, a Gaussian random variable with $\sigma = V_{n,rms} = \sqrt{k\mathrm{T}/C}$ is added to the each capacitor's final value.

Mitigation The sampling capacitor size sets both the sampler's noise as well as the baseline for the processing noise because the same capacitors are used for computations. The capacitor size is selected to be 200fF/capacitor based on the simulated total output-referred noise.

Processing Noise

Motivation Similar to the sampling operation, noise from the CRAFT core switches corrupt the computations. At the end of every share operation, $k\mathrm{T}/C$ noise power is added to each output copy. Interestingly, this instantaneous sampled noise on the two capacitors arising from a single switch is completely correlated (equal magnitude, opposite signs). Similarly, the noise sampled on the capacitors during a share and multiply operation arise from the same switches, and are therefore partially correlated. Also, any noise of the input operands (for e.g., from the previous stage of operations) get averaged out during a share operation.

Modeling For modeling and simulating the processing noise, Gaussian random variables distributed in magnitude and sign on the output copies in exactly the manner the noise-transfer functions dictate for the particular switch configuration are used. This

generates the proper noise correlation on the output copies, that, averaged over multiple simulations, provide the expected output referred noise power. For example, for a 2-point share operation (row 1 of Fig. 6.7), the integrated rms noise voltages at the outputs are given by $V_{n,rms} = \sqrt{0.5 \times kT/C}$, and are differentially correlated. This configuration is used in stages 1 and 4 of CRAFT, shown in the Type 1 configuration in Fig. 6.10.

Mitigation Noise of the later stages of the CRAFT engine is reduced due to copy averaging. Four copies of each output are available, and are shared to reduce the total noise power by 4. Moreover, correlation, ρ, between these copies that appear on the same final stage output can be exploited for noise reduction if the noise-correlated pairs appear at the same output.

Total Output Noise The noise contribution of each stage is computed analytically and tabulated in Table 6.1. The attenuation (A_v) reduces the output referred noise by A_v^2. The total single-ended, copy-averaged output referred noise, including the noise of the 50Ω source, is computed as shown in the last column yielding a total output referred noise of $0.23 \cdot \frac{kT}{C}$ (-64.4dBFS per differential $\Re$, $\Im$ output for a noise EVM of -62.5dBFS[7]).

[7] $\sqrt{\Re^2 + \Im^2}$ is a chi distribution with 2 degrees of freedom. For $X_i \sim \mathcal{N}(0,1)$ it has mean $\sqrt{\pi/2}$ (=1.96dB).

Table 6.1: Summary of noise contribution in CRAFT

Stage	Noise sources	$\overline{V_n^2}$ per copy	A_v	ρ_{out}	$P_{n,out} = \left(\frac{1}{4}\right) \cdot \overline{V_n^2} \cdot A_v^2 \cdot (1+\rho)$
	Sampler (4-copy)	$1.000 \cdot \frac{kT}{C}$	0.38	0	$0.0366 \cdot \frac{kT}{C}$
1	2pt. share	$0.500 \cdot \frac{kT}{C}$	0.38	0	$0.0183 \cdot \frac{kT}{C}$
2	2pt. sh./scale (m=$1/\sqrt{2}$)	$0.646 \cdot \frac{kT}{C}$	0.54	0	$0.0473 \cdot \frac{kT}{C}$
3	2pt. sh./scale (m=0.541)	$0.729 \cdot \frac{kT}{C}$	1	-0.51	$0.0912 \cdot \frac{kT}{C}$
4	2pt. share	$0.500 \cdot \frac{kT}{C}$	1	-1	0
	Total output referred CRAFT noise				$0.193 \cdot \frac{kT}{C}$
	50Ω source, output referred				$\approx 0.037 \cdot \frac{kT}{C}$
	Total output referred noise				$0.23 \cdot \frac{kT}{C}$

6.4.2 Incomplete Settling

Motivation For an RF discrete time signal processor, very high sampling and processing speeds are mandated. The speed of operations in CRAFT directly trade-off with the settling accuracy of the operations. The switch size can be increased to allow better settling; however, this not only increases the power consumption, but also causes larger charge-injection and clock-feedthrough errors. To optimize the computation accuracy, the time-domain settling characteristics give useful insight.

For the modeling and mitigation sections that follow, two aspects are considered. First, models are developed for a single-ended settling scenario. Later, this model is extended to the pseudo-differential operation utilized in CRAFT. Mitigation techniques are developed based on these models.

Note: indentation takes up space... formatting suggestions?

Modeling

(i) Single-ended settling: Consider 2 operand capacitors, each with a capacitance, C, with voltages of $V_a^+(t)$ and $V_b^+(t)$. A 2-point share operation, implemented as in row 1 in Fig. 6.7 and Fig. 6.10 (Type 1), has a settling response as shown in equation (6.2), and is plotted in Fig. 6.18 (see $V_a^+(t)$ and $V_b^+(t)$). The input capacitors (C), holding initial values of v_{a0} and v_{b0}, are connected by a switch of resistance R_{sw} at $t = 0$. This yields a settling error time-constant: $\tau_d = R_{sw}C/2$.

$$V_a(t) = \tfrac{1}{2}\left(v_{a0} + v_{b0}\right) + \tfrac{1}{2}\left(v_{a0} - v_{b0}\right) e^{-t/\tau_d} \tag{6.2}$$

$$V_b(t) = \underbrace{\tfrac{1}{2}\left(v_{a0} + v_{b0}\right)}_{\text{ideal result}} - \underbrace{\tfrac{1}{2}\left(v_{a0} - v_{b0}\right) e^{-t/\tau_d}}_{\text{settling error}}$$

A 2-point share and multiply operation, as shown in Fig. 6.7, has the settling response of equation (6.3). The input capacitors (C), with initial values v_{a0} and v_{b0}, are connected to the stealing capacitor (C_s, with no initial value: $v_{s0} = 0$) by switches of resistance R_{sw}. The input-to-stealing-capacitor ratio defines the scaling factor of the operation: $m = \frac{2}{(2+C_s/C)}$.

$$V_a(t) = \overbrace{\tfrac{1}{2}\left(v_{a0} + v_{b0}\right) m \left[1 + \left(\tfrac{1}{m} - 1\right) e^{-t/\tau_s}\right]}^{\text{scaled-sum term}} + \overbrace{\tfrac{1}{2}\left(v_{a0} - v_{b0}\right) e^{-t/\tau_d}}^{\text{difference term}} \tag{6.3}$$

$$V_b(t) = \underbrace{\tfrac{1}{2}\left(v_{a0} + v_{b0}\right) m}_{\text{ideal result}} \underbrace{\left[1 + \left(\tfrac{1}{m} - 1\right) e^{-t/\tau_s}\right]}_{\text{scaling error factor}} - \underbrace{\tfrac{1}{2}\left(v_{a0} - v_{b0}\right) e^{-t/\tau_d}}_{\text{difference settling error}}$$

The difference and sum-settling time-constants, τ_d and τ_s respectively, are

$$\tau_d = R_{sw}C \qquad\qquad \tau_s = R_{sw} \cdot (C \parallel C_s/2)$$

$$= R_{sw}C \cdot (1 - m)$$

(ii) Pseudo-differential settling: Pseudo-differential operands are described by the relationships below, where V^+ and V^- are the voltages on separate capacitors for the positive and negative components of the operand.

$$V_A(t) = V_a^+(t) - V_a^-(t) \qquad V_B(t) = V_b^+(t) - V_b^-(t) \tag{6.4}$$

Consequently, charge-domain operations are performed separately on these components. The time-domain settling response of these differential operands is shown in Fig. 6.18. For example, $V_A(t)$ is composed of two share and multiply settling operations as shown below. Using $m' = m\left[1 + \left(\frac{1}{m} - 1\right) e^{-t/\tau_s}\right]$ and $\epsilon = \frac{1}{2}\left(v_{a0} - v_{b0}\right) e^{-t/\tau_d}$, we can write,

$$\begin{aligned} V_A(t) &= \overbrace{\left[\tfrac{1}{2}\left(v_{a0}^+ + v_{b0}^+\right) m'^+ + \epsilon^+\right]}^{V_a^+(t)} - \overbrace{\left[\tfrac{1}{2}\left(v_{a0}^- + v_{b0}^-\right) m'^- + \epsilon^-\right]}^{V_a^-(t)} \\ &= \left(v_{a0}^+ + v_{b0}^+\right) \underbrace{\tfrac{1}{2}(m'^+ + m'^-)}_{m'_d} + \underbrace{(\epsilon^+ - \epsilon^-)}_{\epsilon_d} \end{aligned} \tag{6.5}$$

The differential nature ($v_{a0}^+ = -v_{a0}^-$ and $v_{b0}^+ = -v_{b0}^-$) allows this share and multiply expression to be written in a form similar to the single-ended expressions $V_a^+(t)$ and $V_a^-(t)$ within equation (6.5) ($m' = m'_d = 1$ for share operations).

$$V_A(t) = \left(v_{a0}^+ + v_{b0}^+\right) m'_d + \epsilon_d \qquad V_B(t) = \left(v_{a0}^+ + v_{b0}^+\right) m'_d - \epsilon_d \tag{6.6}$$

For the differential operands, scaling error, $m'_d = m\left[1 + \left(\frac{1}{m} - 1\right) \frac{1}{2}(e^{-t/\tau_s^+} + e^{-t/\tau_s^-})\right]$, and inter-copy error, $\epsilon_d = \left(v_{a0}^+ - v_{b0}^+\right) \frac{1}{2}(e^{-t/\tau_d^+} + e^{-t/\tau_d^-})$, now include the settling effects of the positive and negative components of the pseudo-differential representation. For the case of small operand swings, $\tau_d^+ \approx \tau_d^-$ and the differential error's

relationship to the two operations performed (pos. and neg.) is clearest: $m'_d \approx m'$ and $\epsilon_d \approx 2\epsilon$.

Mitigation

(i) Single-ended settling: The settling accuracy can be improved through the addition of a third switch ($R_{sw,\mathrm{N}}$) for the share and multiply operation, as shown in Fig. 6.10 (Type 2A). This improves the differential settling error on the 2 operands by providing an alternate settling path. Note, from Fig. 6.10 (Type 2A), the widths of the main sharing switches (having a resistance, R_{sw}) effectively occur in series for the differential settling equation. As a result, increasing the third switch width (having a resistance, $R_{sw,3}$) improves the differential settling twice as power-efficiently compared to increasing the main switch width (see equation (6.7) below). For our switch-sizing choice, the settling-power product ($(e^{\tau_d})^2 \times P_{sw}$) improves $\sim$ 3dB.

$$\tau'_d = (2R_{sw} \parallel R_{sw,3}) \cdot (C/2) = R_{sw}C \cdot \left(\frac{1}{1+2R_{sw}/R_{sw,3}}\right) \tag{6.7}$$

(ii) Pseudo-differential settling: Considering these settling expressions and the situation under small operand swings, a method for RC settling error cancellation (RCX) is developed. Recognizing that V_a^- and V_b^- settle to the same value but have ϵ error of opposite sign, they can be interchanged to form the new differential

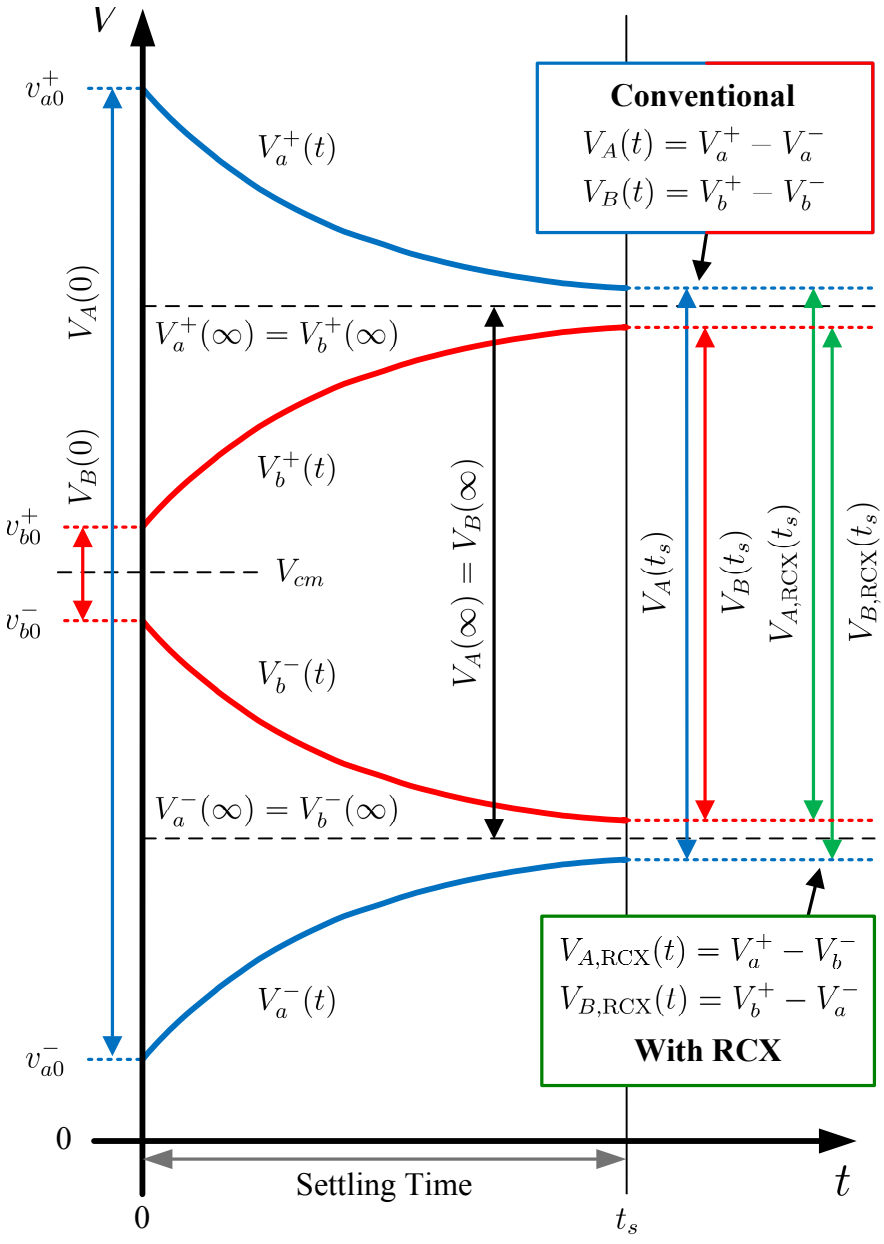

Figure 6.18: Reduction in differential settling error using RCX

operands below (compare with equation 6.4).

$$V_{A,\mathrm{RCX}}(t) = V_a^+(t) - V_b^-(t) \qquad V_{B,\mathrm{RCX}}(t) = V_b^+(t) - V_a^-(t)$$

This is implemented using only wire swapping as shown in Fig. 6.18. As shown, swapping the output wires between the copies greatly reduces differential settling error. While the pseudo-differential "inputs" $(+,-)$ to the operation are (V_a^+, V_a^-) and (V_b^+, V_b^-), the capacitors holding the "output" results (2 copies) are (V_a^+, V_b^-) and (V_b^+, V_a^-). This rewiring causes only a slight change to equation (6.6) for the two output copies as seen below.

$$V_{A,\mathrm{RCX}}(t) = \left(v_{a0}^+ + v_{b0}^+\right) m_d' + \epsilon_{d,\mathrm{RCX}} \qquad V_{B,\mathrm{RCX}}(t) = \left(v_{a0}^+ + v_{b0}^+\right) m_d' - \epsilon_{d,\mathrm{RCX}} \tag{6.8}$$

This changes the $(e^{-t/\tau_d^+} + e^{-t/\tau_d^-})$ term in ϵ_d into $(e^{-t/\tau_d^+} - e^{-t/\tau_d^-})$, allowing the differential-settling error to be canceled when $\tau_d^+ \approx \tau_d^-$. Also, it can be shown that $\epsilon_{d,\mathrm{RCX}} < \epsilon_d$:

$$\epsilon_d = (\epsilon^+ - \epsilon^-) \qquad \epsilon_{d,\mathrm{RCX}} = (\epsilon^+ + \epsilon^-) = (\epsilon^+ - \epsilon^-)\left[\frac{\epsilon^+ + \epsilon^-}{\epsilon^+ - \epsilon^-}\right]$$
$$= \epsilon_d \cdot \tanh\left(t\left(\frac{\tau_d^+ - \tau_d^-}{2\tau_d^+\tau_d^-}\right)\right)$$

This means that RCX always yields a net improvement[8] . RCX is used in all the CRAFT butterflies as seen in Fig. 6.10. Simulations show a net improvement of about 10dB in the FFT settling error from RCX.

