

Low-Power RFIC Design Techniques for Self-Powered Wireless CMOS Circuits with Integrated Antennas

by

Peter Harris Robert Popplewell, B.Eng., M.A.Sc.

A thesis submitted to the
Faculty of Graduate Studies and Research
in partial fulfillment of the requirements
for the Degree of

Doctor of Philosophy

in

Electrical Engineering

Ottawa-Carleton Institute for Electrical and Computer Engineering

Department of Electronics

Faculty of Engineering

Carleton University

Ottawa, Canada

May, 2010

© Copyright, 2010

Peter Harris Robert Popplewell



Library and Archives
Canada

Published Heritage
Branch

395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque et
Archives Canada

Direction du
Patrimoine de l'édition

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file *Votre référence*
ISBN: 978-0-494-67899-2
Our file *Notre référence*
ISBN: 978-0-494-67899-2

NOTICE:

The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

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

AVIS:

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

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

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.

■ ■ ■
Canada

Abstract

This thesis focuses on system level design strategies and techniques for circuit level implementations that facilitate low-power RFICs in bulk CMOS. Ultimately, the goal of this work is to enable the design of inexpensive and completely integrated circuits that consume so little power that they can be self-powered while communicating by means of an integrated antenna. As an application example, the design and implementation of a unique low-power integrated FM receiver is presented. The receiver is a completely new topology, using a modified PLL operated in both open-loop and closed-loop configurations, and using oscillator injection locking to accomplish FM demodulation with a minimum of circuitry. The receiver communicates at 5.2 GHz while consuming 285 μ W when duty cycled in a typical application.

The receiver represents one half of a collaborative research project which developed a novel integrated transceiver suitable for short range wireless applications such as RFID tagging or the transmission of data from medical sensors. The circuit is unique in that it is virtually completely integrated, optionally making use of an on-chip antenna, and has such low power consumption that it could be self-powered by a thin film ultracapacitor and solar cell stacked on top of the chip. Both the transmitter and receiver consist of PLLs which initially phase lock VCOs, and then allow them to “roll” in order to transmit and receive the signal. The VCO in the receiver is injection locked by the incoming signal. The current design has a communication range of 6.5 cm when integrated antennas are used for both ends of the link, which can be increased at the expense of the data rate or increased power consumption in the receiver. When one end of the communication link uses a 6.7 dBi off-chip patch antenna, the communication range increases to 1.75 m.

The appropriate background theory and calculations necessary to understand the design of the circuits are presented, along with the details of the circuits themselves and their simulated and measured behaviours. A brief discussion on the design and behaviour of the transmitter circuit is also included, as this discussion fosters understanding of the receiver design and the novel transceiver topology.

Acknowledgments

I would like to thank all those who have helped me to achieve my academic goals, whether you have helped me financially, technically, or emotionally, I am grateful.

I am pleased to acknowledge financial assistance received in the form of scholarships from the Natural Sciences and Engineering Research Council, The IEEE Solid-State Circuits Society, and from Skyworks Solutions Inc.

To Calvin Plett and John Rogers, my academic supervisors, I thank you for your guidance and patience. You are both true researchers at heart, eager to participate in all technical discussions, and have shown great patience with me when I was forced to postpone my academic advancement from time to time because of work or family priorities.

To Victor Karam and Atif Shamim, I am thrilled that we were able to combine our research efforts together. We've come a long way from our early day brainstorming sessions, and I know no finer gentlemen with whom I'd rather conduct research.

To Justin Fortier, Chris DeVries, Eugene Ivanov, Christina Young, Doug Beards and the rest of my former colleagues at Kleer Semiconductor (now SMSC), our day to day discussions and technical victories provided me with a technical foundation that cannot be obtained from any textbook. I am grateful to know and have worked with all of you.

To Florin Balteanu and my colleagues at Skyworks Solutions, thank you for supporting me during the final stages of my campaign to complete this thesis, and for tolerating my split focus.

My family is always supportive and I owe thanks to my parents and my brother for my success in both academic and personal life.

To my three year old son Luke, you are always so incredibly happy and positive that merely your presence in my life is a huge emotional boost and I look forward to having more spare time to spend with you now.

To my amazing daughter Alice, you are only one year old now and are already my hero.

Finally, I thank my wife Vicki who is always beside me, supportive, patient, and willing to defer much for the sake of my academic pursuits.

Contents

Abstract	ii
Acknowledgments	iii
List of Figures	ix
List of Tables	xii
List of Abbreviations and Symbols	xiii
1 Introduction	1
1.1 Motivation	1
1.2 Thesis Objectives	3
1.3 Thesis Organization	5
2 Background	8
2.1 Recognizing the Tradeoffs in Low-Power RFICs	8
2.2 Applying the Tradeoffs – Receiver Examples from Literature	10
2.2.1 Low-Power at the Expense of Cost, Size and Integration	10
2.2.2 High Frequency and Power with the Benefit of Integrated Antennas	11
2.2.3 Minimizing the BOM while Balancing other Metrics	12
2.3 Circuit Topologies Using Oscillators as Gain Elements	15
2.3.1 The Theoretical Transfer Function of an Oscillator	15
2.3.2 Q - Enhanced Filters	17
2.3.3 The Super-Regenerative Receiver	18
2.3.4 Injection-Locked and PLL Based Receivers and Transmitters	19
2.4 On-Chip Antennas	22
2.4.1 Coupling Inductors – a Convenient Accident	22

2.4.2	Designs on High-Resistivity Substrates	22
2.5	Generating and Storing Power On Chip	23
2.5.1	Thin Film Ultracapacitors	23
2.5.2	Thin Film Photocells, Thermogenerators, Inductive and RF Power Transmission	23
2.6	Small-Form-Factor Crystals and On-Chip Self-Referenced <i>LC</i> Clocks	25
2.7	Background Summary	26
3	The Proposed Lock-and-Roll Transceiver	28
3.1	The Lock-and-Roll Transmitter	29
3.1.1	TX Overview	29
3.1.2	TX Power Consumption	32
3.2	The On-Chip Inductive Antenna	32
3.2.1	Rectangular Antenna/Inductor Equivalent Circuit Model . . .	34
3.2.2	Antenna Efficiency	35
3.3	System Link Budget and the Friis Equation	36
3.4	The Lock-and-Roll Receiver	37
3.4.1	RX Overview	37
3.4.2	RX Power Consumption	39
3.4.3	PLL Loop Component Selection	40
3.4.4	RX VCO's Injection Locking Bandwidth	41
3.5	FM Modulation Considerations	42
3.6	Tradeoffs Between Power, Range, and Data Rate	43
3.7	Lock-and-Roll Transceiver Summary	44
4	Injection-Lockable VCO Design for Low-Power Applications . . .	47
4.1	A Review of Oscillator Design Fundamentals	49
4.1.1	The Barkhausen Criteria and Gain Margin	49
4.1.2	Current-Limited vs. Voltage-Limited Oscillator Designs	50
4.1.3	Designing for Minimal MOSFET Noise Contribution	51
4.1.4	<i>LC</i> Resonance and Tank <i>Q</i>	52
4.1.5	Phase Noise, Oscillator Pulling, and Injection Locking Bandwidth	53
4.1.6	Oscillator Fundamentals Summary	56
4.2	VCO Design for the Lock-and-Roll Receiver	57

4.2.1	Advantages of the Complementary Differential Topology . . .	58
4.2.2	Inductor Selection – Minimizing Process Variation Regardless of Q	60
4.2.3	Transistor Sizing – Trading-Off High Gain Margin and Increased Tunability for Higher Phase Noise	61
4.2.4	Tank Circuit and Tunability	62
4.2.5	Optimizing for Injection Locking Bandwidth and Low Power .	63
4.2.6	Designing a VCO with Margin – VCO Layout with Laser Cut Options	64
4.2.7	Simulated and Measured Results	66
4.3	VCO Design Summary	71
5	Injection-Locking Circuit Design	74
5.1	Coupling Voltage vs. Steering Current	75
5.1.1	Differential vs. Pseudo-Differential, Cascode Topologies and Tail Currents	76
5.1.2	Optimizing for Efficiency with Hard Switched Inputs	77
5.1.3	Minimizing Disruption on the VCO Core	78
5.1.4	Designing for Low Voltage Supply and Weak Input Swing . . .	79
5.2	The Lock-and-Roll Receiver LNA	79
5.2.1	Circuit Topology and Design	79
5.2.2	Trading-Off Noise, Current, Transistor Size and Output Impedance	81
5.2.3	Simulated LNA Performance	83
5.2.4	LNA Output Impedance and the Effect on the VCO	84
5.2.5	Low- Q On-Chip Input Match Design	84
5.2.6	Match Variation over Process and Temperature	86
5.2.7	LNA Circuit Layout Including Input Match	88
5.2.8	Measured Results	90
5.3	Injection-Locking Circuits Summary	93
6	PLL Component Designs that Enable Open-Loop Operation . . .	95
6.1	Highly Adjustable Loop Filter Design	98
6.1.1	Achieving a Balance Between Fast Acquisition, Increased Sta- bility, and Leakage Robustness	99

6.1.2	Loop Filter Layout with Laser Trim Tunability	102
6.2	Loop Filter Switch Design	104
6.2.1	Charge Injection and Mitigating the Effects with Dummies . . .	104
6.3	PFD Circuit Design and Behaviour	106
6.3.1	Trading-Off Lower Acquisition Time for Reduced Loop Filter Leakage	107
6.3.2	PFD Dead-zone	108
6.4	High-Impedance Charge Pump Design	109
6.4.1	Matching Up/Down Pump Profiles, Simulated Output Current Response	109
6.4.2	Minimizing Charge Injection and Leakage Through Design and Layout	112
6.4.3	Measured Output Current Response	114
6.5	Unity-Gain Loop-Filter Buffer Design	115
6.6	Simulated vs. Measured VCO Drift Rates	116
6.7	Suspected Leakage from Tie-Down Diodes	119
6.8	PLL Components that Enable Open-Loop Operation Summary . . .	120
7	The Lock-and-Roll Receiver Test Chip	123
7.1	High-Speed, Low-Power Divider Design with TSPC Input Stages . . .	126
7.2	Digital Up/Down Pulse Mux Design and Use	129
7.3	Buffer Circuits that Enable Testability	130
7.3.1	VCO Output Buffer Design	130
7.3.2	Buffering to Interface with the 50 Ω Domain	131
7.3.3	Output Bitstream Buffer Design	132
7.4	Measurement Methodology	132
7.4.1	Probe De-embedding Structures, Arrangement and Use	134
7.4.2	Enabling Chip-on-Board System Level Testing	136
7.5	Measured Receiver Output	139
7.5.1	Integrating Capacitor Size and Noise on the Final Output . .	140
7.5.2	Data Rate Limitations Revisited	143
7.6	Measured Receiver Power Consumption Breakdown	146
7.7	Comparing the Lock-and-Roll RX to State-of-the-Art Alternatives . .	147

7.8	Lock-and-Roll Receiver Test Chip Summary	152
8	Conclusion	155
8.1	Thesis Contributions	157
8.2	Publications and Major Recognition/Awards Resulting from this Work	159
8.3	Future Work	161
Appendix A Oscillator Design Fundamentals – Supplementary Infor-		
	mation	163
A.1	Barkhausen Criteria Derivation using the Linear Model	163
A.2	The $-G_m$ Oscillator and Gain Margin	164
	A.2.1 Transistor Operating Point and the Effect on Output Impedance	168
A.3	MOSFET Noise Theory	169
	A.3.1 CMOS Thermal Noise	169
	A.3.2 CMOS Shot Noise	171
	A.3.3 CMOS Flicker Noise	172
A.4	LC Resonance, Unloaded and Loaded Tank Q	173
	A.4.1 Parallel LC Resonance	173
	A.4.2 Unloaded Q Factor	174
	A.4.3 Loaded Q Factor	176
	References	177