[8] It is important to understand that the inter-copy error is not actually being cancelled, it is just partially translated into a common-mode component. In case a stage with common-mode rejection follows, this settling error is rejected.

6.4.3 Operand-dependent Processing-switch Resistance

Motivation The on-resistance of the nMOS only processing-switches are dependent on the operand value. This variation in R_{sw} causes a range of time-constants to be realized. Furthermore, the time-variation of node voltages cause R_{sw} to vary during the computations. For example, a 2-point share has a switch with resistance, $R_{sw} = R_{\text{triode}}(V_G, V_D, V_S) = R_{\text{triode}}(V_{dd}, V_a(t), V_b(t))$. Consequently, the system of differential equations no longer have solutions as in equations (6.2) and (6.3).

Modeling Short of full circuit-level simulation, settling behavior must at least be modeled using iterative numerical evaluation of the differential equations with a small time-step between re-computation of switch resistance values. Simulated node voltages can then be used to calculate effective settling time-constants $\tau_{s,\text{eff}}$ and $\tau_{d,\text{eff}}$ (of equations (6.2) and (6.3)) at $t = t_s$. These values, heretofore referenced as simply τ_s and τ_d, lump the effects of voltage and time variation, as well as resistance differences in multi-switch configurations rather than referring to the instantaneous value of $R_{sw}(t) \cdot C$. This reinterpretation allows the analysis and notation of equation (6.6) and equation (6.8) to be retained.

6.4.4 Clock Feed-through

Motivation Clock signals used to toggle the CRAFT core switches feeds through the nMOS-switch overlap capacitances to the operand capacitors, C or C_s, causing a

charge-domain error on the operand values. Each charge-domain operation consists of one rising and one falling edge applied to the switches. In the case of complete settling, share operations (equal operand capacitors, C) generate equal and opposite disturbances ($+\Delta$ and $-\Delta$) on the operands as the switch closes and opens causing no overall error. However, for share and multiply operations, the operand capacitors are unequal (2 Cs and 1 C_s). As a result, the operands suffer unequal disturbances (say, Δ_1 and Δ_2) during a switching event. The rising-edge clock-feedthrough charge gets shared during the operation phase. However, when the switches are opened, the $-\Delta_1$ and $-\Delta_2$ falling-edge clock-feedthrough charge is not re-distributed, causing errors of $\pm\frac{(\Delta_1-\Delta_2)}{2}$.

Modeling A MATLAB®model of the CRAFT clock-feedthrough is utilized as follows. A rising-edge feedthrough voltage step is applied. The perturbed voltages then perform sharing. Another falling-edge feedthrough voltage step is then applied. The magnitude of each voltage step is a function of switch sizes, node capacitances, and the supply voltage.

Mitigation Clock feedthrough is minimized by adding half-dummy switches. This comes at the cost of doubling switching power and requiring differential clocks. Isolating the error sources, simulating their propagation in MATLAB®, and analyzing their effect on the total error is therefore critical for optimization.

6.4.5 Charge Injection and Absorption

Motivation During switched capacitor operations, turning on MOS switches causes charge absorption perturbing the source and drain nodes, while turning off the MOS switches causes channel charge injection to these nodes.

The charge injection error depends on the impedance on each side, the clock transition time, the channel transient time constant, and the node time constants, as discussed in [124]. For charge absorption, apart from the factors mentioned, the operand node voltages play an important part as shown in Fig. 6.19. Note, however, that as a result of the sharing operation that follows turning on the switch, the charge absorption error is distributed among the sharing capacitors causing a constant error. On the contrary, the charge injection error, occurring at the end of the share operation is not shared, and causes impedance and fall-time dependent errors. In CRAFT, however, due to the high speed of operations, incomplete settling causes some of the charge absorption error to be retained. As a result, this error needs to be modeled for high dynamic range performance.

Modeling For modeling the charge injection and absorption, the clock edge is assumed to be much slower than the minority carrier time constants. For charge injection, an equal split is assumed at the end of the sharing operation (equal voltages and equal impedances on both nodes) as shown in Fig. 6.19. For the share and multiply operation, the model developed in [124] is utilized. For charge absorption, the effect of

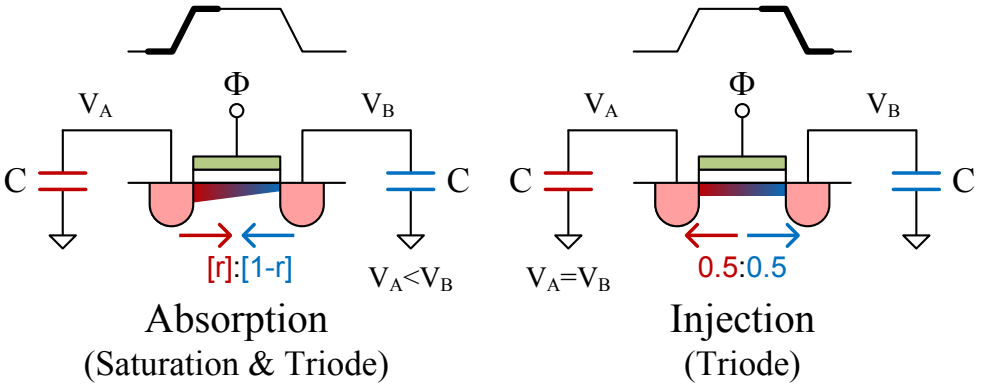

Figure 6.19: Charge injection and absorption for a sharing operation

the instantaneous node voltages is included in the model. This error is then operated upon, and the remnant of this error due to incomplete settling is included in the output values.

Mitigation Like clock feedthrough, half-dummies are an effective way of reducing the effects of channel charge[9] . However, apart from doubling the power consumption, their primary limitation is that they work best only when the charge-split is even. This occurs only for share operations with small input operand swings. In other cases, the dummy weighting could be adjusted to compensate for unequal charge split due to unequal impedances (stealing capacitor). However, this would reduce the effectiveness of the clock-feedthrough cancellation.

To minimize the effect of the residual charge absorption disturbance, improving t_s/τ without increasing the switch channel charge Q_{ch} is desired. Here t_s is the settling time,

[9] Note, that in passive charge-domain operations, switch channel charge must be handled without the advantages available to traditional switched-capacitor circuits such as virtual grounds or low impedance nodes.

and τ is the settling time constant. Without increasing τ, this can be accomplished by: increasing t_s, decreasing switch length, or choosing a V_{dd} for a minimum C_{ch}/R_{sw}, as described in Section 6.5.

6.5 Design Methodology and Optimization

For the CRAFT design, power, speed, and dynamic range directly trade-off with each other. Moreover, the analysis of non-idealities, discussed in Section 6.4, represents a complex design space with different trade-offs associated with the mitigation techniques. This section outlines a design and optimization methodology used in the CRAFT design for superior performance. For this specific implementation, the following specifications/constraints were assumed:

1. A 5GS/s input rate with an interleave by 2 CRAFT engine for processing contiguous windows: This provides a total processing time of $16 \times \frac{1}{5\text{GHz}} = 3.2\text{ns}$.

2. A 60dB (10bit) dynamic range is chosen as a target specification.

Using these specifications, the CRAFT engine is optimized for power. The design methodology is divided into an architecture choice based on the constraints listed, followed by energy optimization.

6.5.1 Architecture Choice

For the architecture choice the processing time is assumed to be shared equally among the 5 clock phases. This assumption is revisited during energy optimization later. Also, a nominal V_{dd} for all the stages is assumed as a first guess. Based on these assumptions, the following choices are made:

1. *Input swing:* The maximum input swing, $V_{sig,\max}$ for the sampler is chosen to achieve better than -60dBFS non-linearity floor for the sampler running at 5GS/s. This represents the peak-to-peak input swing. As a result, the common mode voltage, V_{cm} is set at $\frac{1}{2} \cdot V_{sig,\max}$.

2. *Capacitor size:* The sampling capacitor size is selected such that the noise floor from the sampling operation is at least 10dB lower than required for the target SNDR of 60dB. This provides a sampling capacitor size of 200fF.

3. *Attenuation:* The attenuation in the FFT degrades performance. Fortunately, the techniques to minimize the attenuation, as discussed in Section 6.3, do not trade-off significantly with power. Consequently, all techniques that mitigate attenuation are incorporated for improved performance.

4. *Dummy switches:* The effect of clock-feedthrough and charge injection for each stage on the overall SNDR is simulated. Dummy switches are selected for stages where the overall SNDR is otherwise not met.

5. *Settling optimizing switches:* The overall settling error is simulated, and settling optimization switches are chosen for stages where the SNDR specification is otherwise note met.

6. *Sampler switch size:* The minimum switch size that provides adequate sampler settling for the required SNDR is determined.

7. *Settling:* The minimum per stage settling required for the overall SNDR is selected.

6.5.2 Energy Optimization

The exact switch sizing in each stage, as well as the V_{dd} employed, trade-off with the energy consumption. The energy optimization algorithm is outlined below.

1. *Supply voltage:* The triode resistance of the switch is approximated for short channel devices in deep triode[10] as:

$$R_{\text{triode}} \propto \left(\frac{1+U_0 V_{ov}}{V_{ov}}\right) \approx \left(\frac{1}{V_{ov}}\right)^p$$

Using this approximation, the switch resistance is $R_{sw} \propto W^{-1}\left(V_{ov}\right)^{-p}$. For a constant R_{sw}, $W \propto \left(V_{ov}\right)^{-p} = \left(V_{dd} - \left(V_{tn} + V_{cm}\right)\right)^{-p}$. In order to calculate the energy per switch operation, we compute $\frac{1}{2} C_{gs} V_{dd}^2 \propto W \cdot V_{dd}^2$. Using these equations, we compute the energy per switch operation for a constant switch on-resistance:

$$E_{sw}\Big|_{\text{const.}R_{sw}} \propto \frac{V_{dd}^2}{\left(V_{dd} - \left(V_{tn} + V_{cm}\right)\right)^p} \tag{6.9}$$

[10] For the devices used in CRAFT, $p = 0.50$ provided an accurate empirical fit.

This is plotted in Fig. 6.20. As seen, this curve has a unique minimum energy. The minimum occurs at $V_{dd,\text{opt}} = \frac{2}{2-p}(V_{tn} + V_{cm})$. This optimum V_{dd} is then customized per stage depending on the varying voltage swings as a result of attenuation[11] .

2. *Switch size:* For calculating the optimal switch width note, that each V_{dd} corresponds to a particular switch width (for a given resistance) on the constant-resistance plot. The maximum allowable nominal resistance can be calculated based on the required settling and allocated time chosen in Section 6.5.1. Therefore, from the V_{dd} chosen above, and the maximum allowable resistance, the optimal switch width, W, for each stage, is calculated.

3. *Time allocation:* Note that the energy per stage is dependent on the required switch resistance, and consequently, the time allocated per stage. The per-stage time allocated is now considered as the last optimization variable, and is redistributed (instead of the equal distribution assumed earlier) to optimize the total energy further. For the new allocated times, new optimal switch widths are determined.

4. *Half-dummies:* Finally, the use of half-dummies in specific stages is revisited. Although unlikely, it might be possible to hit the required SNDR by providing

[11] For the optimization described, it is assumed that 2 different V_{dd}s are available for optimization to cover a general case. The optimization algorithm can be easily modified for the specific case using a single or multiple V_{dd}s, based on availabilitiy.

more time (reducing switch width and resulting clock-feedthrough) to the SNDR-limiting stage instead of adding half-dummies. This possibility is also verified before finalizing the design.

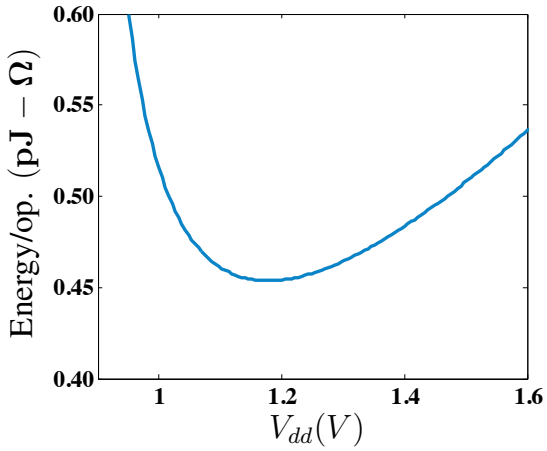

Figure 6.20: Optimum V_{dd} for minimum energy per operation

6.6 Measurement Results

The design has been implemented in the IBM 65nm CMOS process. Measurement results are shown in Fig. 6.22−6.26(a). The measurements shown have been calibrated to compensate for systematic offsets due to parasitics. The calibration process used, and a suggested improvement in the process are detailed below.

6.6.1 Calibration

The accuracy of the CRAFT operations is dependent on the matching between the capacitors used to realize it. Any systematic mismatch between the capacitors causes

computation errors that reduce the dynamic range of the system. However, since these errors are systematic, and input independent, it is possible to calibrate out the errors after conversion to digital. For our measurements, we use a simple calibration technique described below.

First, the non-idealities in the CRAFT operation are represented using a modified FFT matrix, $\mathbf{F'}$, that includes the effects of mismatch and represents the non-ideal CRAFT operation. The resultant outputs are represented by $\mathbf{X'} : \mathbf{X'} = \frac{1}{k}\,\mathbf{F'}\,\mathbf{x}$. A simple calibration scheme is constructed by observing that the FFT matrix, $\mathbf{F}$ is comprised of 256 elements. A set of inputs (16 samples) provide a set of 16 outputs corresponding to 16 linear equations. Consequently, a set of 16 mutually independent sets of inputs: vectors $\mathbf{x_i}$, with entries $x_i(t)$ constitute 256 independent equations. Using the measured results, $\mathbf{X'_i}$, with entries $X'_i(n)$, the 256 elements of the non-ideal $\mathbf{F'}$ matrix can be determined.

To generate the linearly independent inputs, 15 separate single-tone inputs, each centered on a non-DC bin i, are applied to the system. The inputs are at -6dBFS (dB of full-scale) so as to avoid introducing circuit non-linearity effects during calibration. The 16th independent input constituted an output on the DC bin, which can not be stimulated due to the ac coupled nature of the system. Instead, an ideal DC response is assumed. Synchronous averaging of 500 measurements runs, for each $\mathbf{X'_i}$, is used to reduce noise.

The ideal time-domain inputs, $\mathbf{x}$, are calculated based on the magnitude and phase

of the measurement bin peak $\mathbf{X}'(i)$. Since the actual measured phase and magnitude of the on-bin terms are used, the attenuation caused by the FFT implementation (k term of non-ideal FFT: $\mathbf{X}' = \frac{1}{k}\mathbf{F}'\mathbf{x}$) is incorporated. Also, the test setup and sampler frequency response effect are partially removed.

After determining an estimate of $\mathbf{F}'$, a correction matrix, $\hat{\mathbf{H}}$ is computed.

$$\mathbf{F}' = \hat{\mathbf{X}}'\hat{\mathbf{x}}^{-1} \quad \Rightarrow \hat{\mathbf{H}} = \mathbf{F}(\mathbf{F}')^{-1} = \mathbf{F}\hat{\mathbf{x}}(\hat{\mathbf{X}}')^{-1}$$

where $\hat{\mathbf{x}}$ is the calibration input and $\hat{\mathbf{X}}'$ is the calibration output.

All subsequent measurements across different magnitude and frequency inputs are then calibrated by applying $\hat{\mathbf{H}}$ to $\mathbf{X}'$ $\Rightarrow \mathbf{X} = \mathbf{F}\mathbf{x} = \hat{\mathbf{H}}\mathbf{F}'\mathbf{x} = \hat{\mathbf{H}}\mathbf{X}'$.

6.6.2 Test Setup

The test setup is shown in Fig. 6.17. I and Q inputs are generated using the Tektronix AWG-7122B arbitrary waveform generator and input to the CRAFT engine using 50Ω terminated probe-pads, which feed the sampler array. The latched outputs are externally buffered, and captured by external ADCs controlled by an FPGA (NI-7811R) programmed using LabVIEW®.

Note that the CRAFT processor runs at an input/output rate of 5GS/s per I and Q channel. Additionally, the outputs are analog values. Due to the very high speed of operation, and the limitations in the number of I/O pins, the outputs are first latched using the OTA-based analog latches described in Section 6.3.3, and multiplexed out at a slower rate limited to 40MS/s. The CRAFT speed is not compromised due to the

output read-out limitation. This output multiplexing and read-out is controlled by an FPGA (NI-7811R) programmed using LabVIEW®. The outputs are first buffered, then digitized by external ADCs, and finally captured by the FPGA. Using this scheme, RF inputs that are synchronous as well as asynchronous to the clock, can be captured and phase aligned. The latter is particularly important to emulate practical scenarios.

The entire output data chain was designed to have the lowest possible non-ideality floor. However, even with the best available discrete components, the output test setup non-linearity floor is limited to only −48dBFS. Note that the AWG resolution is also limited to 8 bits, limiting the measurement resolution.

The time domain input and output characteristics for a signal on the first bin is shown in Fig. 6.21. The input signal is sampled on multiple phases causing the output to rotate as shown in the figure.

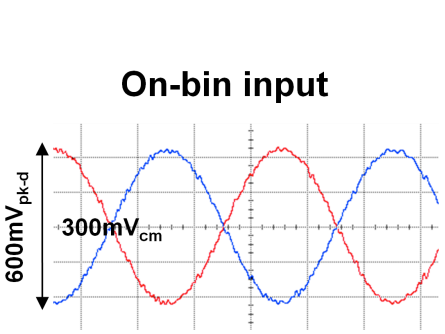

(a) Example on-bin input

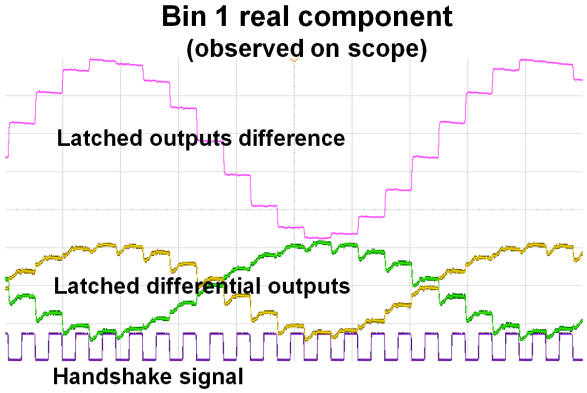

(a) Output (rotation due to asynchronous triggering)

Figure 6.21: Time domain latched output of CRAFT, single bin, I, differential

6.6.3 On-bin 1-tone Measurements

Fig. 6.22 shows the CRAFT output magnitudes ($\sqrt{\Re^2 + \Im^2}$) with a 312.5MHz ($= \frac{5\text{GHz}}{16}$) single-tone input at 5GS/s. Curve I shows the measured uncalibrated output magnitude across 16 bins. Curve II in Fig. 6.22 shows the calibrated plot depicting the circuit noise floor (including the integrated noise of the analog latches) at -46dB. The test setup noise floor is at -48dB as shown. To explore the non-linearity floor, a synchronous average over 500 measurements is used. The resultant Curve III shows the non-linearity floor with 43dB SFDR[12] and 48dB SNDR[13] . The non-linearity predicted by our system level simulations (not including twiddle factor errors) is shown in Curve IV. The discrepancy with Curve III is attributed to non-idealities in the AWG generated inputs that provide only a 8-bit resolution. Unfortunately, despite the use of state-of-the-art test equipment, the limitations of the input and output test setup limit the observable non-idealities in CRAFT.

[12] SFDR for a one-tone test is calculated as the difference between a full-scale on-bin signal and the largest off-bin output caused by the CRAFT non-linearity/calibration errors.

[13] SNDR is calculated as follows:

$$SNDR = 20 \times \log_{10}\left(\frac{\sqrt{\sum_{k=1}^{N} V_{ideal}^2(k)}}{\sqrt{\frac{1}{N}\sum_{k=1}^{N}\{V_{meas}(k) - V_{ideal}(k)\}^2}}\right) \tag{6.10}$$

where N is the number of FFT bins. Since the outputs will be observed/digitized on a per bin basis, the total (noise + distortion) in the denominator is averaged over the N bins to yield the average (noise + distortion) per bin.

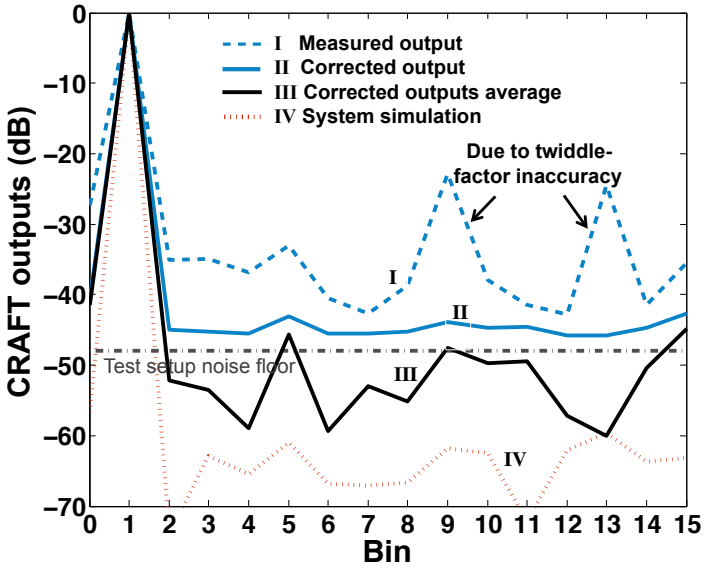

Figure 6.22: CRAFT outputs with a 362.5MHz single-tone input sampled at 5GHz

SNDR Variation

The SNDR at 1, 3, and 5 GS/s for a single-tone input frequency placed at different bins is shown in Fig. 6.23(a). The average SNDR across bins is ≈ 50dB at 1 & 3 GHz and degrades to 47dB at 5GHz; SNDR better than 45dB is maintained across all frequencies. The SNDR and SFDR measurements are tabulated in Table 6.2. This provides 7-8 bits of spectrum sensing resolution over a 5GHz (2.5GHz×2 due to I, Q inputs) frequency range in the digital back-end. The output magnitudes across bins for these single-tone tests are shown in Fig. 6.23(b).

The CRAFT magnitude outputs with varying input amplitudes for a single-tone, on-bin 312.5MHz input at 5GS/s are shown in Fig. 6.25(a). The input is varied over a 18dB range. As seen, the circuit is limited by the noise floor for low input amplitudes,

while it becomes non-linearity limited at larger input amplitudes.

Fig. 6.24(a) plots the output SNDR versus the input amplitude. A fourth-order fit shows a linear SNDR improvement with increasing amplitude before being limited by a combination of the circuit non-linearities, and the AWG resolution (8 bits).

Table 6.2: Table showing the CRAFT output SNDR and SFDR with full-scale inputs for single-tone on-bin inputs

	1GS/s		3GS/s		5GS/s	
Bin	**SNDR**	**SFDR**	**SNDR**	**SFDR**	**SNDR**	**SFDR**
1	46.6	37.7	45.8	35.8	48.1	42.6
2	50.9	42.8	50.7	43.8	48.4	39.7
3	51.5	44.7	49.8	44.8	45.8	41.8
4	48.9	37.9	51.4	40.2	51.1	41.1
5	50.2	44.4	47.7	40.7	46.2	42.5
6	51.4	44.8	48.0	40.1	45.0	36.3
7	48.1	40.2	48.2	42.6	45.1	40.1
8	56.5	51.2	61.2	57.2	49.9	42.8
9	48.5	41.3	47.5	41.2	45.2	39.6
10	53.7	47.8	46.6	39.6	44.5	35.6
11	50.3	44.9	47.6	40.0	46.6	43.3
12	49.2	37.9	52.2	41.4	50.9	41.6
13	52.0	46.2	50.0	45.5	45.0	39.9
14	53.8	46.6	50.9	43.0	48.5	38.8
15	47.1	39.0	45.4	35.9	48.9	42.3
Min	46.6	37.7	45.4	35.8	44.5	35.6
Max	56.5	47.8	61.2	45.5	51.1	43.3
Mean	**50.6**	**43.2**	**49.5**	**42.1**	**47.3**	**40.5**
Std. dev.	**2.7**	**4.1**	**3.8**	**5.0**	**2.3**	**2.3**

6.6.4 On-bin 2-tone Measurements

Results from a two-tone test with tones on adjacent bins are shown in Fig. 6.25(b). Two single-tone measurements and their superposition (adjusted to the sampler input full-scale) are also shown. The difference between the superposition and the measured two-tone output is indicative of the additional non-linearity introduced. The relative increase in bins 13 and 14 is likely to be due to uncorrected twiddle factor errors in

stage 2 of the CRAFT engine. Blue Note that the effect of these twiddle factor errors, leaking signal on to the negative of the signal frequency, is identical to the effect of I-Q mismatch errors in a homodyne receiver.

6.6.5 Power Consumption

The CRAFT core consumes 12.2pJ/conv. of energy. At 5GS/s, it consumes a total power of 3.2mW. Measurements of the energy consumption versus supply voltage and frequency are depicted in Fig. 6.26(a). These measurements clearly show the expected digital-like relationship of the CRAFT energy with frequency and supply voltage. This further corroborates our premise that CRAFT is expected to respond favorably to technology scaling.

6.7 Conclusion

This chapter describes a wideband ultra-low power RF front-end channelizer based on a 16 point FFT. The design is based on a charge re-use technique that enables it to run at speeds of 5GS/s with a 47dB SNDR, while consuming only 12.2pJ/conv.

The chapter details the design of the CRAFT engine. It also describes the interface circuitry, and the test setup required to test such a high-speed, high dynamic range system. Non-idealities in the CRAFT computations are discussed, analytical models derived, and new circuit techniques developed for reducing these non-idealities. These techniques can be easily extended to other passive switched capacitor circuits to improve

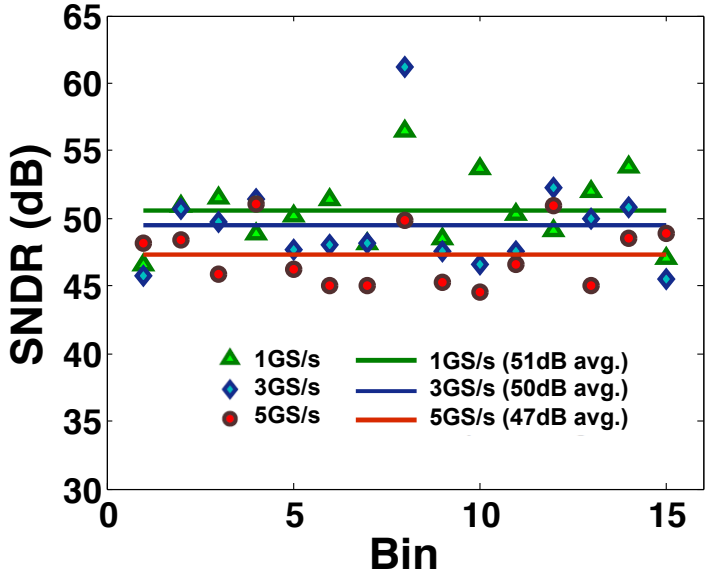

(a) SNDR variation across bins at 3 sampling rates

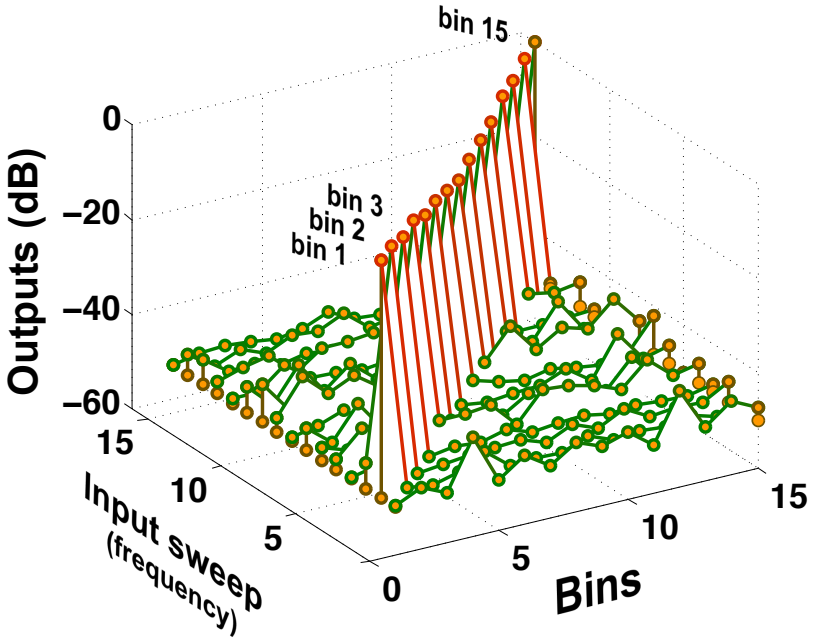

(b) Single-tone tests for each bin for the 15 non-zero bins

Figure 6.23: Measurement results including SNDR variation across bins, and frequencies: (a), and 1-tone measurements: (b)