List of Figures

1.1	Status Quo Wired Dosimeters in Use	3
1.2	Patient Awaiting Treatment using Wired Dosimeters	4
2.1	The RFIC Design Star of Tradeoffs	10
2.2	Chu's Low-Frequency, Low-Power FSK Receiver at 176 kHz	11
2.3	Bergveld's Low-IF Receiver Topology at 2.4 GHz	13
2.4	Oscillator Frequency Response	16
2.5	DeVries' Q -Enhanced Filter with Digital Tuning	17
2.6	DeVries' RF Back End with Q -Enhanced Filter	18
2.7	Plessey's Injection-Locked FM Demodulator	20
2.8	CIT's PLL-Based FM Demodulator	21
3.1	Lock-and-Roll TX Topology	29
3.2	Lock-and-Roll TX FM Signal Generation	31
3.3	Single Turn Inductors/Antennas	33
3.4	Antenna Inductances and Q	33
3.5	Antenna/Inductor Lumped Element Equivalent Circuit	35
3.6	Communication Range vs. Antenna Configurations	37
3.7	Lock-and-Roll RX Topology	38
3.8	Coupled LNA/VCO Locking Bandwidth Verification	42
3.9	Secondary CP Output vs. Phase Difference at the PFD	43
3.10	Data Rate, Locking Bandwidth, and Communication Range Tradeoff	45
4.1	NMOS, $-G_m$, LC Oscillator	51
4.2	9.94 MHz Oscillator with Weak 9.82 MHz Injection	54
4.3	9.94 MHz Oscillator with Strong 9.82 MHz Injection	55
4.4	9.94 MHz Oscillator Injection-Locked at 9.82 MHz	55
4.5	Free-Running and Injection-Locked Oscillator Phase Noise	57

4.6	Lock-and-Roll Receiver VCO Schematic	59
4.7	Lock-and-Roll Receiver VCO Layout	65
4.8	Simulated VCO+LNA Tuning Range using Coupled Extraction . . .	67
4.9	Simulated vs. Measured Tuning Range	69
4.10	Simulated Extracted VCO+LNA Locking Bandwidth Check	70
4.11	Measured VCO Spectrum, Injection Locked to Modulated RX Input .	71
5.1	Locking Circuit Schematic with Coupled Output Voltage	75
5.2	Locking Circuit Schematic using Current-Steering Approach	76
5.3	Locking Circuit Schematic for the Lock-and-Roll Receiver	80
5.4	LNA to Antenna Match Translation on a Smith Chart	86
5.5	S11 Simulation with Extracted LNA over Process and Temperature .	88
5.6	LNA Layout Including Input Match	89
5.7	Measured vs. Simulated Differential Input Impedance	92
6.1	Lock-and-Roll RX Components Enabling Open-Loop Mode	97
6.2	Lock-and-Roll RX Tunable Loop Filter Schematic	100
6.3	Lock-and-Roll RX Tunable Loop Filter Layout	103
6.4	Lock-and-Roll RX Loop Switch Schematic	104
6.5	Channel Charge Injection Effect in NMOS	105
6.6	Lock-and-Roll RX Tristate PFD Schematic	107
6.7	Lock-and-Roll RX High-Impedance Charge Pump Schematic	109
6.8	Simulated Charge Pump Output Current vs. V_{CNTL}	111
6.9	Lock-and-Roll RX High-Impedance Charge Pump Layout	113
6.10	Measured Charge Pump Output Current vs. V_{CNTL}	114
6.11	Lock-and-Roll RX Unity-Gain Loop-Filter Buffer Schematic	116
6.12	Lock-and-Roll RX Measured Drift and Point of Failure	118
7.1	Lock-and-Roll RX Test Chip Block Diagram	124
7.2	Lock-and-Roll RX Die Microphotograph	125
7.3	Lock-and-Roll RX Divider Schematic	128
7.4	Lock-and-Roll RX Up/Down Mux Schematic	129
7.5	Lock-and-Roll RX VCO Buffer Schematic	131
7.6	Lock-and-Roll RX 50 Ω Output Buffer Schematic	132

7.7	Receiver Test Chip Probe De-Embedding Options	135
7.8	Receiver Test Chip Bonded to PCB	137
7.9	Measured Receiver Output Signal with 1 kb/s Data Rate	139
7.10	Measured Output Signal Noise Period	140
7.11	The Cycle Slip Phenomenon	141
7.12	Measured Transition Delay with $\Delta f = 500$ kHz	144
7.13	Simulated RX Demodulation Showing Delay	145
7.14	Receiver Comparison Strictly Considering Energy/Bit	151
A.1	Simple Model of a Feedback System	163
A.2	$-G_m$ LC Oscillator	165
A.3	Hartley LC Oscillator	165
A.4	Colpitts LC Oscillator	165
A.5	$-G_m$ LC Oscillator, Open-Loop Analysis	166
A.6	NMOS Modes of Operation	168
A.7	Parallel RLC Resonant Tank	173

List of Tables

2.1	Bergveld's Low-IF Receiver Power Budget	14
3.1	Optimized Antenna/Inductor Lumped Element Model Parameters . .	35
5.1	Lock-and-Roll Receiver LNA Extracted Performance	84
6.1	Lock-and-Roll RX Loop Filter Laser Options	101
6.2	Open-Loop Performance, Simulation vs. Measurement	117
7.1	Lock-and-Roll Receiver Measured Power Breakdown	146

List of Abbreviations and Symbols

C_{ox}	refers to the oxide capacitance per unit of gate area (of a MOSFET)
λ	refers to the channel length modulation factor (used in MOSFET transistor modeling)
μ_n	refers to electron mobility (here in silicon)
ADC	Analog-to-Digital Converter
ADS	Advanced Design System (a software package from Agilent Technologies)
AGC	Automatic Gain Control
BAW	Bulk Acoustic Wave (an off-chip filter or resonator)
BEOL	Back End Of Line (here refers to a high-quality integrated resistor implemented in top level, thick metal)
BiCMOS	Bipolar Complementary Metal Oxide Semiconductor (a manufacturing process enabling both bipolar and CMOS transistors on the same chip)
BFSK	Binary Frequency Shift Keying (an FM method of encoding data in a carrier signal)
BNC	Bayonet Neill-Concelman (a type of signal connector, typically 50 Ω and used for DC or low frequency connections up to 10 MHz)
BOM	Bill Of Materials
BPF	Band-Pass Filter
CMOS	Complementary Metal Oxide Semiconductor
CP	Charge Pump
DC	Direct Current (zero frequency)
DRC	Design Rule Check
EM	Electro-Magnetic (fields or waves)
ESD	Electro-Static Discharge

F	refers to the noise factor of a circuit
FDM	Frequency Down-Conversion Mixer (in a receiver)
FM	Frequency Modulation (a method of encoding data in a carrier signal)
FSK	Frequency Shift Keying (a form of digital FM)
FUM	Frequency Up-Conversion Mixer (in a transmitter)
GaAs	Gallium Arsenide (a more expensive, higher speed alternative to silicon as a substrate for ICs)
HFSS	High-Frequency Structure Simulator (a 3-D EM simulator from Hewlett Packard/Ansoft)
I-phase	In phase (one of two phases in quadrature modulated signals)
IC	Integrated Circuit
IF	Intermediate Frequency (a mid conversion frequency used in heterodyne transceivers)
IP3	refers to the third-order Intercept Point (a measure of amplifier linearity, can be input or output referred)
LC	refers to an inductive capacitive (typically parallel connected) resonant circuit
LNA	Low-Noise Amplifier
LO	Local Oscillator
LPF	Low-Pass Filter
MIM	Metal Insulator Metal (a type of capacitor)
MOS cap	refers to a MOSFET device with drain and source connected together to form a capacitor
MOSFET	Metal-Oxide Semiconductor Field-Effect Transistor
mux	a short form for “multiplexer”
NA	Network Analyzer (an instrument used to measure S-parameter performance)
NF	Noise Figure (F in decibels)
NPN	Refers to alternating negatively, positively, negatively doped regions of silicon, as in an NPN bipolar transistor
OOK	On-Off Keying (a method of encoding data in a

	carrier signal)
OCRI	Ottawa Center for Research and Innovation
P1dB	refers to the 1-dB compression point (where an amplifier's gain drops from its linear extrapolation by 1 dB, indicating gain compression, can be input or output power referred)
PA	Power Amplifier
PAL	Phase Alternating Line (a colour television encoding scheme primarily used in Europe)
PCB	Printed Circuit Board
PCT	Patent Cooperation Treaty
PFD	Phase Frequency Detector
PLL	Phase-Locked Loop
PN	refers to the junction between positively doped and negatively doped regions (of silicon)
ppm	parts per million
Q-phase	Quadrature phase (one of two phases in quadrature modulated signals)
QAM	Quadrature Amplitude Modulation (a method of encoding data in a carrier signal)
RF	Radio Frequency
RFIC	Radio Frequency Integrated Circuit
RFID	Radio Frequency Identification
RX	Receiver
SAW	Surface Acoustic Wave (an off-chip filter or resonator)
SiGe	Silicon Germanium (a higher performance substrate material than that of bulk CMOS)
SMA	Sub-Miniature type 'A' (a type of cable connector used for RF connections up to 18 GHz)
SNR	Signal-to-Noise Ratio
SoC	System on a Chip
SoI	Silicon on Insulator (an alternative substrate material with better isolation properties than that of bulk CMOS)

TX	Transmitter
USB	Universal Serial Bus
VCC or VDD	refers here to highest on-chip voltage supply rail
VCO	Voltage Controlled Oscillator
VHF	Very-High Frequency (the 30 MHz to 300 MHz frequency band)
VSS	refers here to lowest on-chip voltage supply rail

Chapter 1

Introduction

1.1 Motivation

Today's society is adopting new technology at a phenomenal rate and is demanding feature rich wireless devices capable of communicating signals with high data rates over long distances. What some people might find surprising, however, is that the demand for short range wireless devices which are only capable of transmitting signals with low data rates, but that do so while consuming extremely little power, is also increasing.

The idea of using radio frequency identification (RFID) tags for corporate or personal asset management or for scanning commercial items during shipping, while counting inventory, or at time of sale has ballooned in popularity in recent years. The 2010 Ford F-series pick-up trucks, for example, have an innovative "tool inventory" option that allows contractors to toss their tools (having RFID tags installed) into the bed of their truck and to later take inventory from the driver's seat, using RFID tags to eliminate the chance of leaving something valuable behind while quickly moving between job sites. At the same time, medical sensors which can be embedded in the body, ingested, or simply placed on the surface of the body from where they relay vital information over a wireless link are also gaining acceptance. Both of these application examples, to name but two, require low-power radio frequency (RF) circuits and would benefit from a solution that is so power efficient that it could be self-powered, or make use of power scavenging techniques. In fact in the case of an embedded or ingestible sensor, operating from a battery power source is typically impractical, and poses a possible health risk due to the chemical makeup of most batteries.