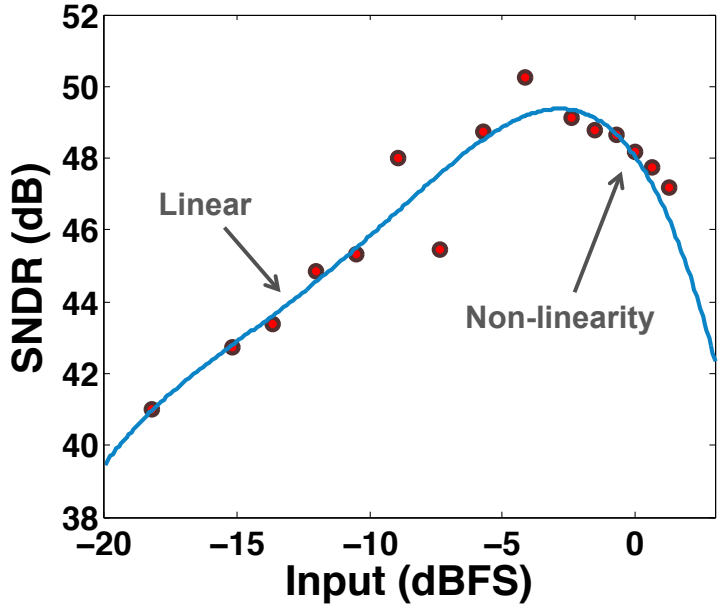

(a) SNDR variation vs. input amplitudes at 5GS/s

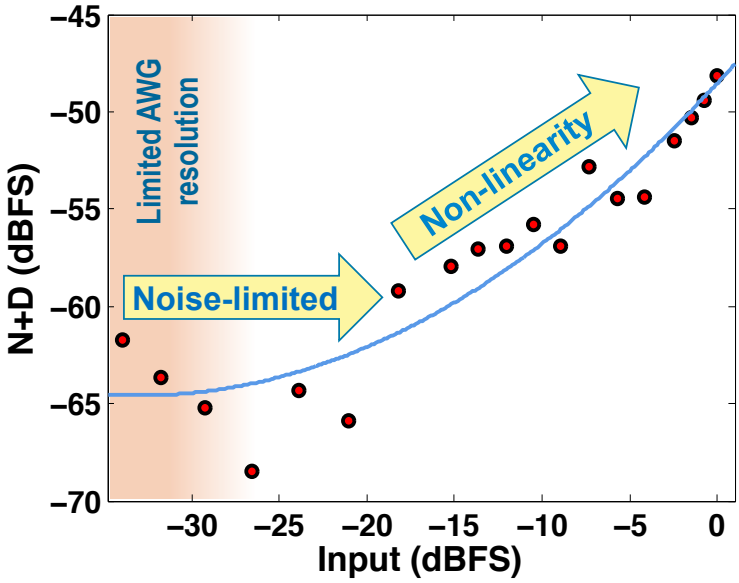

(b) Total noise and distortion with varying input amplitude

Figure 6.24: Measurement results including SNDR, and noise and distortion across input amplitudes: (a), (b)

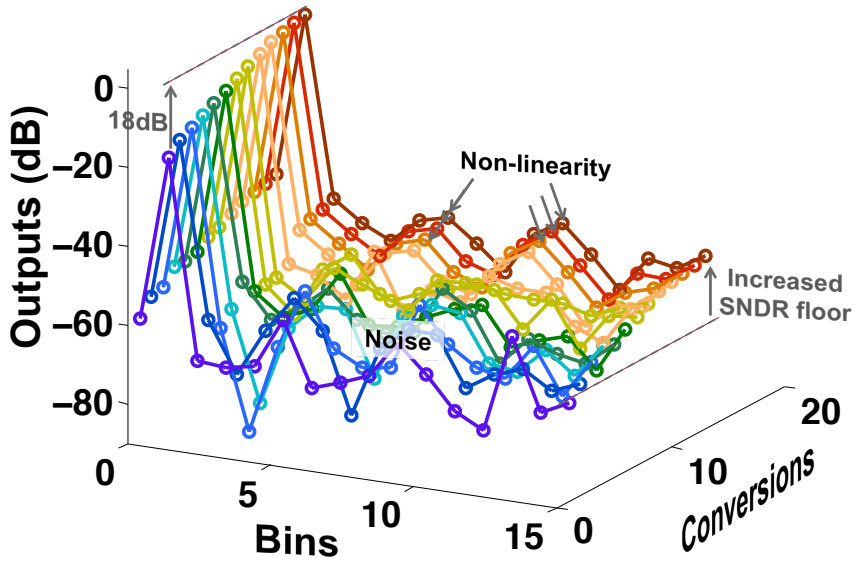

(a) The noise and non-linearity in the CRAFT output with increasing amplitude of a single-tone on-bin input at 5GS/s

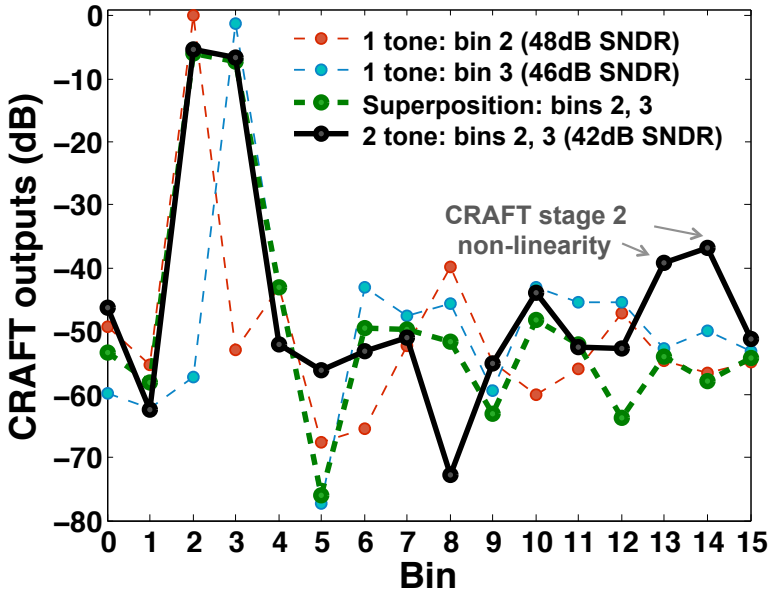

(b) A two-tone non-linearity test on CRAFT

Figure 6.25: Measurement results including noise and distortion across input amplitudes: (a), and results of a 2-tone test: (b)

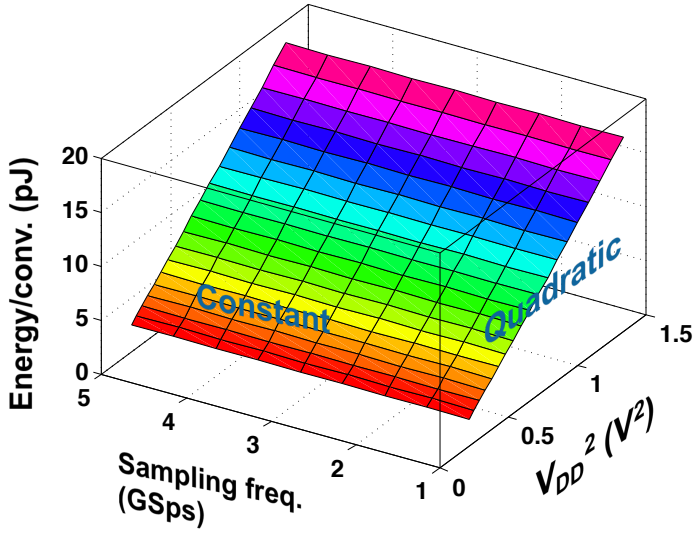

(a) The digital-like energy relation with f_{samp} and V_{dd}

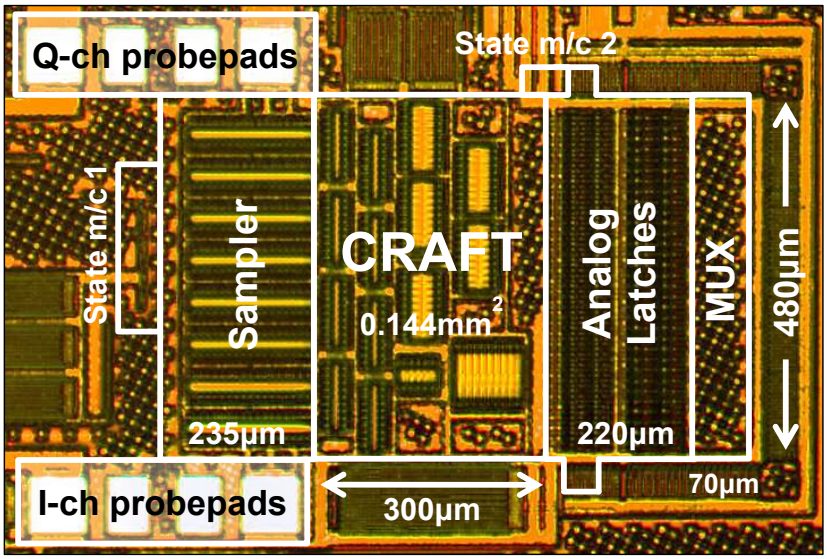

(b) Die photo of CRAFT and supporting circuits

Figure 6.26: The energy consumption relation with frequency and supply: (a), and a die photograph: (b)

their performance. Measurement results are then presented.

Table 6.3 compares the CRAFT performance with one digital and two analog domain FFT implementations. As shown, CRAFT operates at speeds 5X faster than previous state-of-the-art designs. Additionally, it consumes better than 28X lower energy. As an RF channelizer, it is expected to reduce digitization requirements enabling wideband digital spectrum sensing. As a result, it helps advance the state-of-the-art for wideband SDR architectures.

Moreover, this technique of performing charge domain operations prior to sampling opens up a large number of possibilities. The CRAFT engine enables low power wideband digitization of RF signals, that can be utilized for spectrum sensing and real-time analysis, and can be extended to signal intelligence applications, as well as electronic warfare. Moreover, by virtue of its orthogonal frequency domain transform characteristics, the CRAFT engine can be used in phased arrays utilizing the FFT for beamforming, or in spatio spectral beamformers [125].

Table 6.3: Table for comparison with other FFT implementations

Ref	Domain	# of bins	SNDR (dB)	Power (mW)	Samp. speed (GS/s)	Scaled* E/conv. (pJ/conv.)	E/conv. ratio
[114]	Current	64	__	389	1.2	345.8	28X
[115]	Current	8	36	19	1	405.3	33X
[120]	Digital	128	51†	175	1	1600	131X
This work	**Charge**	**16**	**47‡**	**0.92**	**5**	**12.2**	**1X**

*Scaled for complexity similar to the scaling used in [115]
†8.5bit ENOB assumed for 10bit internal word length
‡After twiddle factor correction

Chapter 7

Conclusions & Contributions

The realization of the alluring vision of a cognitive radio requires novel and disruptive RF circuit architectures, and innovative circuits. This has motivated a deluge of exciting research in this area in the past decade, and will continue to spur further research into the next. At this time, a highly reconfigurable RF architecture seems feasible by employing new wideband circuits that have been recently developed. However, the realization of the ideal cognitive radio, with the level of versatility described by early descriptions [6], requires further innovation, probably spanning the next decade.

In this thesis, a number of circuit architectures and techniques in recent literature are discussed. Their suitability for cognitive radio applications is briefly reviewed. New architectures and circuits are described, and results from early prototypes are presented for demonstration.

7.1 Remarks

Based on the current research directions, the following observations regarding on-chip SDR architectures for portable cognitive radios can be distilled. These remarks are summarized in Fig. 7.1 and Fig. 7.2. The contributions described in this thesis are highlighted in bold font.

SDR Signaling:

1. For flexibility through reconfigurability in the frequency domain, frequency agility is an important characteristic. Wide-tuning frequency agile architectures and circuits are critical for signaling applications.

2. To avoid the reciprocal mixing of large blockers onto small signals in wideband architectures, low LO phase noise is important.

3. Programmable RF filters are required in the front-end to provide the required frequency agility while relaxing the linearity requirements of the circuits that follow.

4. Programmable baseband filters are required to handle multiple radio-access technologies.

SDR Spectrum Sensing:

1. For the cognitive radio transceiver with its emphasis on flexibility and versatility,

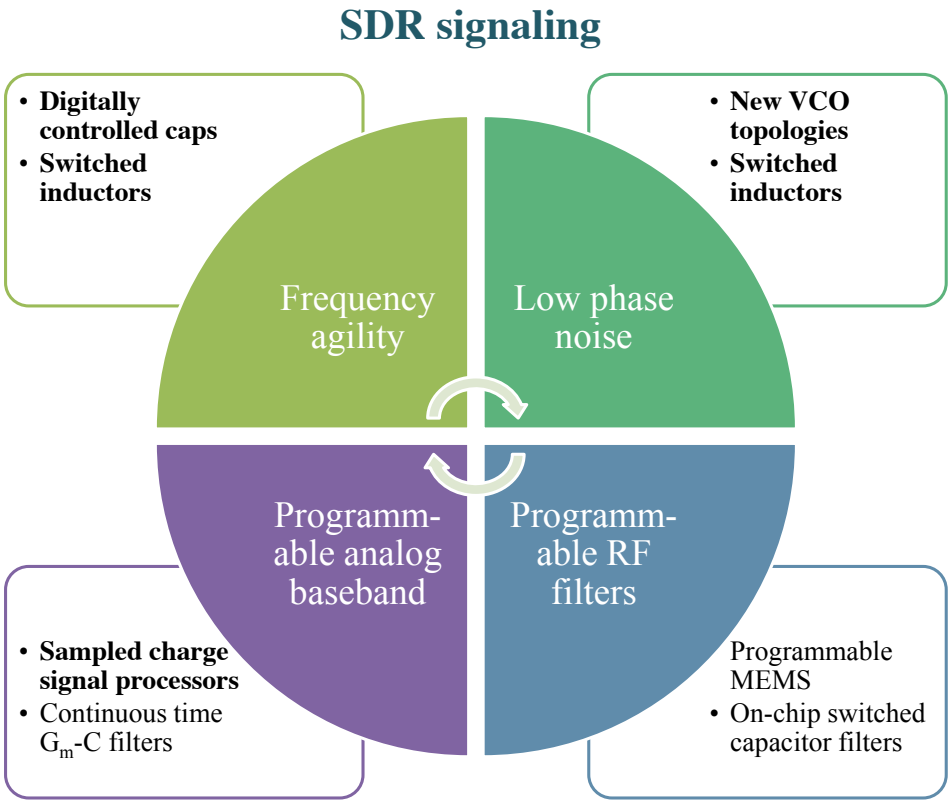

Figure 7.1: Summary of requirements for spectrum sensing in cognitive radios

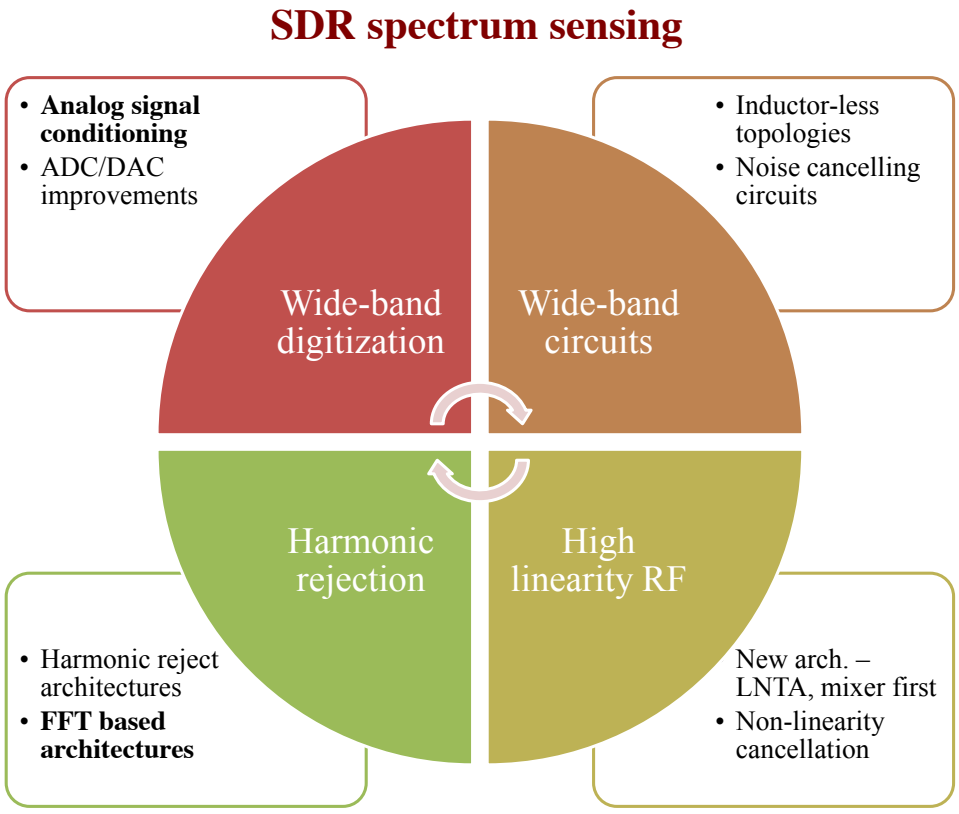

Figure 7.2: Summary of requirements for signaling in cognitive radios

the paradigm of performing most functions in the digital domain still remains the popular direction. As a result, the wideband digitization of the RF signal becomes one of the most significant bottle-necks in the realization of cognitive radios.

2. To tackle the wideband digitization problem, analog domain signal conditioning prior to digitization is being considered. Discrete time signal processing is emerging as a popular option.

3. For spectrum sensing, RF circuits and architectures with wideband characteristics, i.e., relatively constant performance (matching, noise-figure) over a large bandwidth are essential.

4. As a result of the large bandwidth of interest, circuit linearity becomes extremely critical. At the receiver, large in-band blockers undergo distortion in non-linear circuits destroying smaller wanted signals (inter-modulation, cross-modulation, gain-compression, etc.).

5. Due to the large bandwidth of operation, the LO harmonics fall within band; consequently, harmonic mixing is a critical issue in both the transmitter and receiver.

7.2 Contributions

In this thesis, the following major contributions were made:

7.2.1 Signaling: Frequency Agility using Wide-tuning VCOs

1. Switched inductor based VCOs were proposed as a viable method to obtain very wide tuning ranges while maintaining good phase noise and power consumption performance across the entire tuning range.

2. The first comprehensive analytical model of a switched inductor VCO was developed, and a design framework based on it was introduced.

3. Techniques for large switched capacitor array design and layout were introduced to overcome the effect of interconnect parasitics. Switch sizing techniques for optimal tuning range were explored.

4. Two prototype designs, one with a single inductor switch, and one with three inductor switches were measured. The results obtained were shown to outperform state-of-the-art LC VCOs for SDR applications.

5. A family of low phase noise VCO topologies based on capacitive feedback (LiTVCOs) were introduced. These LiTVCOs achieve large oscillation amplitudes by decoupling the tank from the active devices. They also achieve low active device noise injection into the tank. Consequently, superior phase noise performance is achieved.

6. A prototype PLL incorporating a LiTVCO based on the capacitive feedback technique was designed. The LiTVCO achieved excellent phase noise and FOM over

a large tuning range.

7.2.2 Spectrum Sensing: Charge Domain RF Front-end for Low-power Digitization

1. RF front-end sampled charge processing was explored using analytical system models as a mechanism for signal conditioning to ease digitization.

2. A wideband spectrum sensing architecture based on a passive switched capacitor DFT based analog signal processing front-end was developed. This architecture was shown to be capable of wideband, low-power digitization in the frequency domain.

3. Several non-idealities in passive switched capacitor were explored and modeled, and techniques to mitigate these in sampled charge systems were developed.

4. An ultra-low power, high speed charge domain FFT prototype for an SDR receiver front-end was designed and demonstrated in measurement. The prototype demonstrates orders of magnitude improvement in performance over other competing architectures/circuits.

A summary of the thesis work is shown in Fig. 7.3.

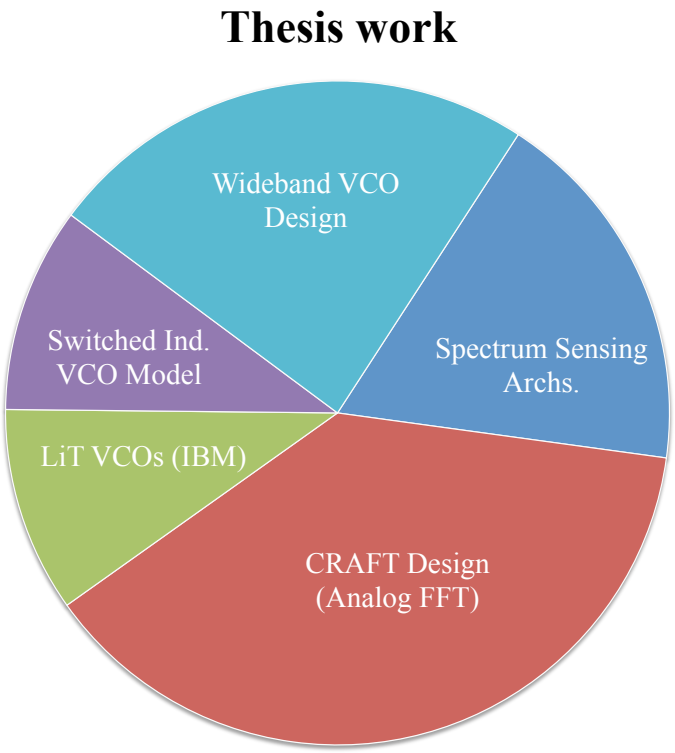

Figure 7.3: Summary of thesis work

References

[1] B. Fette, *Cognitive Radio Technology*, Newnes, 2006.

[2] NTIA, "U.S. frequency allocations," *http://www.ntia.doc.gov/osmhome/allochrt.pdf*.

[3] D. Cabric, I.D. O'Donnell, M.S.-W. Chen, and R.W. Brodersen, "Spectrum sharing radios," *IEEE Circuits and Systems Magazine*, vol. 6, no. 2, pp. 30 – 45, 2006.

[4] Federal Communications Commission, "In the matter of facilitating opportunities for exible, efcient, and reliable spectrum use employing cognitive radio technologies, report and order," *FCC 05-57, ET Docket No. 03-108*, 2005.

[5] T. Ulversoy, "Software defined radio: Challenges and opportunities," *IEEE Communications Surveys Tutorials*, vol. 12, no. 4, pp. 531 – 550, 2010.

[6] J. Mitola III, "Cognitive radio: an integrated agent architecture for software defined radio dissertation," *Royal Institute of Technology, Stockholm*, 2000.

[7] J. Mitola III, "Software radios−survey, critical evaluation and future directions," in *Telesystems Conference, 1992. NTC-92., National*, May 1992.

[8] R. Bagheri, A. Mirzaei, S. Chehrazi, M. E. Heydari, M. Lee, M. Mikhemar, W. Tang, and A. A. Abidi, "An 800-MHz to 6-GHz software-defined wireless receiver in 90-nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 41, no. 12, pp. 2860 – 2876, November 2006.

[9] J. Gambini, O. Simeone, Y. Bar-Ness, U. Spagnolini, and T. Yu, "Packet-wise vertical handover for unlicensed multi-standard spectrum access with cognitive radios," *IEEE Transactions on Wireless Communications*, vol. 7, no. 12, pp. 5172 – 5176, December 2008.

[10] O. Gaborieau, S. Mattisson, N. Klemmer, B. Fahs, F.T. Braz, R. Gudmundsson, T. Mattsson, C. Lascaux, C. Trichereau, W. Suter, E. Westesson, and A. Nydahl, "A SAW-less multiband WEDGE receiver," in *IEEE International Solid-State Circuits Conference*, February 2009, pp. 114 –115,115a.

[11] T. Sowlati, B. Agarwal, J. Cho, T. Obkircher, M. El Said, J. Vasa, B. Ramachandran, M. Kahrizi, E. Dagher, Wei-Hong Chen, M. Vadkerti, G. Taskov, U. Seckin, H. Firouzkouhi, B. Saeidi, H. Akyol, Yunyoung Choi, A. Mahjoob, S. D'Souza, Chieh-Yu Hsieh, D. Guss, D. Shum, D. Badillo, I. Ron, D. Ching, Feng Shi, Yong He, J. Komaili, A. Loke, R. Pullela, E. Pehlivanoglu, H. Zarei, S. Tadjpour, D. Agahi, D. Rozenblit, W. Domino, G. Williams, N. Damavandi,

S. Wloczysiak, S. Rajendra, A. Paff, and T. Valencia, "Single-chip multiband WCDMA/HSDPA/HSUPA/EGPRS transceiver with diversity receiver and 3G DigRF interface without SAW filters in transmitter / 3G receiver paths," in *IEEE International Solid-State Circuits Conference*, February 2009, pp. 116 –117,117a.

[12] Sang-June Park, Kok-Yan Lee, and G.M. Rebeiz, "Low-loss 5.15–5.70-GHz RF MEMS switchable filter for wireless LAN applications," *IEEE Transactions on Microwave Theory and Techniques*, vol. 54, no. 11, pp. 3931 – 3939, November 2006.

[13] K. Entesari and G.M. Rebeiz, "A differential 4-bit 6.5-10-GHz RF MEMS tunable filter," *IEEE Transactions on Microwave Theory and Techniques*, vol. 53, no. 3, pp. 1103 – 1110, March 2005.

[14] Chih-Chieh Cheng and G.M. Rebeiz, "High-Q 4–6-GHz suspended stripline RF MEMS tunable filter with bandwidth control," *IEEE Transactions on Microwave Theory and Techniques*, vol. 59, no. 10, pp. 2469 – 2476, October 2011.

[15] G. Rebeiz, K. Entesari, I. Reines, S.-j. Park, M. El-tanani, A. Grichener, and A. Brown, "Tuning in to RF MEMS," *IEEE Microwave Magazine*, vol. 10, no. 6, pp. 55 – 72, October 2009.

[16] R.M. Young, J.D. Adam, C.R. Vale, T.T. Braggins, S.V. Krishnaswamy, C.E. Milton, D.W. Bever, L.G. Chorosinski, Li-Shu Chen, D.E. Crockett, C.B. Freidhoff,

S.H. Talisa, E. Capelle, R. Tranchini, J.R. Fende, J.M. Lorthioir, and A.R. Tories, "Low-loss bandpass RF filter using MEMS capacitance switches to achieve a one-octave tuning range and independently variable bandwidth," in *IEEE MTT-S International Microwave Symposium Digest*, June 2003, vol. 3, pp. 1781 – 1784 vol.3.

[17] Yun Zhu, R.W. Mao, and C.S. Tsai, "A varactor and FMR-tuned wideband band-pass filter module with versatile frequency tunability," *IEEE Transactions on Magnetics*, vol. 47, no. 2, pp. 284 – 288, February 2011.

[18] M. Sanchez-Renedo, R. Gomez-Garcia, J.I. Alonso, and C. Briso-Rodriguez, "Tunable combline filter with continuous control of center frequency and bandwidth," *IEEE Transactions on Microwave Theory and Techniques*, vol. 53, no. 1, pp. 191 – 199, January 2005.

[19] J. Nath, D. Ghosh, J.-P. Maria, A.I. Kingon, W. Fathelbab, P.D. Franzon, and M.B. Steer, "An electronically tunable microstrip bandpass filter using thin-film Barium-Strontium-Titanate (BST) varactors," *IEEE Transactions on Microwave Theory and Techniques*, vol. 53, no. 9, pp. 2707 – 2712, September 2005.

[20] Z. Ru, E.A.M. Klumperink, C.E. Saavedra, and B. Nauta, "A 300–800 MHz tunable filter and linearized LNA applied in a low-noise harmonic-rejection RF-sampling receiver," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 5, pp. 967 – 978, May 2010.

[21] L. E. Franks and I. W. Sandberg, "An alternative approach to the realizations of network functions: N-path filter," *Bell Systems Technical Journal*, vol. 59, no. 1, pp. 1321 – 1350, 1960.

[22] A. Mirzaei, H. Darabi, and D. Murphy, "Architectural evolution of integrated M-phase high-Q bandpass filters," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 59, no. 1, pp. 52 – 65, January 2012.

[23] A. Ghaffari, E.A.M. Klumperink, M.C.M. Soer, and B. Nauta, "Tunable high-Q N-path band-pass filters: Modeling and verification," *IEEE Journal of Solid-State Circuits*, vol. 46, no. 5, pp. 998 – 1010, may 2011.

[24] M.C.M. Soer, E.A.M. Klumperink, P.-T. de Boer, F.E. van Vliet, and B. Nauta, "Unified frequency-domain analysis of switched-series- passive mixers and samplers," *IEEE Transactions on Circuits and Systems I*, vol. 57, no. 10, pp. 2618 – 2631, October 2010.

[25] J.A. Weldon, R.S. Narayanaswami, J.C. Rudell, Li Lin, M. Otsuka, S. Dedieu, Luns Tee, King-Chun Tsai, Cheol-Woong Lee, and P.R. Gray, "A 1.75-GHz highly integrated narrow-band CMOS transmitter with harmonic-rejection mixers," *IEEE Journal of Solid-State Circuits*, vol. 36, no. 12, pp. 2003 – 2015, December 2001.

[26] H. Darabi, "Highly integrated and tunable RF front-ends for reconfigurable multiband transceivers," in *IEEE Custom Integrated Circuits Conference*, September

2010, pp. 1 – 8.

[27] A. Bevilacqua and A.M. Niknejad, "An ultrawideband CMOS low-noise amplifier for 3.1-10.6-GHz wireless receivers," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 12, pp. 2259 – 2268, December 2004.

[28] A. Ismail and A.A. Abidi, "A 3-10-GHz low-noise amplifier with wideband LC-ladder matching network," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 12, pp. 2269 – 2277, December 2004.

[29] F. Bruccoleri, E.A.M. Klumperink, and B. Nauta, "Generating all two-MOS-transistor amplifiers leads to new wide-band LNAs," *IEEE Journal of Solid-State Circuits*, vol. 36, no. 7, pp. 1032 – 1040, July 2001.

[30] J. Janssens, M. Steyaert, and H. Miyakawa, "A 2.7 volt CMOS broadband low noise amplifier," in *IEEE Symposium on VLSI Circuits*, June 1997, pp. 87 – 88.

[31] F. Bruccoleri, E.A.M. Klumperink, and B. Nauta, "Wide-band CMOS low-noise amplifier exploiting thermal noise canceling," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 2, pp. 275 – 282, February 2004.

[32] F. Tzeng, A. Jahanian, and P. Heydari, "A multiband inductor-reuse CMOS low-noise amplifier," *IEEE Transactions on Circuits and Systems II*, vol. 55, no. 3, pp. 209 – 213, March 2008.

[33] H. Hashemi and A. Hajimiri, "Concurrent multiband low-noise amplifiers-theory, design, and applications," *IEEE Transactions on Microwave Theory and Techniques*, vol. 50, no. 1, pp. 288 – 301, January 2002.

[34] M. El-Nozahi, E. Sanchez-Sinencio, and K. Entesari, "A CMOS low-noise amplifier with reconfigurable input matching network," *IEEE Transactions on Microwave Theory and Techniques*, vol. 57, no. 5, pp. 1054 – 1062, May 2009.

[35] B. Sadhu, J. Kim, and R. Harjani, "A CMOS 3.3-8.4 GHz wide tuning range, low phase noise LC VCO," *Proceedings of IEEE Custom Integrated Circuits Conference*, pp. 559 – 562, September 2009.

[36] B. Razavi, "Cognitive radio design challenges and techniques," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 8, pp. 1542 – 1553, August 2010.

[37] P. Andreani and S. Mattisson, "On the use of MOS varactors in RF VCOs," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 6, pp. 905 – 910, June 2000.

[38] R.L. Bunch and S. Raman, "Large-signal analysis of MOS varactors in CMOS -G_m LC VCOs," *IEEE Journal of Solid-State Circuits*, vol. 38, no. 8, pp. 1325 – 1332, August 2003.

[39] A. Kral, F. Behbahani, and A. A. Abidi, "RF-CMOS oscillators with switched tuning," *Proceedings in IEEE Custom Integrated Circuits Conference*, May 1998.

[40] A. D. Berny, A. M. Niknejad, and R. G. Meyer, "A 1.8-GHz LC VCO with 1.3-GHz tuning range and digital amplitude calibration," *IEEE Journal of Solid-State Circuits*, pp. 909 – 917, April 2005.

[41] Z. Ru, N.A. Moseley, E. Klumperink, and B. Nauta, "Digitally enhanced software-defined radio receiver robust to out-of-band interference," *IEEE Journal of Solid-State Circuits*, vol. 44, no. 12, pp. 3359 – 3375, December 2009.