Consider the example of an RFID tag being used to identify items as inexpensive as a loaf of bread or a package of chewing gum at a local grocery store. Assuming that the tag would stay with the item it is identifying or be disposed of after the item is purchased, the tag would have to be extremely cheap to manufacture. This economic reality leads to two reasonable conclusions. Firstly, the circuit should be completely integrated onto a single chip minimizing the use of expensive external passive components. This could also be carried as far as to implement the antenna on chip, minimizing the size of the overall solution and again reducing the number of off-chip components required. Secondly, the added cost of using a battery to power the solution would be substantial, and as such, a self-powered circuit or one that practices power scavenging would be of economic benefit. Indeed the additional chip space required for implementing an on-chip antenna and to enable power scavenging will increase the complexity and fabrication costs of the integrated circuit (IC) itself, but in most cases where a particular chip is being manufactured in high volume, the added cost per chip should be far less than what's saved by eliminating the need for off-chip antennas and power supplies.

Another application example with surprisingly similar requirements is a wireless dosimeter used to measure the radiation dosage received by cancer patients during treatment. Current generation metal-oxide semiconductor field-effect transistor (MOSFET) dosimeters [1] are placed on the body and wired to a central hub located near the patient's waist which then relays the data to a computer for analysis. Figure 1.1 shows a patient wired up with two of today's status quo wired dosimeters, one on her cheek and one on her eye lid. Figure 1.2 shows a typical setup with a patient laying on a bed awaiting treatment with 16 individual wired dosimeters placed on his body. Other than the obvious discomfort associated with having a wire pulling on sensitive tissue like an eye lid, the metal wires that connect the sensors to the hub run across the body and can block the radiation – an undesired side-effect. For this reason, a wireless version of the sensor would be preferred, and need only be capable of transmitting very little data over a very small distance. At the same time, a wireless solution that makes use of a battery power source could lead to additional problems. Batteries typically contain heavy metals which can deflect radiation. Consequently, a wireless sensor whose circuits consume so little power that the solution can be self-powered or make use of power scavenging techniques is highly desired.



Figure 1.1: Status Quo Wired Dosimeters in Use

While the case for a battery-free RFID tag solution might be an economic one, clearly the argument for a self-powered medical radiation dosimeter is fueled by the need to provide what's best for the patient. Ultimately, the requirements are quite similar and can be addressed by using a combination of power efficient system level design strategies and circuit level implementation techniques.

1.2 Thesis Objectives

The intention of this thesis is to present new and novel ways of achieving inexpensive, low-power RFICs which consume so little power that they can be self-powered, and therefore completely integrated onto a single chip including the power supply. As circuits of this type are ideal for meeting the needs of RFID applications as well as for transmitting data from medical sensors, two applications where including the antenna on-chip is also beneficial, methods of designing systems and circuits that complement an on-chip antenna are also explored. In order to validate the system level topologies which are suggested, along with their circuit level implementations,




Figure 1.2: Patient Awaiting Treatment using Wired Dosimeters

the design, simulation, layout, and measured results of a novel “lock-and-roll” receiver (RX) are referred to throughout. The thesis objectives can be summarized as follows:

1. To explore previous low-power, completely (or mostly) integrated RFIC topologies for short range, low-speed data reception and to suggest new approaches for overcoming their limitations.
2. To demonstrate the feasibility of the new approaches proposed in 1 using a uniquely modified and completely integrated phase-locked loop (PLL) (nicknamed the lock-and-roll receiver), operated in the open and closed states, as an FM demodulator to an input signal with a center frequency of 5.2 GHz.
3. To demonstrate the feasibility of a system, namely the lock-and-roll receiver, making use of an on-chip antenna for the purpose of short range communications at 5.2 GHz.
4. To achieve such low power consumption from the lock-and-roll receiver design that powering the integrated circuits with ultracapacitors, which can be charged using a solar cell, is theoretically feasible, therefore proving, to a first order, that the entire receiver (and potentially the entire transceiver) can be integrated onto a single chip including the antenna and power supply.

1.3 Thesis Organization

This thesis is organized to provide background information on low-power, completely integrated RFIC designs, followed by a discussion of new and novel techniques for improving the state of the art. Paralleling the presentation of these techniques is the proposal of a completely new receiver architecture, referred to as the lock-and-roll receiver.

Chapter 2, *Background*, focuses on the inevitable performance metric tradeoffs faced by RFIC transceiver designers, and explores previously published novel techniques of achieving low-power RFICs and identifies the strengths and weaknesses of past designs. Additionally, Chapter 2 provides an overview of the previous work of the author and that of other researchers in the field which have demonstrated the use of oscillators as high-gain, narrow-band amplifiers, commonly used at the core of many low-power RF receiver architectures. Similarly, previous works by the author and others which have demonstrated the potential of using on-chip coil inductors to couple signals, thus serving as antennas, are also presented. Finally, Chapter 2 introduces ultracapacitors as on-chip power sources and discusses different approaches for scavenging power from the IC's working environment. This background information serves to inform the reader of previous work in the field such that the novelty of the system and circuit design strategies (for achieving low power consumption and completely integrated RFICs) which are presented and practiced in the following chapters can be appreciated.

Chapter 3, *The Proposed Lock-and-Roll Transceiver*, presents a completely new RFIC transceiver design that requires virtually no external passive components, consumes little power, and complements the use of on-chip antennas. Though the author's research focuses primarily on the design of the RX portion of the transceiver, Chapter 3 provides a necessary overview of the entire system. Only by carefully planning and understanding the overall architecture of an RFIC at the system level can one ease the task of designing circuits that meet the block and level specifications set forth by the system architect and simultaneously achieve low power consumption with all passive components and antennas being integrated on chip. The transceiver system design and circuit implementation, of which the receiver represents roughly one half, is the culmination of the research work of the author and one other Ph.D. student.

The author takes credit for the system and circuit level design of the receiver (RX), while acknowledging that the transmitter (TX) is primarily the work of his colleague. Block level analyses of both the transmitter and receiver sections are presented, along with link budget calculations and a discussion of the tradeoffs that can be made between data rate, communication range, and power consumption. Chapter 3 outlines the division of work amongst members of the research team, and clarifies exactly which work is the subject of the section of this thesis that focuses on circuit design. The design and simulation of the on-chip antenna, which represents the work of a third team member, is summarized in Chapter 3 (with appropriate references and permissions) as the results are necessary to understanding the link budget calculations and the design of the low-noise amplifier (LNA) circuit which is part of the receiver.

Chapter 4, *Injection-Lockable VCO Design for Low-Power Applications*, discusses the theoretical tradeoffs between designing a VCO that has good phase noise and large signal swing, which is the target of a typical VCO design, with designing one that is easily injection locked while consuming minimal power. Putting the theory into practice, Chapter 4 presents the design of the lock-and-roll receiver's VCO which must be injection locked by the incoming signal. Detailed discussions of the VCO circuit schematic and layout are included and simulation results which are pertinent to the predicted communication range of the topology, presented in Chapter 3, are shown.

Chapter 5, *Injection-Locking Circuit Design*, focuses on different methods of introducing injection-locking signals into the core of an LC oscillator without disrupting the natural resonant frequency of the tank circuit, and consequently the frequency response of the VCO, or excessively lowering the oscillator's Q factor. As an example, the design of the low-noise amplifier in the lock-and-roll receiver (which is impedance matched to the on-chip antenna and is used to amplify the incoming RF signal to a level that is sufficiently large to injection lock the VCO) is presented. Both the LNA's schematic and layout are scrutinized, and simulation results validating the design and reaffirming the calculations of Chapter 3 are presented.

Chapter 6, *PLL Component Designs that Enable Open-Loop Operation*, addresses the need for designing modified PLL components which differ from their traditional counterparts in that they allow the loop to be operated in open-loop mode, consequently enabling low-power RF transmission and reception and complete system

integration. The design of the adjustable loop filter, the loop filter switch, the phase frequency detector (PFD), the charge pump (CP), and the loop-filter buffer which are used on both the RX and TX of the lock-and-roll transceiver are presented and discussed as examples, highlighting key layouts, as well as simulated and measured results. As some of these blocks were designed with the help of a fellow team member, the designs are presented with the mutual understanding of both designers that these circuits are important to the functionality of both the RX and the TX, and that they are the result of many hours of collaborative design efforts.

Chapter 7, *The Lock-and-Roll Receiver Test Chip*, summarizes the lock-and-roll receiver test chip which is the ultimate application example of the system and circuit design strategies that are discussed in Chapters 2-6, showing top level schematics and layouts, along with simulated and measured results. This chapter explains aspects of the test chip design that served only to facilitate testability (such as probe de-embedding structures, laser trim options, etc.) and which were not captured elsewhere in the preceding chapters. The divider, the up down pulse multiplexer (mux), and the pre-pad buffer circuits are also presented here. Limitations of the test chip are discussed, as well as strategies that were employed to overcome certain limitations in order to make the most of the testable silicon.

In Chapter 8, *Conclusion*, the research work is summarized and the key achievements are revisited and highlighted. Future work and possible extensions of the research experiment are discussed. Publications and patent filings that were a direct result of this work are listed.

Chapter 2

Background

2.1 Recognizing the Tradeoffs in Low-Power RFICs

Designing a completely integrated transceiver that achieves ultra-low power consumption, thus enabling the use of on-chip power sources, and can communicate via on-chip antennas while achieving acceptable range and data rates requires a delicate balancing of numerous well known tradeoffs.

An antenna's gain (and efficiency) is directly proportional to its size relative to the wavelength (λ) of the carrier it communicates. The wavelength of a signal is inversely proportional to its frequency according to $\lambda_0 = c/f$ where λ_0 is the signal's wavelength in air, and c is the speed of light in air. As a result, higher carrier frequencies (> 5 GHz) are conducive to integrated antennas with adequate gains. Unfortunately, the power consumed by an integrated circuit is directly proportional to the frequency at which it operates. Standard complementary metal oxide semiconductor (CMOS) logic, for example, has near zero static power consumption and only consumes appreciable current at the moment it switches. Therefore, a circuit that experiences more switching events over time, because it is processing a signal with a higher frequency, will inevitably consume more average current than a circuit operating at a lower frequency. Thus, there is a clear system tradeoff between antenna gain and overall power consumption due to the operating frequency.

Similarly, there is an indirect tradeoff between a system's overall cost and the frequency of operation. As previously mentioned, higher carrier frequencies enable the use of on-chip antennas which reduce the bill of materials (BOM) for a particular solution and hence its cost. This is only true assuming the added chip space (required for the antenna) is cheaper than the off-chip alternative, but assuming a standard

high-volume CMOS process this is a reasonable conclusion. Also, higher frequency circuit designs typically make use of smaller sized capacitors (for decoupling, filtering, etc.) than their lower frequency counterparts, and reduced component sizes result in less layout area, more chips per wafer, and ultimately translate into lower individual product costs.

Yet another tradeoff, but one that sets an upper bound on the frequency of operation, arises when one strives for a self-powered design that uses on-chip energy storage. Modern ultracapacitors are an effective means of storing relatively large quantities of energy on chip (where the ultracapacitor is manufactured on top of the chip), yet they are limited in terms of the speed at which they can deliver that energy to an integrated circuit. Standard metal-insulator-metal (MIM) capacitors can serve as a buffer, providing charge storage closer to the circuit and delivering smaller quantities of energy at faster rates where the MIM capacitors are recharged by the higher density ultracapacitors, yet these MIM capacitors require chip area (and added cost), thus bringing a third dimension into the tradeoff.