[42] M. Soer, E. Klumperink, Z. Ru, F.E. van Vliet, and B. Nauta, "A 0.2-to-2.0GHz 65nm CMOS receiver without LNA achieving >11dBm IIP3 and <6.5 dB NF," in *IEEE International Solid-State Circuits Conference*, February 2009, pp. 222 – 223, 223a.

[43] H. Khorramabadi and P.R. Gray, "High-frequency CMOS continuous-time filters," *IEEE Journal of Solid-State Circuits*, vol. 19, no. 6, pp. 939 – 948, December 1984.

[44] J. Silva-Martinez, M. Steyaert, and W. Sansen, *High-performance CMOS continuous time filters*, Kluwer Academic Publishers, 1993.

[45] S. Pavan and Y. Tsividis, *High frequency continuous time filters in digital CMOS processes*, Kluwer Academic Publishers, 2000.

[46] R. Bagheri, A. Mirzaei, M.E. Heidari, S. Chehrazi, Minjae Lee, M. Mikhemar, W.K. Tang, and A.A. Abidi, "Software-Defined Radio receiver: dream to reality," *IEEE Communications Magazine*, vol. 44, no. 8, pp. 111 – 118, August 2006.

[47] K. Muhammad, Y.C. Ho, T. Mayhugh, C.M. Hung, T. Jung, I. Elahi, C. Lin, I. Deng, C. Fernando, J. Wallberg, S. Vemulapalli, S. Larson, T. Murphy, D. Leipold, P. Cruise, J. Jaehnig, M.C. Lee, R.B. Staszewski, R. Staszewski, and K. Maggio, "A discrete time quad-band GSM/GPRS receiver in a 90nm digital CMOS process," in *IEEE Custom Integrated Circuits Conference*, September 2005, pp. 809 – 812.

[48] K. Muhammad, D. Leipold, B. Staszewski, Y.-C. Ho, C.M. Hung, K. Maggio, C. Fernando, T. Jung, J. Wallberg, J.-S. Koh, S. John, I. Deng, O. Moreira, R. Staszewski, R. Katz, and O. Friedman, "A discrete-time bluetooth receiver in a 0.13 μm digital CMOS process," in *IEEE International Solid-State Circuits Conference*, February 2004, pp. 268 – 527 vol.1.

[49] K. Muhammad, R.B. Staszewski, and D. Leipold, "Digital RF processing: toward low-cost reconfigurable radios," *IEEE Communications Magazine*, vol. 43, no. 8, pp. 105 – 113, August 2005.

[50] R. van de Plassche, "A sigma-delta modulator as an A/D converter," *IEEE Transactions on Circuits and Systems*, vol. 25, no. 7, pp. 510 – 514, July 1978.

[51] M. Bolatkale, L.J. Breems, R. Rutten, and K.A.A. Makinwa, "A 4 GHz continuous-time $\Sigma - \Delta$ ADC with 70 dB DR and -74 dBFS THD in 125 MHz BW," *IEEE Journal of Solid-State Circuits*, vol. 46, no. 12, pp. 2857 –2868, December 2011.

[52] H. Shibata, R. Schreier, Wenhua Yang, A. Shaikh, D. Paterson, T. Caldwell, D. Alldred, and Ping Wing Lai, "A DC-to-1GHz tunable RF $\Delta\Sigma$ ADC achieving DR = 74dB and BW = 150MHz at f_0 = 450MHz using 550mW," in *IEEE International Solid-State Circuits Conference (ISSCC)*, February 2012, pp. 150 – 152.

[53] A.M.A. Ali, A. Morgan, C. Dillon, G. Patterson, S. Puckett, P. Bhoraskar, H. Dinc, M. Hensley, R. Stop, S. Bardsley, D. Lattimore, J. Bray, C. Speir, and R. Sneed, "A 16-bit 250-MS/s IF sampling pipelined ADC with background calibration," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 12, pp. 2602 – 2612, December 2010.

[54] S. Devarajan, L. Singer, D. Kelly, S. Decker, A. Kamath, and P. Wilkins, "A 16-bit, 125 MS/s, 385 mW, 78.7 dB SNR CMOS pipeline ADC," *IEEE Journal of Solid-State Circuits*, vol. 44, no. 12, pp. 3305 – 3313, December 2009.

[55] B. Murmann and B.E. Boser, "A 12-bit 75-MS/s pipelined ADC using open-loop residue amplification," *IEEE Journal of Solid-State Circuits*, vol. 38, no. 12, pp. 2040 – 2050, December 2003.

[56] C.C. Lee and M.P. Flynn, "A 12b 50MS/s 3.5mW SAR assisted 2-stage pipeline ADC," in *IEEE Symposium on VLSI Circuits*, June 2010, pp. 239 – 240.

[57] C-C. Liu, S-J. Chang, G-Y. Huang, and Y-Z. Lin, "A 10-bit 50-MS/s SAR ADC with a monotonic capacitor switching procedure," *IEEE Journal of Solid-State*

Circuits, vol. 45, no. 4, pp. 731 – 740, April 2010.

[58] Y. Zhu, C-H. Chan, U-F. Chio, S-W. Sin, S-P. U, R.P. Martins, and F. Maloberti, "A 10-bit 100-MS/s reference-free SAR ADC in 90 nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 6, pp. 1111 –1121, june 2010.

[59] S.M. Louwsma, A.J.M. van Tuijl, M. Vertregt, and B. Nauta, "A 1.35 GS/s, 10 b, 175 mW time-interleaved AD converter in 0.13μm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 4, pp. 778 – 786, April 2008.

[60] B. Murmann, "ADC performance survey 1997-2012," *http://www.stanford.edu/∼murmann/adcsurvey.html*.

[61] P.B. Kenington, "Linearized transmitters: an enabling technology for software defined radio," *IEEE Communications Magazine*, vol. 40, no. 2, pp. 156 – 162, Feb 2002.

[62] R. Shrestha, E.A.M. Klumperink, E. Mensink, G.J.M. Wienk, and B. Nauta, "A polyphase multipath technique for Software-Defined Radio transmitters," *IEEE Journal of Solid-State Circuits*, vol. 41, no. 12, pp. 2681 – 2692, December 2006.

[63] S-J. Kim and G.B. Giannakis, "Sequential and cooperative sensing for multi-channel cognitive radios," *IEEE Transactions on Signal Processing*, vol. 58, no. 8, pp. 4239 – 4253, August 2010.

[64] W.A. Gardner, "Signal interception: a unifying theoretical framework for feature detection," *IEEE Transactions on Communications*, vol. 36, no. 8, pp. 897 – 906, August 1988.

[65] S. M. Mishra, A. Sahai, and R. W. Brodersen, "Cooperative sensing among cognitive radios," *IEEE International Conference on Communications*, pp. 1658 – 1663, June 2006.

[66] Y. Ding and R. Harjani, "A +18 dBm IIP3 LNA in 0.35 μm CMOS," in *IEEE International Solid-State Circuits Conference*, 2001, pp. 162 – 163, 443.

[67] V. Aparin and L.E. Larson, "Modified derivative superposition method for linearizing FET low-noise amplifiers," *IEEE Transactions on Microwave Theory and Techniques*, vol. 53, no. 2, pp. 571 – 581, February 2005.

[68] W-H. Chen, G. Liu, B. Zdravko, and A.M. Niknejad, "A highly linear broadband CMOS LNA employing noise and distortion cancellation," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 5, pp. 1164 – 1176, May 2008.

[69] I. D. O'Donnel and R. W. Brodersen, "An ultra-wideband transceiver architecture for low power, low rate, wireless systems," *IEEE Transactions on Vehicular Technology*, pp. 1623 – 1631, September 2005.

[70] S. Hoyos, B.M. Sadler, and G.R. Arce, "Analog to digital conversion of ultra-wideband signals in orthogonal spaces," in *IEEE Conference on Ultra Wideband*

Systems and Technologies, November 2003, pp. 47 – 51.

[71] S. Hoyos, B.M. Sadler, and G.R. Arce, "Broadband multicarrier communication receiver based on analog to digital conversion in the frequency domain," *IEEE Transactions on Wireless Communications*, vol. 5, no. 3, pp. 652 – 661, March 2006.

[72] S. R. Valezquez, T. Q. Nguyen, S. R. Broadstone, and J. K. Roberge, "A hybrid filter bank approach to analog-to-digital conversion," *Proceedings of the IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis*, pp. 116 – 119, October 1994.

[73] W. Namgoong, "A channelized digital ultrawideband receiver," *IEEE Transactions for Wireless Communications*, pp. 502 – 510, May 2003.

[74] T-L Hsieh, P. Kinget, and R. Gharpurey, "A rapid interference detector for ultra wideband radio systems in 0.13μm CMOS," *IEEE Radio Frequency Integrated Circuits Symposium*, pp. 347 – 350, April 2008.

[75] M. Elbadry, B. Sadhu, J. Qiu, and R. Harjani, "Dual channel injection-locked quadrature LO generation for a 4GHz instantaneous bandwidth receiver at 21GHz center frequency," *IEEE Radio Frequency Integrated Circuits Symposium*, June 2012.

[76] S. Kalia, M. Elbadry, B. Sadhu, S. Patnaik, J. Qiu, and R. Harjani, "A simple, unified phase noise model for injection-locked oscillators," *IEEE Radio Frequency Integrated Circuits Symposium*, June 2011.

[77] B. Sadhu, U. Omole, and R. Harjani, "Modeling and synthesis of wide-band switched-resonators for VCOs," *Proceedings of IEEE Custom Integrated Circuits Conference*, pp. 225 – 228, September 2008.

[78] S-M. Yim and K. K. O, "Switched resonators and their applications in a dual-band monolithic CMOS LC-tuned VCO," *IEEE Transactions on Microwave Theory and Techniques*, vol. 54, no. 1, pp. 74 – 81, January 2006.

[79] S-M. Yim and K. K. O, "Demonstration of a switched resonator concept in a dual-band monolithic CMOS LC-tuned VCO," in *IEEE Conference on Custom Integrated Circuits*, 2001, pp. 205 – 208.

[80] M. Kossel, T. Morf, J. Weiss, P. Buchmann, C. Menolfi, T. Toifl, and M.L. Schmatz, "LC PLL with 1.2-octave locking range based on mutual-inductance switching in 45-nm SOI CMOS," *IEEE Journal of Solid-State Circuits*, vol. 44, no. 2, pp. 436 – 449, February 2009.

[81] J. J. Rael and A. Abidi, "Physical processes of phase noise in differential LC oscillators," *Proceedings of IEEE Custom Integrated Circuits Conference*, pp. 569 – 572, May 2000.

[82] D. Ham and A. Hajimiri, "Concepts and methods of optimization of integrated LC VCOs," *IEEE Journal of Solid-State Circuits*, pp. 896 – 909, June 2001.

[83] A. Hajimiri and T. H. Lee, "Design issues in CMOS differential LC oscillators," *IEEE Journal of Solid-State Circuits*, pp. 717 – 724, May 1999.

[84] J. Craninckx and M. J. Steyaert, "A 1.8-GHz low-phase-noise CMOS VCO using optimized hollow spiral inductors," *IEEE Journal of Solid-State Circuits*, pp. 736 – 744, May 1997.

[85] J. Borremans, S. Bronckers, P. Wambacq, M. Kujik, and J. Craninckx, "A single-inductor dual-band VCO in 0.06mm^2 5.6GHz multi-band front-end in 90nm digital CMOS," *Digest of Technical Papers, IEEE International Solid-State Circuits Conference*, pp. 324 – 616, Feb 2008.

[86] C. P. Yue and S. S. Wong, "On-chip spiral inductors with patterned ground shields for Si-based RF ICs," *IEEE Journal of Solid-State Circuits*, pp. 743 – 752, May 1998.

[87] Integrand Software, Inc., *EMX user's manual*, 2010.

[88] N. H. W. Fong, J. O. Plouchart, N. Zamdmer, D. Liu, L. Wagner, C. Plett, and N. G. Tarr, "Design of wide-band CMOS VCO for multiband wireless LAN applications," *IEEE Journal of Solid-State Circuits*, pp. 1333 – 1342, August 2003.

[89] A. Bevilacqua, F. P. Pavan, C. Sandner, A. Gerosa, and A. Neviani, "Transformer-based dual-mode voltage-controlled oscillators," *IEEE Transactions on Circuits and Systems II*, pp. 293 – 297, April 2007.

[90] B. Sadhu, J.O. Plouchart, S.K. Reynolds, A.V. Rylyakov, and J. Tierno, "Fully decoupled LC tank based oscillators for low phase noise and high oscillation amplitude applications," *U.S. Patent (filed)*, September 2011.

[91] D.B. Leeson, "A simple model of feedback oscillator noise spectrum," *Proceedings of the IEEE*, vol. 54, no. 2, pp. 329 – 330, February 1966.

[92] A. Hajimiri and T.H. Lee, "A general theory of phase noise in electrical oscillators," *IEEE Journal of Solid-State Circuits*, vol. 33, no. 2, pp. 179 – 194, Feb 1998.

[93] Y. Wachi, P. Nagasku, and H. Kondoh, "A 28GHz low-phase-noise CMOS VCO using an amplitude-redistribution technique," in *IEEE International Solid-State Circuits Conference*, 2008, pp. 482 – 483.

[94] B. Sadhu, M. A. Ferriss, J-O. Plouchart, A. S. Natarajan, A. V. Rylyakov, A. Valdes-Garcia, B. D. Parker, S. Reynolds, A. Babakhani, S. Yaldiz, L. Pileggi, R. Harjani, J. Tierno, and D. Friedman, "A 21.8−27.5GHz PLL in 32nm SOI using Gm linearization to achieve −130dBc/Hz phase noise at 10MHz offset from a 22GHz carrier," *IEEE Radio Frequency Integrated Circuits Symposium*, June 2012.

[95] Barry Gilbert, "Automatic biasing scheme for reducing oscillator phase noise," , no. 6064277, May 2000.

[96] B. Floyd, "A 16−18.8-GHz sub-integer-N frequency synthesizer for 60-GHz transceivers," *IEEE Journal of Solid-State Circuits*, pp. 1076 – 1086, May 2008.

[97] D. Ham and A. Hajimiri, "Concepts and methods of optimization of integrated LC VCOs," *IEEE Journal of Solid-State Circuits*, pp. 896–909, June 2001.

[98] M. Ferriss, J-O. Plouchart, A. Natarajan, A. Rylyakov, B. Parker, A. Babakhani, S. Yaldiz, B. Sadhu, A. Valdes-Garcia, J. Tierno, and D. Friedman, "An integral path self-calibration scheme for a 20.1−26.7GHz dual-loop PLL in 32nm SOI CMOS," *IEEE Symposium on VLSI Circuits*, June 2012.

[99] J-O. Plouchart, M. Ferriss, A. Natarajan, A. Valdes-Garcia, B. Sadhu, A. Rylyakov, B. Parker, M. Beakes, A. Babakani, S. Yaldiz, L. Pileggi, R. Harjani, S. Reynolds, J. Tierno, and D. Friedman, "A 23.5GHz PLL with an adaptively biased VCO in 32nm SOI-CMOS," in *IEEE Custom Integrated Circuits Conference*, September 2012.

[100] J.F. Osorio, C.S. Vaucher, B. Huff, E. v.d. Heijden, and A. de Graauw, "A 21.7-to-27.8GHz 2.6-degrees-rms 40mW frequency synthesizer in 45nm CMOS for mm-wave communication applications," in *IEEE International Solid-State Circuits Conference*, February 2011, pp. 278 – 280.

[101] O. Richard, A. Siligaris, F. Badets, C. Dehos, C. Dufis, P. Busson, P. Vincent, D. Belot, and P. Urard, "A 17.5-to-20.94GHz and 35-to-41.88GHz PLL in 65nm CMOS for wireless HD applications," in *IEEE International Solid-State Circuits Conference*, 2010, pp. 252 – 253.