Additionally, low antenna gain can be compensated for by a receiver design with high sensitivity to maintain a desired communication range. This usually comes at the expense of more complicated circuits that consume more power. Furthermore, many circuit designers struggle to manage additional tradeoffs while trying to maximize specifications like sensitivity and data rate. Figure 2.1 summarizes visually the tangled relationship among these numerous system considerations. In evaluating the state of the art with respect to ultra-low power communication solutions, one must constantly consider these tradeoffs, recognizing that no solution will excel at all metrics and that the best solution is likely one that is also the most balanced.

Worthy of note is that Razavi [2] summarizes the general design tradeoffs encountered using his own “RF Design Hexagon”, and while his hexagon is very relevant when considering general RF design concerns, the tradeoffs outlined in Figure 2.1 are more specific to self-powered, low-cost, and completely integrated short range communication systems.




Figure 2.1: The RFIC Design Star of Tradeoffs

2.2 Applying the Tradeoffs – Receiver Examples from Literature

2.2.1 Low-Power at the Expense of Cost, Size and Integration

An interesting example of the tradeoffs depicted in Figure 2.1 is the 176 kHz receiver for frequency shift keyed (FSK) signals described in [3] and summarized in Figure 2.2. The receiver design is completely made up of discrete components, where an antenna receives the FSK signal which is then band-pass filtered and amplified. The signal is then applied to a low-pass filter (LPF) and a high-pass filter (HPF) in parallel to detect energy at either of the two frequencies representing a high (digital bit 1) or a low (digital bit 0), and the outputs of the two filters are connected to a comparator that determines the receiver's output.

The receiver is very simple, but the total BOM costs add up to about \$5 worth of discrete components. An additional drawback of the design is that in order to enable the system to make use of a relatively inexpensive crystal reference with a stability of ± 300 ppm, the designers chose a data format that limits the receiver to a data rate of only 100 b/s. Due to its simplicity the receiver is low-power, consuming only 6 mW while receiving.




Figure 2.2: Chu's Low-Frequency, Low-Power FSK Receiver at 176 kHz

Needless to say, the frequency of operation demands a large antenna, and with many discrete components the overall size is large which limits the receiver in terms of its potential applications. No reference is made in [3] to the design of the transmitter circuit which communicates with the receiver, to the communication range of the system, or to the system's application. The designers claim a 2.5 year battery life from the receiver owing to the fact that the circuit is enabled by a low-power timing device, which enables the receiver for only 148 seconds per day.

This design achieves low power at the expense of integration and size, making tradeoffs between overall cost and the maximum data rate.

2.2.2 High Frequency and Power with the Benefit of Integrated Antennas

The designers of the receiver topology described in [4] chose to tradeoff the metrics of Figure 2.1 much differently than [3]. The result is a design that receives 5 GHz signals, utilizes an antenna which is integrated onto the printed circuit board (PCB), but not the chip, and processes data rates of up to 50 Mb/s. The antenna design is described in more detail in [5].

The topology is based on the standard heterodyne receiver [2]. Although [4] claims a system-on-a-chip (SoC) solution that consumes 20 mW, careful inspection reveals that assembling a working receiver requires interconnecting the LNA with a power splitter, a VCO, filters including at least one surface acoustic wave (SAW) filter,

amplifiers, and a mixer fabricated in an expensive gallium arsenide (GaAs) process all onto a PCB that contains the “integrated” antenna. In addition, the claimed 20 mW power consumption only applies to the LNA, leaving the reader to wonder just how much power is actually required to receive a signal. A benefit of the topology is that it is apparently capable of receiving data encoded using 64 QAM [6] at a range of 100 m. There is no reference in [4] to the transmitter circuit that communicates with the receiver and no mention is made as to the cost of the receiver either, although one can assume it’s much greater than the \$5 reported by [3].

Clearly the ideal solution for a communication system that is ultra-low power is unlikely to be based on the traditional heterodyne topology – simply based on the numerous building blocks, which all consume power, that make up a heterodyne receiver. That said, [4] attests to the fact that as the frequency of communication increases, the physical size of the antenna can be reduced. In the case of [4], a 100 m communication link is claimed using a 29 mm x 34 mm sized antenna that was integrated onto the PCB – an impressive result regardless of the system’s apparent drawbacks.

2.2.3 Minimizing the BOM while Balancing other Metrics

A receiver design that is cheap to manufacture and assemble (in terms of its BOM) in high volume must require pairing with a minimum of off-chip components, and should be manufacturable in a standard CMOS process which is competitively priced. This theory was definitely on the minds of the designers who produced the receiver described in [7].

The receiver in [7] communicates at 2.4 GHz while consuming a total of 32 mW and is fabricated in a standard 0.18 μm process. The solution is completely integrated onto a single chip except for the antenna, the band-pass filter, the crystal reference and the battery. As a result, the solution could be manufactured in high volume at a competitive price point although [7] does not allude to the system’s cost.

The receiver handles quadrature modulated signals, which have in-phase (I) and quadrature-phase (Q) components, and is based on a low intermediate frequency (IF) topology as shown in Figure 2.3. The external band-select filter chooses the 2.4 GHz band and impedance matches the antenna to the LNAs. Two LNAs are

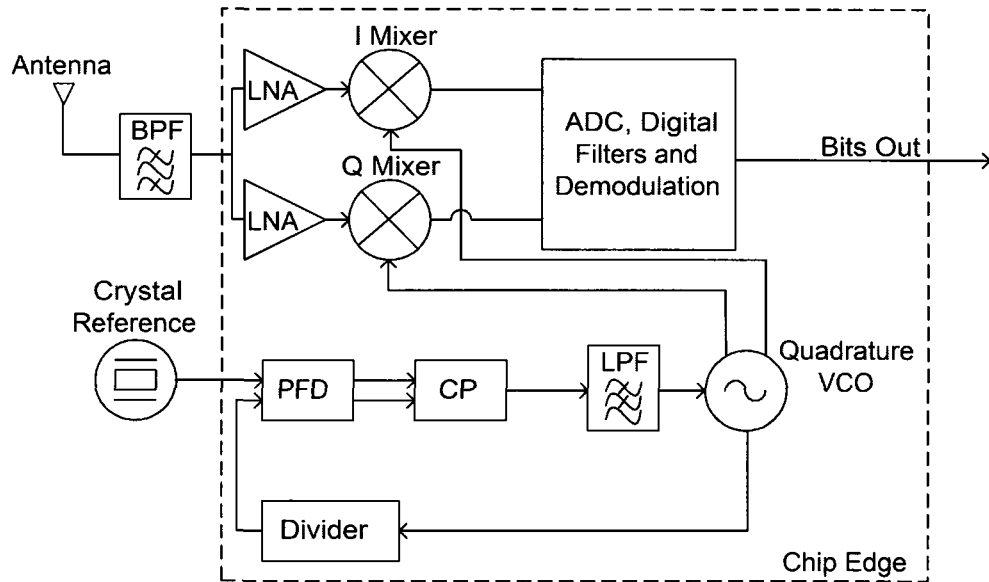


Figure 2.3: Bergveld's Low-IF Receiver Topology at 2.4 GHz

used in parallel to maximize I-phase and Q-phase path isolation. Passive mixers down-convert the I-phase and Q-phase signals to a 500 kHz IF and a complex $\Sigma\Delta$ analog-to-digital converter (ADC), followed by additional digital circuitry, processes the output bitstream. In order to down-convert the 2.4 GHz I-phase and Q-phase signals to IF, the local oscillator (LO) ports of the mixers are driven by a quadrature VCO whose phase is accurately controlled by a PLL that is referenced to the off-chip crystal.

The topology is unique in that it relies heavily on its digital circuitry in order to eliminate some of the filtering that typically takes place in the analog domain of a heterodyne receiver. The authors claim that the receiver chain is linear, enabling it to be used with a wider range of modulation schemes including those that don't have a constant signal envelope. There is no mention in [7] as to the maximum data rate that the receiver can process, as to what kind of transmitter it was tested/designed with, nor to the communication range of the system. Consuming 32 mW, [7] claims that the receiver is "a factor of two lower than *the* state-of-the-art CMOS receivers" which may be true when compared to receivers with the same data rate and communication range (which are unknown to the reader), yet for short range, low data rate communications

there are less power hungry design alternatives than the topology shown in Figure 2.3.

In addition to the drawback that the topology has relatively high power consumption compared to some alternatives, [7] identifies a concern that is common to many mixed signal chips where digital and analog circuits share the same substrate. As the topology in Figure 2.3 consists of a substantial number of noisy digital circuits, the designers had to take extra steps to try and improve the isolation between the digital and analog portions of the chip including the use of strictly differential circuits in the analog domain, guard rings in the layout, numerous separate supplies, and separate analog and digital pad rings. The last two isolation techniques, unfortunately, greatly increased the number of decoupling capacitors that were required off chip, which will increase the BOM. This subtlety is not emphasized in [7] where the system is claimed as requiring only three off-chip components, namely the antenna, the band filter, and the crystal. Given the power requirements, one also assumes an off-chip battery as the power source that is also omitted from the list in [7].

Lastly, breaking down the power consumption budget for the system described in [7] provides insight into which aspects of the design should be avoided if one is to attempt a more power efficient design. The power budget is summarized in Table 2.1.

Table 2.1: Bergveld’s Low-IF Receiver Power Budget

Receiver Element	Power Consumption
LNA	$0.6 \text{ mA} * 1.8 \text{ V} = 1.1 \text{ mW}$
PLL	$7.6 \text{ mA} * 1.8 \text{ V} = 13.7 \text{ mW}$
ADC	$2.3 \text{ mA} * 1.8 \text{ V} = 4.1 \text{ mW}$
Bandgap	$0.1 \text{ mA} * 1.8 \text{ V} = 0.2 \text{ mW}$
Crystal Oscillator	$1.2 \text{ mA} * 1.8 \text{ V} = 2.2 \text{ mW}$
Digital Filters + Demodulators	$7.4 \text{ mA} * 1.4 \text{ V} = 10.4 \text{ mW}$
Total	31.7 mW

Note that to conserve power, the digital sections of the chip are powered from a 1.4 V supply rather than the 1.8 V supply used for the analog sections. This adds further complexity to the design in terms of power regulation circuitry requirements if both are to be supplied from the same battery. This matter is not discussed in [7].

From Table 2.1 the two most power hungry aspects of the design are the ADC and digital circuits (14.5 mW combined) and the PLL (13.7 mW). If the ADC and digital circuits could somehow be eliminated from the receiver design, the power consumption would drop by nearly 50 %. Also, as is apparent in subsequent sections of this thesis, certain low-power design strategies can be employed to decrease the power consumption of an RF PLL. The PLL in [7], an integer-N design at 2.4 GHz consumes roughly twice the power consumption of the 5.2 GHz integer-N PLL required for the lock-and-roll receiver described in the subsequent chapters of this thesis.

2.3 Circuit Topologies Using Oscillators as Gain Elements

A receiver design with good sensitivity, which can therefore support a communication link over an adequate range, inevitably requires a high-gain element that consumes much power. One of the most promising design strategies, in terms of achieving high gain with minimal power consumption, is to use oscillators in nontraditional roles as gain elements.

2.3.1 The Theoretical Transfer Function of an Oscillator

Previous works by the author [8], [9] and others [10], [11] have focused on simplified linear models of the oscillator that predict its behaviour as a high-gain element. This thesis will only summarize such oscillator behaviour as is necessary for understanding the details of Q -enhanced filters and regenerative circuits.