[102] D. Murphy, Q.J. Gu, Y-C. Wu, H-Y. Jian, Z. Xu, A. Tang, F. Wang, and M.-C.F. Chang, "A low phase noise, wideband and compact CMOS PLL for use in a heterodyne 802.15.3c transceiver," *IEEE Journal of Solid-State Circuits*, vol. 46, no. 7, pp. 1606 – 1617, July 2011.

[103] K. Scheir, G. Vandersteen, Y. Rolain, and P. Wambacq, "A 57-to-66GHz quadrature PLL in 45nm digital CMOS," in *IEEE International Solid-State Circuits Conference*, February 2009, pp. 494 – 495, 495a.

[104] S. Pellerano, R. Mukhopadhyay, A. Ravi, J. Laskar, and Y. Palaskas, "A 39.1-to-41.6GHz $\Delta\Sigma$ fractional-N frequency synthesizer in 90nm CMOS," in *IEEE International Solid-State Circuits Conference*, February 2008, pp. 484 – 630.

[105] V.J. Arkesteijn, E.A.M. Klumperink, and B. Nauta, "Jitter requirements of the sampling clock in software radio receivers," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 53, no. 2, pp. 90 – 94, February 2006.

[106] A. Mirzaei, S. Chehrazi, R. Bagheri, and A.A. Abidi, "Analysis of first-order anti-aliasing integration sampler," *IEEE Transactions on Circuits and Systems I*, vol. 55, no. 10, pp. 2994 – 3005, November 2008.

[107] Z. Ru, E.A.M. Klumperink, and B. Nauta, "On the suitability of discrete-time receivers for Software-Defined Radio," in *IEEE International Symposium on Circuits and Systems*, May 2007, pp. 2522 – 2525.

[108] D. Jakonis, K. Folkesson, J. Dbrowski, P. Eriksson, and C. Svensson, "A 2.4-GHz RF sampling receiver front-end in 0.18-μm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 6, pp. 1265 – 1277, June 2005.

[109] R. Bagheri, A. Mirzaei, S. Chehrazi, M.E. Heidari, Minjae Lee, M. Mikhemar, Wai Tang, and A.A. Abidi, "An 800-MHz–6-GHz Software-Defined wireless receiver in 90-nm CMOS," December 2006, vol. 41, pp. 2860 – 2876.

[110] Z. Ru, E.A.M. Klumperink, and B. Nauta, "Discrete-time mixing receiver architecture for RF-sampling Software-Defined Radio," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 9, pp. 1732 – 1745, September 2010.

[111] L.R. Carley and T. Mukherjee, "High-speed low-power integrating cmos sample-and-hold amplifier architecture," May 1995, pp. 543 – 546.

[112] G. Xu and J. Yuan, "Comparison of charge sampling and voltage sampling," in *IEEE Midwest Symposium on Circuits and Systems*, 2000, vol. 1, pp. 440 – 443 vol.1.

[113] S. Karvonen, T.A.D. Riley, and J. Kostamovaara, "Charge-domain FIR sampler with programmable filtering coefficients," *IEEE Transactions on Circuits and*

Systems II, vol. 53, no. 3, pp. 192 – 196, March 2006.

[114] J. Yuan, "A charge sampling mixer with embedded filter function for wireless applications," in *International Conference on Microwave and Millimeter Wave Technology*, 2000, pp. 315 – 318.

[115] N. J. Guilar, F. Lau, P. J. Hurst, and S. H. Lewis, "A passive switched-capacitor finite-impulse-response equalizer," *IEEE Journal of Solid-State Circuits*, vol. 42, no. 2, pp. 400 – 409, February 2007.

[116] A. Abidi, "The path to the Software-Defined Radio receiver," *IEEE Journal of Solid-State Circuits*, pp. 954 – 966, May 2007.

[117] F. Montaudon, R. Mina, S. Le Tual, L. Joet, D. Saias, R. Hossain, F. Sibille, C. Corre, V. Carrat, E. Chataigner, J. Lajoinie, S. Dedieu, F. Paillardet, and E. Perea, "A scalable 2.4-to-2.7GHz Wi-Fi/WiMAX discrete-time receiver in 65nm CMOS," in *Solid-State Circuits Conference, 2008. ISSCC 2008. Digest of Technical Papers. IEEE International*, February 2008, pp. 362 – 619.

[118] F. Harris, Chris Dick, and Michael Rice, "Digital receivers and transmitters using polyphase filter banks for wireless communications," *IEEE Transactions on Microwave Theory and Techniques*, pp. 1395 – 1412, April 2003.

[119] F. Rivet, Y. Deval, J.-B. Begueret, D. Dallet, P. Cathelin, and D. Belot, "The experimental demonstration of a SASP-based full software radio receiver," *IEEE*

Journal of Solid-State Circuits, pp. 979 – 988, May 2010.

[120] M. Lehne and S. Raman, "A 0.13-μm 1-GS/s CMOS discrete-time FFT processor for ultra-wideband OFDM wireless receivers," *IEEE Transactions on Microwave Theory and Techniques*, pp. 1639 – 1650, November 2000.

[121] B. Sadhu, M. Sturm, B. M. Sadler, and R. Harjani, "A 5GS/s 12.2pJ/conv. analog charge-domain FFT for a software defined radio receiver front-end in 65nm CMOS," *IEEE Radio Frequency Integrated Circuits Symposium*, June 2012.

[122] R. Harjani, B. Sadhu, and M. Sturm, "Multi-stage charge re-use analog circuits," *U.S. Patent (filed)*, June 2012.

[123] B. Murmann, "EE315B: VLSI Data Conversion Circuits," *Stanford University*, 2011.

[124] Y. Ding and R. Harjani, "A universal analytic charge injection model," in *IEEE International Symposium for Circuits and Systems*, 2000, vol. 1, pp. 144 – 147.

[125] S. Patnaik, S. Kalia, B. Sadhu, M. Sturm, M. Elbadry, and R. Harjani, "An 8GHz multi-beam spatio-spectral beamforming receiver using an all-passive discrete time analog baseband in 65nm CMOS," in *IEEE Custom Integrated Circuits Conference*, September 2012.

[126] Y.-W. Lin, H.-Y. Liu, and C.-Y. Lee, "A 1-GS/s FFT/IFFT processor for UWB applications," *IEEE Journal of Solid-State Circuits*, pp. 1726 – 1735, August 2005.

[127] A. Tkachenko, A.D. Cabric, and R.W. Brodersen, "Cyclostationary feature detector experiments using reconfigurable BEE2," in *IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, April 2007, pp. 216 – 219.

[128] B. Sadhu and R. Harjani, "Capacitor bank design for wide tuning range LC VCOs: 850MHz-7.1GHz (157%)," *Proceedings of IEEE International Symposium on Circuits and Systems*, pp. 1975 – 1978, May 2010.

[129] A. Goel and H. Hashemi, "Frequency switching in dual-resonance oscillators," *IEEE Journal of Solid-State Circuits*, pp. 571– 582, March 1999.

Appendix A

SCF

A popular method to extract cyclostationary features is to implement the spectral correlation function (SCF) [64]. In cyclostationary feature detectors, the incoming signal is modeled as a cyclostationary random process with multiple periodicities, as compared to the traditional wide sense stationary model for energy detection where the cyclostationary information is discarded. Then, the autocorrelation function for a zero-mean process $x(t)$ is defined as

$$R_X\left(t+\frac{\tau}{2}, t-\frac{\tau}{2}\right) = E\left[x\left(t+\frac{\tau}{2}\right)x^*\left(t-\frac{\tau}{2}\right)\right]$$

which, in the 't' domain, exhibits the multiple periodicities of the incoming signal. The Fourier transform of the autocorrelation over 't' as expressed in

$$R_X^{\alpha}(\tau) = \lim_{Z\to\infty} \frac{1}{Z} \int_{-\frac{Z}{2}}^{\frac{Z}{2}} R_X\left(t+\frac{\tau}{2}, t-\frac{\tau}{2}\right) e^{-j2\pi\alpha t} \mathrm{d}t$$

is called the cyclic autocorrelation in the cycle frequency (α) domain, and is expected to capture these periodicities. For a cycloergodic process, the definition simplifies to

$$R_X^{\alpha}(\tau) = \lim_{Z\to\infty} \frac{1}{Z} \int_{-\frac{Z}{2}}^{\frac{Z}{2}} x\left(t+\frac{\tau}{2}\right) x^*\left(t-\frac{\tau}{2}\right) e^{-j2\pi\alpha t} \mathrm{d}t$$

Then, using the cyclic Wiener-Khinchin relation on $R_X^{\alpha}(\tau)$, the spectral correlation function (SCF) can then be defined as

$$S_X^{\alpha}(f) = \int_{-\infty}^{\infty} R_X^{\alpha}(\tau) e^{-j2\pi f\tau} \mathrm{d}\tau$$

Here, we note that the cyclic autocorrelation function, $R_X^{\alpha}(\tau)$, and the SCF, $S_X^{\alpha}(f)$, reduce to the conventional definitions of autocorrelation function $R_X(\tau)$, and the power spectral density, $S_X(f)$, for the case when $\alpha = 0$.

Again, an alternative, and more convenient equivalent definition of the SCF can be expressed in terms of the short time Fourier transform (STFT) of the received signal $x(t)$. The STFT of a signal is just the Fourier transform at the particular instance 't', defined as

$$X_T(t,f) = \int_{t-\frac{T}{2}}^{t+\frac{T}{2}} x(u) e^{-j2\pi fu} \mathrm{d}u$$

The SCF can then be defined in terms of the STFT as

$$S_X^{\alpha}(f) = \lim_{T\to\infty} \lim_{\Delta t\to\infty} \frac{1}{\Delta t} \int_{\frac{-\Delta t}{2}}^{\frac{\Delta t}{2}} \frac{1}{T} X_T\left(t, f+\frac{\alpha}{2}\right) X_T^*\left(t, f-\frac{\alpha}{2}\right) \mathrm{d}t$$

thus better justifying its nomenclature. As is evident, this second definition lends itself to an easier execution in hardware and has been used in hardware implementations of cyclostationary feature detectors [127]. For discrete time implementation, the discrete and finite time equivalent definition of the SCF is derived as

$$\hat{S}_X^{\alpha}(f) = \frac{1}{N}\frac{1}{T} \sum_{n=0}^{N} X_{T_{DFT}}\left(n, f+\frac{\alpha}{2}\right) X_{T_{DFT}}^*\left(n, f-\frac{\alpha}{2}\right)$$

where $X_{T_{DFT}}(n, f)$ represents the N point DFT around sample n.

The superiority of cyclostationary feature detectors over energy detection can be visualized in Fig. A.1, where the power spectral density (PSD) and SCF of a BPSK modulated signal are simulated in MATLAB®. The PSD, as would be calculated by an energy detector is shown on the left. The energy detection technique is based on comparing the signal energy with a threshold that is dependent on an estimation of the noise power. As evident in the Fig. A.1(a), any estimation error due to interference or changing noise variance could easily change the absolute energy detected and cause errors, especially in the case of weak signals. Also, the PSD cannot discriminate between signal energy and noise/interferer energy. However, when the SCF is calculated as in Figure 3(b), the additive white Gaussian noise appears only along the main diagonal,

which represents the PSD as estimated by energy detectors. Since the cross-correlation of the noise tends to zero for long observation times, the BPSK signal is clearly visible along the 'α' axis (second diagonal). Moreover, the energy along the 'α' axis is independent of the noise variance.

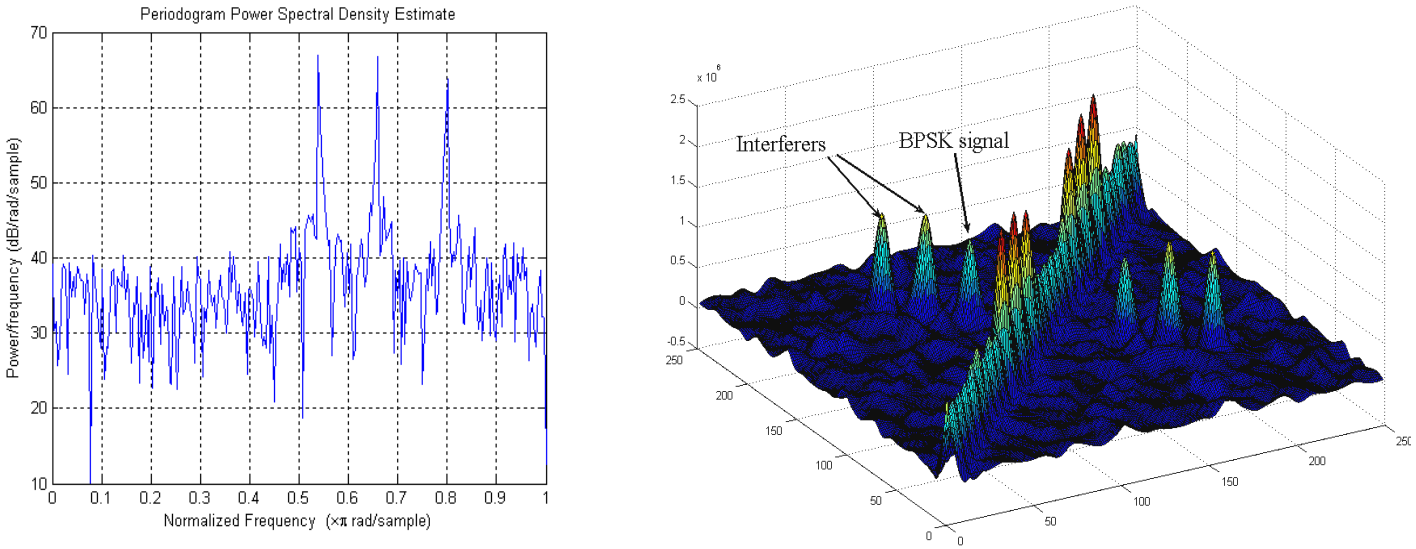

(a) (b)

Figure A.1: (a) PSD of a BPSK signal with -60dB SNR and two interferers, and (b) SCF of the same signal, where the diagonal represents the PSD along the f axis as shown in (a)

Appendix B

Further VCO Capacitor Array Optimization

The VCO tuning range can be further increased by optimizing the switched capacitor array used in the VCO implementation [128]. A new design based on these techniques has been demonstrated via extracted simulation. In particular, two novel techniques are proposed: capacitor array switch sizing discussed in Section 3.5.2, and the use of interconnect inductance for capacitance boosting, discussed below. Post-layout simulation results from Cadence® and ADS Momentum® for a design based on these principles are reported. Based on the close matching between simulation results and measurement data for previous versions of this design discussed above, the simulation results being reported for this design are expected to be a good predictor.

B.0.3 Capacitance Boosting

Interconnect inductance, that is usually viewed as a parasitic and inevitable nuisance in LC oscillator circuits, can be used to further increase the tuning range and maintain tuning monotonicity. To understand this, let us consider the reactance looking into the branch shown in Fig. B.1. Calculating the reactance as shown in (B.1) we note that the magnitude of X is reduced by the interconnect parasitic inductance. Therefore, the effective capacitance can now be written as shown in (B.2).

$$X = -\frac{1}{\omega C} + \omega(2L_{int}) \tag{B.1}$$

$$C_{eff} = -\frac{1}{\omega X} = \frac{C}{1 - 2\omega^2 L_{int} C} \tag{B.2}$$

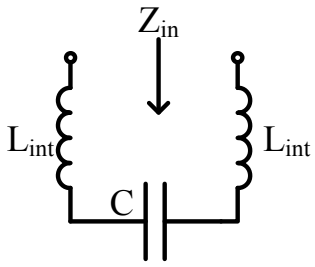

Figure B.1: Model of the capacitor bank with inductive interconnects

The capacitance is therefore effectively magnified using the parasitic interconnect inductance. Also, because of the relative magnitudes, the parasitic capacitance when the capacitor bit is switched **off**, is hardly affected by the interconnect inductance. This can be verified by using a small value of C in (B.2). This feature can be used to increase the tuning range as well as to make the tuning characteristic monotonic. To achieve the

latter, longer interconnects are used for the larger capacitor bits (Fig. 3.24 and B.2) to compensate for the smaller switches used.

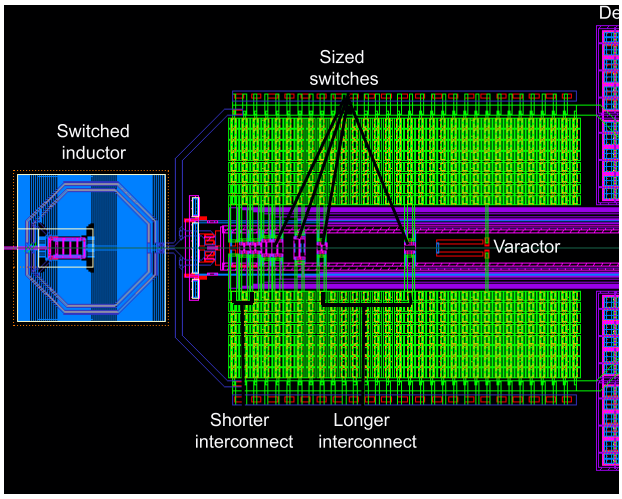

Figure B.2: Layout of the capacitor bank with optimally sized switches

Note that these interconnects are an inevitable part of the layout. In the capacitance boosting scheme, we merely model the interconnects accurately, and manipulate the layout cleverly in order to obtain a performance benefit from this parasitic inductance.

It is important to also recognize that this interconnect inductance gives rise to a higher order resonator that can have multiple modes of oscillation [129]. The resonator should be designed carefully to ensure that these higher order modes are maintained at a much higher frequency such that the startup criteria described in (3.2) fails for these parasitic oscillation modes.

B.0.4 Simulation Results

The previously designed cross-coupled LC oscillator (Fig. 3.5) was updated with the optimized capacitor array and simulated using IBM's 0.13μm CMOS process in Cadence®. Considering the close match between the simulation and measurement results from the fabricated design described earlier, the simulation results discussed below are expected to be a good predictor, with the exception of the unsimulated patterned ground shield caveat discussed below.

Tuning Range

The frequency tuning range (FTR) obtained from simulation spans 6.21GHz from 850MHz to 7.06GHz (157%) as shown in Fig. B.3. The two frequency bands are made to overlap slightly (450MHz overlap on the y-axis) to ensure continuous frequency coverage in the face of process variations. The FTR decreases to 144% (1GHz to 6.2GHz) post RC extraction of the capacitor bank (the inductor is simulated using ADS Momentum®).

This design may be compared to the previous fabricated version [35] described above in a similar process technology which achieved 87% tuning range. For the present design, the design methodology used for the switched-inductor and g_m cells is similar to the previous design. However, optimization of the capacitor bank as described in this section increased the tuning range by an additional 57%. Interestingly, apart from a slight increase in design complexity, there is no price paid for this increased tuning range.

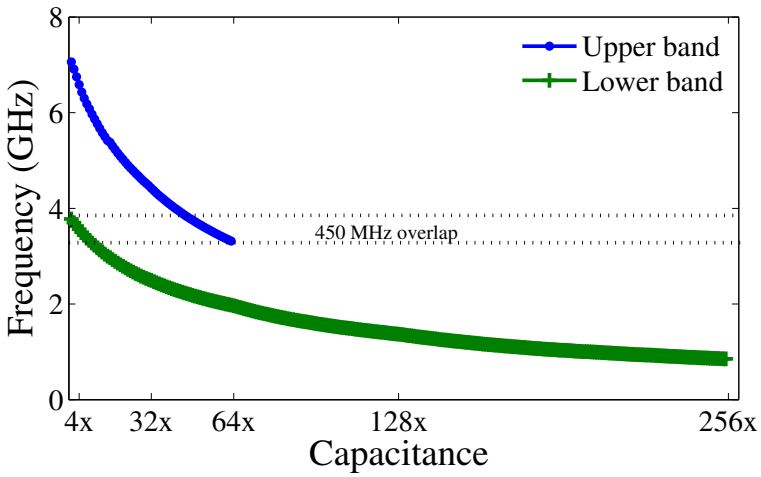

Figure B.3: Frequency tuning range versus capacitance in terms of the 'x', where 'x' is the capacitance of a unit switched capacitor

Phase Noise

The variation in phase noise over the tuning range is shown in Fig. B.4.

The phase noise performance versus frequency is at 1MHz offest is shown in Fig. B.4. As seen in the figure, when a new capacitor is switched in, the phase noise remains approximately constant due to the scheme discussed in Section 3.5.2. This manual tuning accounts for a slight overhead in design as compared to the previous design version.

The phase noise is seen to vary between -119.1 and -107.1 dBc/Hz. The degradation in phase noise performance as compared to the fabricated design described above can be attributed to a lower Q inductor used in simulation. The peak inductor Q for the inductor used in this design was found to be 16 in simulation. A higher inductor Q (≈ 22) can be obtained in this process by using a high resistance substrate and patterned

ground shields, improving the phase noise considerably ($\approx 20log(Q_1/Q_2)^2 \approx 6dB$). This improvement had been verified in the previous version of the design [35] but could not be verified using electromagnetic simulations due to the large memory requirements for simulating patterned ground shields. It is expected that the measured phase noise, and consequently the FOM from fabricated chips are likely to have a similar performance as the Design 1 in Table 3.3.

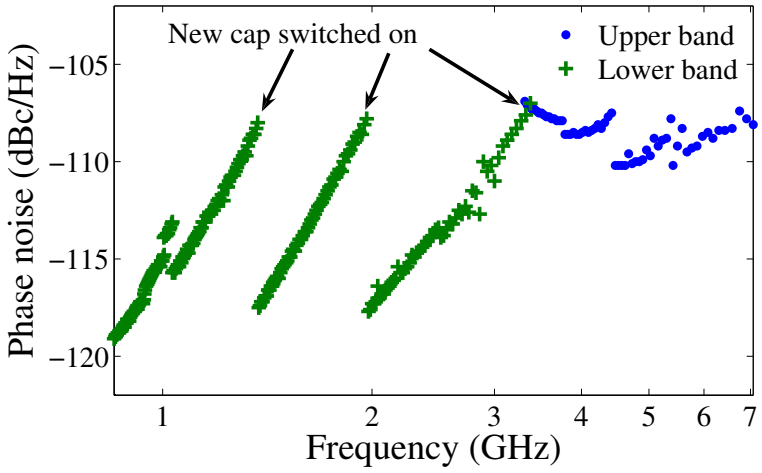

Figure B.4: Variation of phase noise at 1MHz offset with frequency (log scale)

Power

The variation in power is shown in Fig. B.5 with a linear fit in each frequency band. Similar to the technique described in Section 3.4.4, the power dissipation is limited to a maximum of 15mW through g_m cell sizing to improve the tuning range. The overall trends are as expected from the analysis in Section 3.2.

A comparison of this work with previous wide-tuning range single LC tank designs is shown in Table B.1. As seen from the comparison, this technique provides much larger tuning ranges than are currently available in the state of the art designs. Also, good phase noise and power dissipation performance is maintained over the entire tuning range.

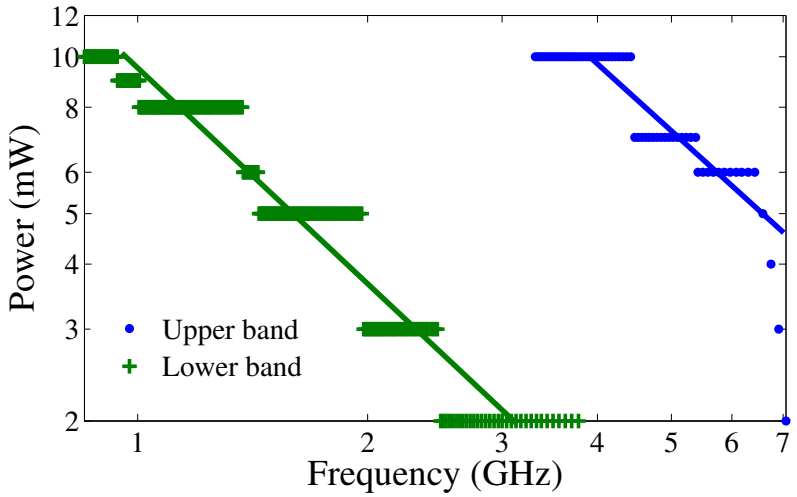

Figure B.5: Variation of VCO core power dissipation with frequency (log log plot)

Table B.1: VCO Performance Comparison

Ref	f_{min} tof_{max} (GHz)	FTR (%)	Phase Noise* (dBc/Hz)@1MHz	FOM_{PFTN} (dB) [82]	Implementation (μm)
[82]	2.0 to 2.6	26	−125.4 to −119.4	−3.1	0.35 BiCMOS
[88]	3.1 to 5.6	58.7	−120.8 to −114.6	5.9 to 10.3	0.13 SOI
[40]	1.1 to 2.5	73	−126.5 (f_0)	5.0 to 8.5	0.18 CMOS
[89]	3.6 to 8.4	74	−104 to −101.5	−4.6 to 4.0	0.13 CMOS
This (simulated)	**1 to 6.2**	**144**	**−119.1 to −107.1**	**1.1 to 15**	**0.13 CMOS**

(*Assuming 20dB/decade drop with offset frequency)

Appendix C

CRAFT matrices

The CRAFT matrix equation (6.1) is written as $\mathbf{X} = \frac{\mathbf{1}}{\mathbf{16}} \cdot \mathbf{S_4 S_3 S_2 S_1 I}_{\text{bitrev}}\mathbf{x}$.

$\mathbf{I}_{\text{bitrev}}, \mathbf{S_1}, \mathbf{S_2}, \mathbf{S_3}, \mathbf{S_4}$ are shown below:

$$\mathbf{I}_{\text{bitrev}} = \begin{bmatrix}
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{bmatrix}$$

$$\mathbf{S_1} = \begin{bmatrix}
1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & W^4 & -W^4 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & W^4 & -W^4 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W^4 & -W^4 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W^4 & -W^4
\end{bmatrix}$$

$$\mathbf{S_2} = \frac{1}{\sqrt{2}} \begin{bmatrix}
1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & W^0 & 0 & W^0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & W^2 & 0 & W^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & W^4 & 0 & -W^4 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & W^6 & 0 & -W^6 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & -1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W^0 & 0 & W^0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W^2 & 0 & W^2 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W^4 & 0 & -W^4 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W^6 & 0 & -W^6
\end{bmatrix}$$

$$\mathbf{S_3} = \left(\frac{1}{2 \cdot \cos\left(\frac{\pi}{8}\right)}\right) \begin{bmatrix}
1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W^0 & 0 & 0 & 0 & W^0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W^1 & 0 & 0 & 0 & W^1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W^2 & 0 & 0 & 0 & W^2 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W^3 & 0 & 0 & 0 & W^3 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W^4 & 0 & 0 & 0 & -W^4 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W^5 & 0 & 0 & 0 & -W^5 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W^6 & 0 & 0 & 0 & -W^6 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W^7 & 0 & 0 & 0 & -W^7
\end{bmatrix}$$

$$\mathbf{S_4} = \begin{bmatrix}
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1
\end{bmatrix}$$

Appendix D

Acronyms

Care has been taken in this thesis to minimize the use of jargon and acronyms, but this cannot always be achieved. This appendix defines jargon terms in a table of acronyms and their expansion.

D.1 Acronyms

Table D.1: Acronyms

Acronym	Meaning
2G	Second Generation (of wireless communications)
3G	Third Generation (of wireless communications)
4G	Third Generation (of wireless communications)

Continued on next page

Table D.1 – continued from previous page

Acronym	Expansion
A/D	Analog to Digital
AC	Alternating Current
ADC	Analog to Digital Converter
AGC	Automatic Gain Control
ASIC	Application-Specific Integrated Circuit
ASP	Analog Signal Processing
AWGN	Additive White Gaussian Noise
BB	BaseBand
BER	Bit Error Rate
BJT	Bipolar Junction Transistor
BPF	Band Pass Filter
BPSK	Binary Phase Shift Keying
BW	BandWidth
CDMA	Code Division Multiple Access
CMOS	Complementary Metal Oxide Semiconductor

Continued on next page

Table D.1 – continued from previous page

Acronym	Expansion
CR	Cognitive Radio
CRAFT	Charge Re-use Analog Fourier Transform
CRAFT	Charge Reuse Analog Fourier Transform
CT	Continuous-Time
D/A	Digital to Analog
DAC	Digital to Analog Converter
dB	Decibel
dBc	Decibel normalized to carrier
dBFS	Decibel normalized to Full Scale
dBm	Decibel normalized to 1mW
DC	Direct Current
DFT	Discrete Fourier Transform
DNL	Differential Non-Linearity
DR	Dynamic Range
DSB	Double Side-Band

Continued on next page

Table D.1 – continued from previous page

Acronym	Expansion
DSP	Digital Signal Processing
DT	Discrete-Time
EM	Electro-Magnetic
EMI	Electromagnetic Interference
EW	Electronic Warfare
FDD	Frequency Division Duplexing
FET	Field Effect Transistor
FFT	Fast Fourier Transform
FIR	Finite Impulse Response
FPGA	Field-Programmable Gate Array
FSK	Frequency Shift Keying
GPRS	General Packet Radio Service
GPS	Global Positioning System
GSM	Global System for Mobile Communications
HPF	High Pass Filter

Continued on next page

Table D.1 – continued from previous page

Acronym	Expansion
I/O	Input/Output
I−Q	In phase − Quadrature phase
IC	Integrated Circuit
IF	Intermediate Frequency
IIR	Infinite Impulse Response
ILFD	Injection Locked Frequency Divider
ILFM	Injection Locked Frequency Multiplier
ILO	Injection Locked Oscillator
IM	Intermodulation Frequency
INL	Integral Non-Linearity
LC	Inductor-Capacitor
LiTVCO	Linear Transconductance VCO
LNA	Low Noise Amplifier
LO	Local Oscillator
LoS	Line of Sight

Continued on next page

Table D.1 – continued from previous page

Acronym	Expansion
LPF	Low Pass Filter
LTE	Long Term Evolution
MOSFET	Metal Oxide Semiconductor Field Effect Transistor
OFDM	Orthogonal Frequency Division Multiplexing
OpAmp	Operational Amplifier
OTA	Operational Transconductance Amplifier
PA	Power Amplifier
PCB	Printed Circuit Board
PLL	Phase Locked Loop
PSK	Phase Shift Keying
Q	Quality factor
QAM	Quadrature Amplitude Modulation
QPSK	Quadrature Phase Shift Keying
RC	Resistor-Capacitor
RF	Radio-Frequency

Continued on next page

Table D.1 – continued from previous page

Acronym	Expansion
Rx	Receiver
SAR	Successive Approximation Register
SASP	Sampled Analog Signal Processor
SAW	Surface Acoustic Wave
SCF	Spectral Correlation Function
SDR	Software Defined Radio
SFDR	Spurious Free Dynamic Range
SigInt	Signal Intelligence
SNDR	Signal to Noise and Distortion Ratio
SNR	Signal to Noise Ratio
SR	Software Radio
SSB	Single Side-Band
TDD	Time Division Duplexing
TV	Television
Tx	Transmitter

Continued on next page

Table D.1 – continued from previous page

Acronym	Expansion
VCO	Voltage Controlled Oscillator
VGA	Variable Gain Amplifier
VLSI	Very Large Scale Integration
WiMAX	Worldwide Interoperability for Microwave Access
WLAN	Wireless Local Area Network