An oscillator is often viewed as a tuned circuit that produces a dominant output tone at the resonant frequency ω_0 of the tank circuit with instantaneous deviations regarded as phase noise. The phase noise power at an offset $\Delta\omega$ from the oscillator's center frequency at ω_0 is well known to roll off at 20 dB/dec, where $\Delta\omega$ is large, as indicated by the commonly accepted phase noise model [12]. Excluding the power supply, the only input to an ideal oscillator circuit is white noise. In order to produce what appears as a single output tone at ω_0 , the oscillator must behave as an extremely high- Q band-pass filter with exceptionally high, but finite, pass-band gain. This response can be derived mathematically [8], [10], and demonstrated through

simulation and measurement [13], [14]. Figure 2.4 summarizes the gain response of the oscillator when it is free running, where Q_U is the quality factor of the unloaded tank circuit, F is the noise factor of the transconductor, k is Boltzmann's constant, T is the temperature in degrees Kelvin, and ω_0 is the resonant frequency of the LC tank, i.e., $\omega_0 = 1/\sqrt{LC}$. If the 3 dB bandwidth of the unloaded tank circuit is represented by B_U then the quality factor of the loaded oscillator is defined [10] by

$$Q_L = \frac{Q_U P_{out}}{F k T B_U} \frac{2}{\pi} \quad (2.1)$$

and the 3 dB bandwidth of the loaded oscillator, B_L , follows. The $1/f$ noise is assumed to be an insignificant contributor for the Q_L of the oscillator in question.




Figure 2.4: Oscillator Frequency Response

This representation of the free-running oscillator supports Leeson's phase noise model [12] and explains the gain of injected tones that are small enough that the oscillator is not injection locked. Once the power in the injected tone is sufficient to lock the oscillator however, the frequency response changes drastically as the time-varying transconductance, $g_m(V_{out}(t)|_{\omega_0})$, is no longer correlated to ω_0 and instead follows $g_m(V_{out}(t)|_{\omega})$. The power at ω_0 falls to the noise floor as that tone is no longer coherently integrated by the large signal oscillator. The amplitude of the injected signal required to injection lock the oscillator was analyzed and quantified by [15], and the topic was recently revisited by [16]. Assuming that the amplitude

of the injected voltage (V_{inj}) is much smaller than the amplitude of the free-running oscillator (V_{osc}), the locking range can be approximated by

$$\omega_L \approx \frac{\omega_0}{2Q_U} \frac{V_{inj}}{V_{osc}} \quad (2.2)$$

where ω_L is the single sided locking bandwidth, i.e., the oscillator can be locked from $\omega_0 - \omega_L$ to $\omega_0 + \omega_L$ [15].

Equation (2.2) attests to the fact that the gain of the oscillator is indeed highest (although it is not infinite) at the center frequency of the unloaded tank circuit, and that the oscillator is therefore easier to injection lock at smaller frequency offsets from ω_0 .

As a result of exhibiting the transfer function depicted in Figure 2.4, the oscillator is suitable to be used as a narrow-band, high-gain filter or amplifier.

2.3.2 Q - Enhanced Filters

Publications demonstrating the successful use of oscillator type circuits, [17], [18], [19], [20], categorized as Q -enhanced filters, are still rare but cite positive results.

The topology in [17] uses an oscillator type circuit to create a high- Q filter with digital tuning, shown in Figure 2.5.

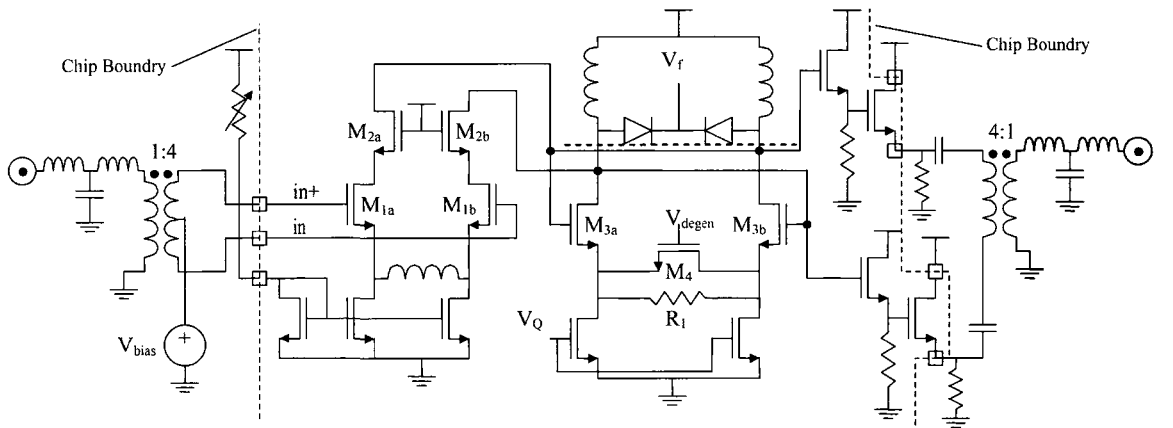


Figure 2.5: DeVries' Q -Enhanced Filter with Digital Tuning

Like an oscillator circuit, the filter circuit is comprised of a gain stage and an LC tank stage with negative resistance. The primary difference between the filter and the oscillator is that the g_m of the filter circuit is smaller than g_0 , the admittance of the tank. To make the circuit an oscillator, g_m would have to be made larger than g_0 . The authors of [17] designed the filter to be used at the IF in a heterodyne RF receiver as shown in Figure 2.6. The filter design showed promising results, demonstrating a $Q \approx 650$ at 500 MHz with a 750 kHz bandwidth while consuming a total power of only 1.02 mW from a 1.8 V supply. However, the author does not allude to the overall power consumption of the receiver which has an LNA, a mixer, and an oscillator all operating at RF. Additionally, the RF filter between the antenna and the LNA is likely implemented using off-chip components, adding to the BOM.

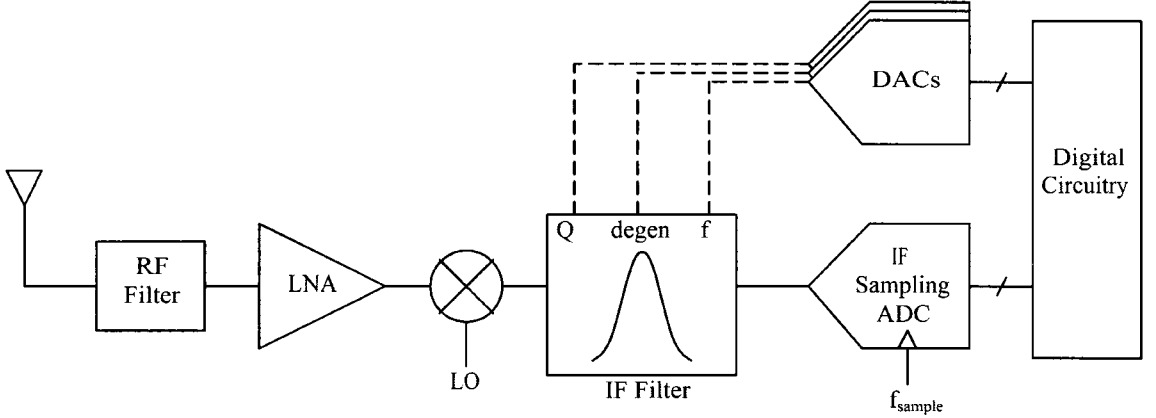


Figure 2.6: DeVries' RF Back End with Q -Enhanced Filter

As mentioned, a Q -enhanced filter differs from an oscillator in that the oscillator has $g_m > g_0$, while $g_m < g_0$ in a Q -enhanced filter, and therefore the filter does not start to oscillate on its own. One of the oldest, most power efficient and low-power receiver topologies on the market today blurs this defining line, putting itself into a category all it's own; the super-regenerative receiver.

2.3.3 The Super-Regenerative Receiver

The concept of a super-regenerative receiver, which was invented by Armstrong [21] in 1922, uses an LC oscillator whose bias is adjusted by an automatic gain control (AGC) loop to keep the oscillator gain low enough that the oscillator does not start

up on its own with no input signal. The incoming signal is injected into the oscillator at its center frequency f_0 defined by

$$f_0 = \frac{1}{2\pi LC} \quad (2.3)$$

and the oscillator is used as a high-gain element to regenerate the weak input signal. This topology traditionally makes use of on/off keying (OOK) to encode the data signal and has seen much success and implementation in the area of low-power, short-range communications required for alarm systems or garage door openers where the receiver must “sniff” or listen for an incoming signal while consuming very little power over long periods of inactivity. As the oscillator circuit only begins to oscillate (and consume appreciable power) when an incoming signal is present, and not otherwise, the topology is extremely power efficient for these applications.

Recent designs [22], [23], have claimed impressive power consumption as low as 1.2 mW, but [22], which is implemented in a standard 0.35 μm CMOS process, makes use of numerous off-chip passive components as does [23], which is implemented in a more expensive bipolar complementary metal oxide semiconductor (BiCMOS) technology. These examples further attest to the tradeoffs depicted in Figure 2.1. In fact all designs of this type to date (to the knowledge of the author) have made use of battery power sources and discrete antennas at the very least, and are therefore not completely integrated solutions. Another drawback of super-regenerative receivers is that they often have poor adjacent channel rejection, and their oscillators can be injection locked by interfering signals. Additionally, OOK modulation, as the name suggests, involves turning the output signal completely on and off at the TX and can cause splatter in the frequency domain which can be very detrimental to other wireless devices operating in the vicinity [2]. This problem can be corrected for in the TX, but at the cost of additional circuits and complexity.

2.3.4 Injection-Locked and PLL Based Receivers and Transmitters

Injection-locked RX topologies have been proposed before, such as [24], making use of the high-gain properties of an integrated VCO at an RX front end with the obvious limitation that the input signal must be narrow band enough to fall inside of the

locking bandwidth defined by (2.2). The topology proposed in [24] may be mostly integrated on chip, though the author does not comment to this effect, but clearly does not include an on-chip antenna or operate from an on-chip power source. In fact the injection-locked oscillator is part of a large, complex and relatively high-power solution for receiving very-high frequency (VHF) television signals using the European phase alternating line (PAL) colour encoding scheme. Shown in Figure 2.7, the solution uses an injection-locked oscillator/mixer to “amplify” the FM input

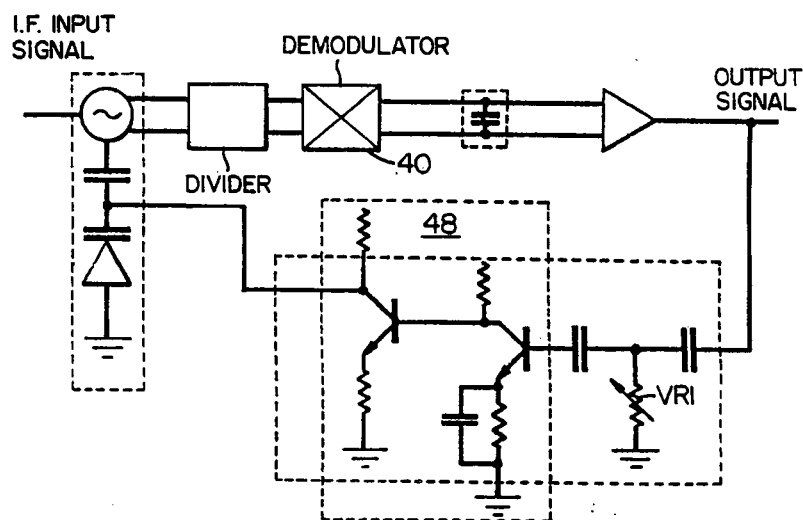


Figure 2.7: Plessey’s Injection-Locked FM Demodulator

signal and uses a correction loop to maintain the oscillator’s center frequency in the RX during reception such that the oscillator’s center frequency does not drift – maintaining the crucial overlap between the locking range and the input signal. The output of the oscillator is divided down in frequency by a frequency divider circuit and the signal is then demodulated. The frequency correction signal is derived from the output bitstream only, which when compared to a traditional PLL correction approach eliminates the need for a reference signal. There is no mention in [24] as to the power consumed by the injection-locking circuit but one presumes that in order for the oscillator’s locking bandwidth to be adequate for video signals the Q of the oscillator circuit must be low and the amplitude of the injection-locking circuit must be high, in both cases adding significantly to the power consumption of the entire

system. Regardless of the system's power consumption, the RX is unique and is said to be suitable for demodulating VHF signals at frequencies below 600 MHz.

The topology proposed in [25] is an example of previous PLL-based receivers. The receiver is shown in Figure 2.8, and uses two monostable multivibrator circuits to

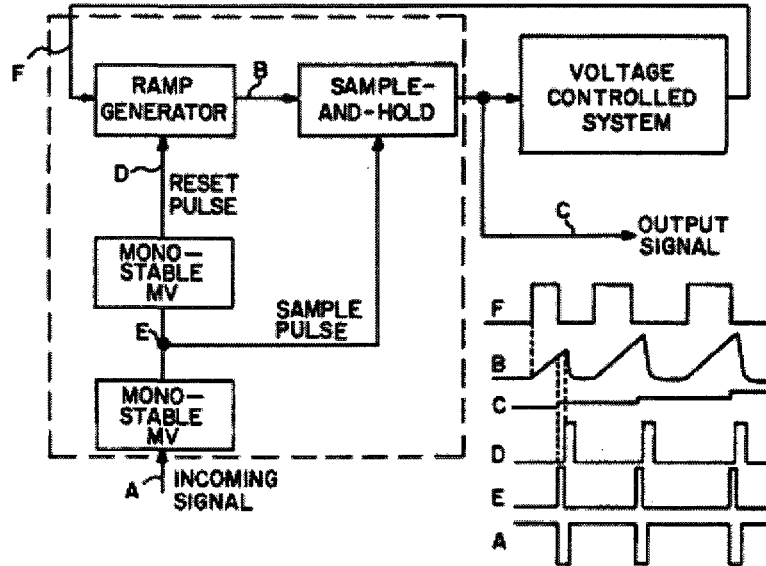


Figure 2.8: CIT's PLL-Based FM Demodulator

condition the FM input signal, essentially inverting it and delaying it in time to create a reset signal for the core block of the circuit – a ramp generator which essentially replaces the PFD found in most standard PLLs. The output of the ramp generator is sampled and held to yield the output signal which also serves as the control signal for a VCO circuit which drives the trigger input of the ramp generator. There is no reference in [25] as to whether the RX was implemented as an IC, and with a claimed operating frequency of 7 kHz for an FM rate of 80 Hz the circuit is definitely not compatible with an on-chip antenna and is unlikely to ever be used in a duty-cycled mode where the average power consumption could be lowered. Nevertheless, the approach is a novel one and demonstrates the potential for PLL-based circuits to perform very low power FM demodulation.

Modified PLL circuits have also been used with varying levels of past success to form FM modulators, or TX topologies. The topology in [26] uses a PLL where the loop is opened and the VCO circuit directly modulated to create the FM signals.

While the loop is open the topology depends on power-inefficient digital circuitry and a complex algorithm to correct for VCO drift. Additionally, the topology does not make use of an on-chip antenna or power source.

2.4 On-Chip Antennas

2.4.1 Coupling Inductors – a Convenient Accident

Previous research projects conducted by the author, [9], [13], successfully used oscillators as high-gain devices for the purpose of measuring and quantifying on-chip inductor coupling. Purely by accident, the author realized that the coupling levels were so strong that signals could be coupled from IC to IC, and not merely between inductors on the same die. As the experimental chips in [9] and [13] were not designed to measure the coupling between separate dice, quantifying the chip to chip coupling was inaccurate, yet the experiments proved that on-chip inductors could serve as antennas and the qualitative results spurred the author towards his current research direction.

2.4.2 Designs on High-Resistivity Substrates

The previously published work of others has also successfully demonstrated the use of on-chip inductive antennas for RF transmission at ranges up to 10 cm, [27], [28], yet all relied on high-resistivity substrates and often required an off-chip receive antenna.

A notable improvement in the state of the art would be an on-chip antenna with acceptable gain that can be manufactured in the low-resistivity substrate that is common among today's high-volume CMOS processes, facilitating integration with a transceiver circuit that could be manufactured easily (without post-processing steps) and in high volumes at a competitive price point.

2.5 Generating and Storing Power On Chip

2.5.1 Thin Film Ultracapacitors

Developments in the design and manufacturing of ultracapacitors, also known as supercapacitors, have made it possible to meet the power supply requirements of small integrated circuits without using a battery. Typical 100 μm thick nanostructured electrode devices can achieve capacitances of up to 1 F/cm^2 [29],[30]. Consider an IC measuring 2 mm by 2 mm of which an efficient RX/TX circuit might occupy roughly one quarter. This would allow for three 1 mm by 1 mm ultracapacitors to be manufactured on top of the remaining quadrants of the chip without covering up the RX/TX circuits. This would result in a 30 mF capacitance which would be capable of $\approx 4.2 \mu\text{Ahr}$ or $\approx 5 \text{ mA}$ for fifteen 200 ms bursts between chargings. Standard integrated MIM capacitors can be fabricated in the regular CMOS process below the ultracapacitors and serve as local charge storage devices because they can deliver charge quicker than the ultracapacitors which recharge them.

Recharging the ultracapacitors without connecting to off-chip components or power sources, thereby maintaining a true SoC solution, can be accomplished by using any one (or combination) of a number of techniques.

2.5.2 Thin Film Photocells, Thermogenerators, Inductive and RF Power Transmission

A solar cell can be manufactured on top of the ultracapacitors to trickle charge the ultracapacitors using ambient light. Previously published thin film solar cell designs have claimed the ability to source current at densities of 14.4 to 16.0 mA/cm^2 at 1.4 V [31], [32]. These results appear adequate to charge an ultracapacitor serving as the power source for a low-power circuit which requires a 1.2 V supply, as is the case for most circuits designed in today's 0.13 μm CMOS processes. The obvious drawback of relying solely on solar energy to recharge an ultracapacitor powered device is the dependence on available light at the time of recharging. Clearly this technique is inappropriate for an ingestible medical sensor that requires recharging while inside the body, or for an RFID tag that must be recharged while being covered by opaque packaging.

Alternatively, ultracapacitors can be recharged by an on-chip antenna structure coupled to a rectifier circuit which produces DC charging current from incoming RF signals. A similar technique is commonly used in RFID tags [33], [34], that communicate with their reader circuits [35], [36], (which transmit the RF signal that powers them) at 900 MHz. Due to the relatively low carrier frequency of these communications systems they make use of off-chip antennas, but because their circuits are switching at 900 MHz the system is relatively low-power such that the power output from the rectifier circuits are adequate – a clear management of the tradeoffs summarized in Figure 2.1. Unfortunately, these RFID systems suffer from a severe self-jamming drawback due to their communication protocol. Because there is no on-chip energy storage on the RFID tag device, such as an ultracapacitor, the RFID reader must transmit the 900 MHz power signal during the whole communication process. The tag essentially “wakes up”, and then retransmits the 900 MHz carrier back to the reader while using OOK or a similar modulation technique to encode and transmit the tag’s data. As a result, the reader has to contend with a 900 MHz jammer, coming from its own output, that is greater in signal strength than the data encoded signal that is coming from the tag (and is also at 900 MHz). Despite the system’s obvious drawbacks, it attests to the suitability of using RF signals to transmit power wirelessly.

Yet another alternative to relying on solar energy is to couple the power necessary to recharge the ultracapacitor on chip inductively. Essentially a transformer application, the secondary winding is placed on chip and a current is induced, by the magnetic field generated from passing current through the off-chip primary winding, and then applied to a load and rectified to yield a DC voltage.

The authors of [37] successfully powered a sensor transmitter from a rectifier circuit that harvested the 0 dBm incident output of a cellular phone at 2 GHz. Unfortunately, it took 5.5 hours of “charging” to yield the required 3.2 V supply voltage for their transmitter which draws a significant 11.4 mA supply current, and the energy harvesting circuit made use of an array of bulky off-chip patch antennas.

Some researchers have had success combining a number of different techniques for transferring power on chip. The system described in [38] combines the output of a thermogenerator circuit with that of an RF coil (and some rectifier/converter circuits) to gather energy from both thermal differences between the chip and the outside of

the package as well as from stray RF signals in the chip’s operating environment. One power supply management circuit handles the outputs of both systems and serves to charge an on-chip ultracapacitor.

2.6 Small-Form-Factor Crystals and On-Chip Self-Referenced *LC* Clocks

Perhaps one of the hardest elements of a communication system to implement on chip is the reference signal which is typically generated from a driver circuit that connects to an off-chip quartz crystal. The accuracy of references generated by means of this standard approach is typically on the order of ± 40 to ± 200 parts per million (ppm) across a temperature range of -40 °C to $+70$ °C, depending on the price and quality of the crystal, and after accounting for silicon process variations which vary the capacitance of the driver circuit. Such variations in the silicon, and consequently to the load on the crystal, translate to a change in the resonant frequency because of the crystal’s finite pullability factor which is typically quoted by the manufacturer in units of ppm/pF [39].

The latest generation of crystals is being made smaller than ever before. Epson/Toyocom manufactures crystals with a tolerance of ± 10 ppm and with a maximum temperature variation of ± 10 ppm from -40 °C to $+85$ °C in packaged form factors as small as 2 mm by 1.6 mm with a height of 0.5 mm [40]. Due to constraints on the physical size of the piece of quartz that can fit into this package, the minimum reference frequency available in this form factor is 24 MHz. A product designed to use this crystal could easily be packaged into a module along with the crystal itself to yield a very compact design.

Alternatively, researchers are actively striving to eliminate the need for off-chip quartz references altogether by designing self-referenced *LC* resonators that are highly accurate. The best published result so far [41] claims overall accuracies on the order of ± 400 ppm over temperature and silicon process variation. Designed to meet the standard for universal serial bus (USB) 2.0 devices, the reference may not be suitable for applications that require very high precision references, but with no external components the result is impressive nonetheless – leading one to conclude

that with continued developments in silicon processing, implementing a self-referenced ± 50 ppm clock might soon be feasible.

2.7 Background Summary

All in all, numerous RFIC designs have been proposed and implemented over the years that balanced the tradeoffs depicted in Figure 2.1 differently and with varying levels of success.

Circuits that operate at lower frequencies generally consume less power than their higher frequency counterparts, while the direct relationships between frequency and antenna gain and between antenna size and antenna gain suggest that a circuit that is designed for use with an on-chip (and therefore a physically small) antenna should operate at a high frequency in order to maximize antenna gain. These two requirements, low power and interoperability with on-chip antennas, would appear to be at odds with each other and any design that is to meet both requirements must clearly find an operating frequency that balances the two goals.

The use of oscillator and oscillator type circuits in nontraditional roles as efficient high-gain elements, such as in super-regenerative receivers and as Q -enhanced filters, has seen much success in previous designs. A notable benefit of the super-regenerative receiver is that its standby current consumption is extremely low, while a known drawback is that the circuit is easily pulled in frequency or injection locked in the presence of a strong interferer signal. Overcoming this drawback, injection-locked receivers have proven to be more robust to interference, but previous designs have so far been lackluster given their complicated feedback circuits, high power consumption, and lower operating frequencies. Modified PLL circuits have proved capable at performing FM demodulation, yet again the existing designs have generally operated at lower frequencies and are clearly incompatible with on-chip antennas.

The use of inductive on-chip antennas has only been explored to date by a handful of researchers and while some of the previous published results look promising, most designs have only achieved reasonable antenna gains with the use of high-resistivity substrates, yielding designs that are not easily transferable to the low-resistivity silicon substrates on which cost efficient CMOS RFICs are being fabricated. That said, previous experiments by the author with inductive structures

manufactured on low-resistance substrates showed, merely by accident, that coupled with the high-gain response of a VCO circuit, RF signals could be transmitted from one IC to another at 5 GHz. These results warrant further investigation as an integrated antenna that achieves reasonable gain and can be manufactured in the same high-volume CMOS process as the radio circuits it communicates with would be a significant improvement in the state of the art.

On-chip ultracapacitors are a relatively new research topic that has garnered much attention of late given the high capacitance/area densities that have been reported. The published results suggest the feasibility of using on-chip ultracapacitors as a power source and with the ability to manufacture solar cells on top of an IC to facilitate recharging of the ultracapacitors, or the implementation of an RF power scavenging/rectifying circuit on chip for the same purpose without the dependence on ambient light, a strong case can be made for a completely integrated communication device.

Off-chip crystal references are being made smaller than ever and can now be integrated into a module with the IC they connect to. Alternatively, research into eliminating the crystal altogether and generating accurate references on chip is showing promising results, leading one to wonder if using an off-chip reference is soon to become a bygone standard.

All these developments in the areas of circuit design, antenna miniaturization, on-chip power generation and storage, and quartz crystal replacement are clearly interesting in their own rights, but surely the greatest achievement of all will come from combining them together to accomplish a true SoC, performing as a wireless communication device, where all aspects of the design are integrated onto a single die (including the antenna, the power source and the reference) such that the BOM consists of only one item – a chip that can be manufactured in high volume at very little cost.

Chapter 3

The Proposed Lock-and-Roll Transceiver

The discussion of previous communication system architectures presented in Chapter 2 attests to the importance of considering many tradeoffs. Recall the tradeoff star from Figure 2.1, when striving to design a low-power, highly integrated, and cost efficient communication solution. If one is striving for a completely integrated solution with an on-chip antenna, the carrier frequency of the signal should be high in order to maximize the size of the on-chip antenna relative to the wavelength of the signal. Recall $\lambda_0 = c/f$ where λ_0 is the signal's wavelength in air, and c is the speed of light in air. Yet if the system is also to be powered wirelessly, or by an on-chip ultracapacitor, the power consumption must be minimized which means that the switching frequency of the circuits should be low. Clearly the two strategies are in contradiction of each other and so a balance must be achieved. Another strategy for minimizing the power consumption of a communication system is to keep it simple, minimizing the number of blocks required, and minimizing (or completely eliminating) the requirement for digital filters, converters, and circuits that, in general, consume more power than well-designed analog circuits (recall Table 2.1).

With these factors in mind, this chapter proposes a completely new and novel low-power transceiver, dubbed the “lock-and-roll transceiver”, comprised of a PLL-based transmitter and a PLL-based receiver. The system makes use of an inductive integrated antenna structure and should be capable of operating from an integrated power supply [29],[30]. The author acknowledges that the implementation details of both the TX and the integrated antenna structure represent the work of his colleagues, while the implementation details of the RX are his own. The system level analysis

of the whole transceiver is the result of much collaboration between the author and his colleagues, namely Victor Karam who designed the TX and Atif Shamim who designed the antenna. All three individuals recognize the need to discuss all aspects of the project when outlining the requirements of the section they were responsible for implementing.

3.1 The Lock-and-Roll Transmitter

3.1.1 TX Overview

The lock-and-roll transmitter is based on a traditional integer-N PLL design where the circuit's VCO is directly modulated to yield a BFSK output signal. The block level diagram of the lock-and-roll TX is shown in Figure 3.1. Modulation via the

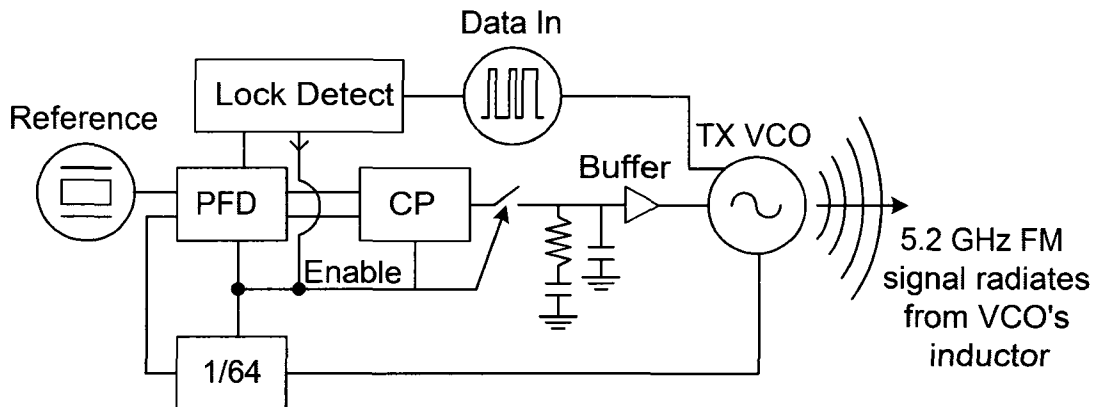


Figure 3.1: Lock-and-Roll TX Topology

VCO's control input is not a new principle [42], and a well known drawback of this approach is that the data signal experiences high-pass filtering through to the output due to the frequency response of the loop, and corruption of low-frequency data results. Modulation via the divider, typically accomplished through a $\Sigma\Delta$ controller [43] results in the opposite problem, as the input data experiences low-pass filtering through to the output and hence the data rate is limited. Alternatively, the PLL loop can be opened and then the VCO directly modulated [44]. With no PLL feedback, the VCO's output frequency is vulnerable to pulling from noise, and to frequency drift caused by charge bleeding off of the loop filter's capacitors. Previous results

[44] claim the frequency drift can be minimized at $2.5 \text{ Hz}/\mu\text{s}$ for a low-voltage VCO in a modern semiconductor process. Though this drift rate seems optimistic, the lock-and-roll TX practices this open-loop modulation principle and by design, has minimal VCO frequency drift while the loop is open. Additionally, the most novel aspect of the design is that the inductor in the core of the LC VCO doubles as the antenna for the TX, and as such the output signal radiates directly from the VCO's tank circuit without the need for a power amplifier (PA) circuit which would add to the TX power consumption considerably. Additionally, when the loop is opened all components are disabled save the VCO in order to conserve power, and the average power consumption is much improved.

Initially, the TX is powered up and the PLL locks the VCO to a multiple of the reference. In this case, the reference signal is at 81.25 MHz and the VCO is locked to 5.2 GHz, or 64 times the reference. For reduced power consumption, 6 fixed divide-by-two prescalers are cascaded to form the divider instead of a multimodulus divider (MMD) design. Channel selection would require an MMD, or the use of different reference frequencies. Once phase locked, a lock-detection circuit triggers the loop to open as shown in Figure 3.1 and the necessary control voltage for a VCO frequency of 5.2 GHz is held on the loop filter capacitor closest to the VCO. The digital bitstream containing the input data is switched onto a second VCO control line, which controls a second varactor to modulate the VCO spectrum using BFSK FM according to the data packet to be transmitted. As mentioned, the VCO inductor doubles as the antenna for the transmitter and no power amplifier is needed. During VCO modulation, power is conserved by turning off the PLL's divider, PFD and CP.

Figure 3.2 a) shows simulated results where the input bitstream is applied to the second control line of the VCO to yield the FM modulated output of the VCO, shown in Figure 3.2 d) in the frequency domain. The enable pulse which turns on all the PLL blocks to lock the VCO to the correct center frequency is shown in Figure 3.2 b), while Figure 3.2 c) shows the transient control voltage signal as the loop acquires lock. Note the different time scales on the three time domain plots in Figure 3.2, where the input bitstream is not applied until well after the loop is locked and the majority of the loop components have been disabled.

There are three aspects to the TX design that maintain the control voltage long enough to enable the transmission of data, namely the use of a unity-gain loop




Figure 3.2: Lock-and-Roll TX FM Signal Generation

buffer, the loop switch design, and the CP design. The loop buffer serves to prevent the bleeding off of charge on the loop filter during open-loop operation through the VCO's varactors. Having unity gain is important such that the characteristics of the loop are not altered. The loop switch is a standard CMOS transmission gate with dummies that prevent channel charge from the switch's transistors from altering the charge held on the loop filter at the moment the loop is opened. Due to the switch's finite impedance, the CP is disabled when the loop is opened to further prevent charge on the loop filter from bleeding off through the CP. Even with these three safeguards against loop filter charge leakage in place, the control voltage for the VCO will decrease over time with the loop held open and the rate at which this happens will ultimately dictate the number of bits that can be transmitted, for a given data rate and Δf , before the loop must be closed again and the VCO re-tuned to 5.2 GHz. Simulated results suggest that the drift rate is about $10 \text{ Hz}/\mu\text{s}$ which, at a data rate of 5 kb/sec and with a $\Delta f=500 \text{ kHz}$, allows for more than 250 bits to be transmitted between the time the loop is opened and the signal being transmitted has drifted so

far off center (from 5.2 GHz) that the RX circuit will not be able to demodulate the signal. This limitation is clarified further in section 3.4.1.

3.1.2 TX Power Consumption

To calculate the power consumption of the lock-and-roll TX circuit one must take into account the expected duty cycle of the transmitter and the amount of time the TX circuit spends in each of its two operating modes, namely the closed-loop and open-loop modes. The average power consumption can be calculated as

$$P_{TX} = \frac{P_{CL}T_{startup} + P_{OL}L_{packet}/R_{data}}{T_{packet}} \quad (3.1)$$

where P_{CL} is the closed-loop power consumption with all blocks enabled, $T_{startup}$ is the startup time (or lock time) of the PLL, P_{OL} is the open-loop power consumption with only the VCO enabled, L_{packet} is the length of the data packet to be transmitted in bits, R_{data} is the data rate, and T_{packet} is the number of packets transmitted each second. The simulated power consumption of the closed-loop PLL is about 7.5 mW while the simulated open-loop power consumption is about 4.2 mW and the simulated startup time is about 1 μ s. Therefore, to transmit a 250 bit packet once a second at 5 kb/s results in an averaged power consumption of about 210 μ W. From (3.1), one recognizes the importance of minimizing the startup time of the loop in order to minimize the proportion of time the loop must operate in the closed-loop mode, where power consumption is highest, each transmit cycle. Careful design of the loop components helps to optimize the lock time as discussed in section 3.4.3.

3.2 The On-Chip Inductive Antenna

The use of an on-chip combined antenna/inductor is elegant, economical, and well suited for short range applications like RFID tags and biomedical sensors, minimizing both the system's physical size and BOM. The on-chip antenna for the lock-and-roll transceiver is fabricated in a standard 0.13 μ m CMOS process with a low-resistance silicon substrate. The same design is used in both the TX and the RX, leading to a communication range of 1.75 m. As the antenna serves double duty as the inductor in the oscillator tank of the TX, there is an important design tradeoff between designing

a structure that is a reasonably high- Q inductor and one that is a good antenna with an appropriate radiation pattern and efficiency. As with any antenna design, the radiation resistance (R_r) should be maximized while the loss resistance (R_L) should be minimized.

A large single-turn loop antenna structure was chosen to optimize the use of chip space, leaving room for the active circuitry of the TX and/or RX to be placed in the antenna's center. An octagonal and a square loop were analyzed using Ansoft's *HFSS* [45], a 3-dimensional EM field simulator. The two antenna structures are shown in Figure 3.3 while Figure 3.4 shows the inductance and Q versus frequency based on the impedance of each structure as simulated in HFSS.




Figure 3.3: Single Turn Inductors/Antennas




Figure 3.4: Antenna Inductances and Q

Both geometries have similar dimensions with an outer diameter of 1 mm by 1 mm, metal width of 0.1 mm and a feeding gap of 0.1 mm. The antennas are fed differentially. A square loop is typically not the best choice for an on-chip inductor because it has sharp 90° bends which increase the series resistance and therefore decrease the Q when compared to an octagonal geometry, as is shown in Figure 3.4. At the same time, the sharp bends in the square loop tend to increase R_r and consequently the gain of the antenna. The square loop has a simulated gain of -22 dBi while the octagonal loop antenna has a gain of -23 dBi. Both the antennas display desirable, smooth, omni-directional radiation patterns in the horizontal plane ($\phi = 0^\circ$) and a broad double-lobe pattern in the elevation plane ($\phi = 90^\circ$) with the null on the chip edges. The octagonal loop antenna has an inductance of 2.3 nH and a Q of 11.6 at 5.2 GHz. The inductance of the rectangular loop is 2.0 nH with a Q of 9.3 at 5.2 GHz. Although the octagonal loop is a better choice for an on-chip inductor, the square loop is more suitable as an on-chip antenna/inductor because it offers 1 dB more gain. Worthy of mention is that these gains were achieved without the use of a patterned ground plane, which is a typical method of improving the Q and inductance [46], but one that decreases the antenna gain considerably.

The HFSS simulator was used to generate a two-port S-parameter output file that characterized the rectangular antenna/inductor over frequency. Using this file, Agilent's *Advanced Design System (ADS)* simulator [47] can be used to estimate an equivalent lumped element circuit model for the antenna/inductor suitable for use in circuit design simulations [48].

3.2.1 Rectangular Antenna/Inductor Equivalent Circuit Model

Figure 3.5 shows the equivalent model of the rectangular antenna/inductor that was developed [48] for the project. In the model R_s , R_{sub} , C_{ox} , C_s and L_s represent the metal coil resistance, substrate resistance, oxide capacitance, substrate capacitance and the inductance associated with the on-chip inductor, respectively.

The optimization routine in ADS was used to vary the values of the lumped model elements until good correlation was achieved between S-parameter simulations of the lumped model and the S-parameter file that was generated with HFSS for the antenna/inductor. The two S-parameter files are closely matched from 5 to 5.6 GHz




Figure 3.5: Antenna/Inductor Lumped Element Equivalent Circuit

for the final optimized circuit element values shown in Table 3.1. The differential input impedance at 5.2 GHz for the equivalent lumped element model simulates as $Z_{in,model} = 7.12 + j66 \Omega$ whereas HFSS predicts $Z_{in,HFSS} = 8.3 + j66 \Omega$. The equivalent model was used extensively for the design and simulation of the LNA (with on-chip input match to the antenna) in the lock-and-roll receiver (see Chapter 5).

Table 3.1: Optimized Antenna/Inductor Lumped Element Model Parameters

Parameter	Value
R_s	7.0 Ω
R_{sub}	4.5 k Ω
C_{ox}	915 fF
C_s	10 fF
L_s	2.0 nH

3.2.2 Antenna Efficiency

A useful parameter for evaluating the effectiveness of an antenna is the antenna efficiency (e_A), which is the ratio of radiated power to total power dissipated by the antenna. The radiation efficiency can be calculated [49] as

$$e_A = \frac{R_r}{R_L + R_r}. \quad (3.2)$$

As the simulated differential impedance of the square inductor at 5.2 GHz is $Z_{in} = 7.12 + j66.00 \Omega$, and thus $R_r + R_L = 7.12 \Omega$. Similarly, the simulated efficiency

is $e_A = 0.67$, and thus we can deduce that $R_r = 4.77 \Omega$ and $R_L = 2.34 \Omega$. Given these known parameters of the antenna, the necessary signal levels for communication between the TX and RX can be calculated.

3.3 System Link Budget and the Friis Equation

In order to ensure successful reception of the transmitted signal, an analysis must be performed that takes into account the gain of the TX and RX antennas, the power at the terminals of the TX antenna (P_T), and the free space losses to predict the power at the terminals of the RX antenna (P_R). The system design must be such that the sensitivity of the receiver is great enough to operate with P_R . The Friis equation [50] is traditionally used for this budgeting purpose. The Friis equation, assuming a conjugate match to both antennas, can be written as

$$P_R = P_T G_T G_R \left[\frac{\lambda_0}{4\pi r} \right]^2 \quad (3.3)$$

where G_T and G_R are the gains of the transmit and receive antennas respectively, λ_0 is the signal's wavelength in free space ($\lambda_0 \approx 57.7$ mm at 5.2 GHz), and r is the distance between the two antennas. Recall that the same antenna is being used in the TX and in the RX, but a typical communication system would have one transceiver making use of the on-chip antenna with a gain of -22 dBi while communicating with another transceiver that could use a 6.7 dBi patch antenna. The peak-to-peak signal swing at the terminals of the transmitting antenna is 1.0 V, and recalling R_r from section 3.2 one can calculate from (3.3) that a communication range of 1.75 m is possible so long as the receiver is sensitive enough to handle a $P_R = 235.5$ pW, or 115.8 μ V peak to peak at the terminals of the conjugately matched antenna. Similarly, if on-chip antennas are used with both transceivers the range decreases to 6.5 cm for the same received power level, and if patch antennas are used for both devices the communication range is theoretically 48 m. Figure 3.6 shows the communication range that can be achieved, assuming the receiver is sensitive enough to operate from $P_R = 235.5$ pW (proven in section 3.4.4), using different combinations of on-chip and off-chip antennas.




Figure 3.6: Communication Range vs. Antenna Configurations

3.4 The Lock-and-Roll Receiver

3.4.1 RX Overview

The receiver topology is based on a traditional third-order PLL with a static divide ratio and a second-order on-chip loop filter. Similar to the TX design, the loop can be opened and closed. This feature allows the VCO to be injection locked by the incoming FM signal, while the divider, phase-frequency detector, and a second charge pump work together to perform FM demodulation. Unlike the TX architecture where many of the PLL's blocks are disabled during open-loop mode to conserve power, only the primary CP in the loop is disabled to facilitate open-loop mode, the rest of the components serve to demodulate the input signal. The basic topology is shown in Figure 3.7.

The receiver loop is initially closed to set the center frequency of the oscillator to 5.2 GHz, which is 64 times the 81.25 MHz reference. The loop is then opened, and the oscillator is injection locked to the incoming FM modulated signal being broadcast by the transmitter. The VCO circuit is an *LC* oscillator but makes use of a standard kit inductor rather than the on-chip antenna/inductor as is the case in the TX circuit's VCO. Rather, the on-chip antenna (or optionally an off-chip patch antenna) is impedance matched to the input of a low-noise amplifier, which has a gain of 20 dB, and couples the FM modulated input into the spectrum of the initially

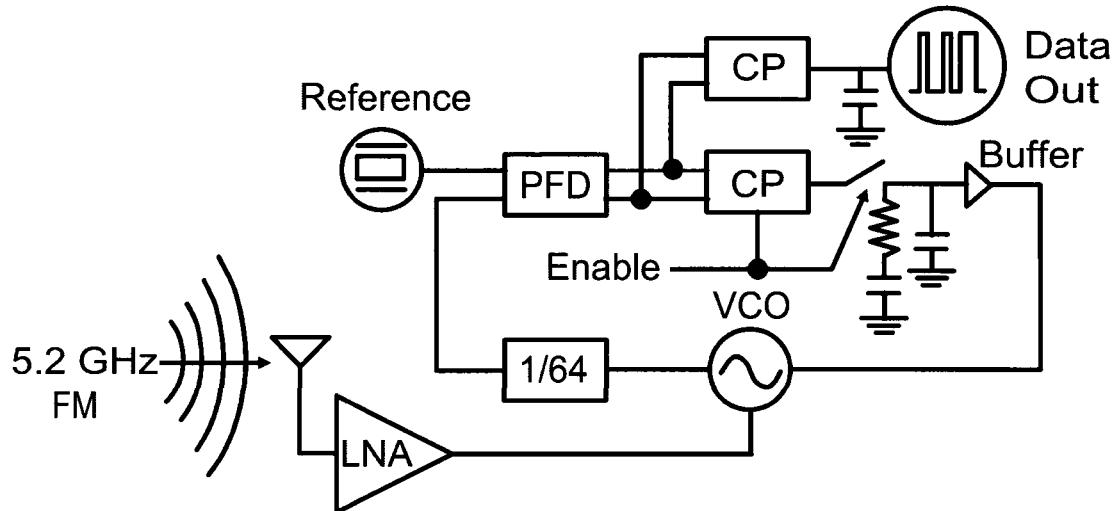


Figure 3.7: Lock-and-Roll RX Topology

free-running oscillator. If the coupled signal is strong enough and if the instantaneous frequency of the FM input is always within the locking bandwidth of the oscillator, recall (2.2), the oscillator is injection locked to the incoming signal.

At first glance the lock-and-roll receiver topology is somewhat similar in nature to a super-regenerative receiver in that it makes use of the VCO circuit in the front end as an efficient high-gain element. Yet the similarities end there, and the lock-and-roll RX, by design, has a few key advantages over the super-regenerative receiver. Firstly, the incoming signal is an FM signal using BFSK modulation. This reduces the worry of generating splatter in the frequency domain from the TX as the output is not toggled on and off as is the case with OOK modulation. Additionally, the oscillator in the RX is intentionally injection locked to the incoming signal and is therefore less likely to be pulled in frequency or to lock to an undesired input tone. Thirdly, both the TX and RX are completely integrated on chip, including the antenna and potentially the power source making the topology as compact, elegant, and cost efficient as possible.

Like the TX in this project, the RX uses a unity-gain buffer between the loop filter and the VCO to eliminate charge leakage through the VCO's varactors while in open-loop mode. The transmission gate between the primary CP and the loop filter has dummy cells that reduce the effects of channel charge injection from the

transmission gate on the loop filter's state at the moment the loop is opened, and the design of the primary CP yields a high-impedance output once the block is disabled. All these attributes help to reduce the rate of charge leakage off the loop filter while in open-loop mode, and consequently they reduce the open-loop drift rate of the VCO. The solution is substantially more power efficient than some of the previous methods, recall section 2.3.4, for overcoming the inevitable VCO drift that occurs in open-loop mode by using digital circuitry to track the drift and correct for it. With a simulated drift rate of about $10 \text{ Hz}/\mu\text{s}$, the loop must be closed periodically to re-center the VCO, yet the drift rate is small enough that 250 bits (recall 3.1.1) can be demodulated without needing to re-center the VCO.

3.4.2 RX Power Consumption

Similar to what was discussed in section 3.1.2 with regards to the power consumption of the TX, one must take into account the expected duty cycle of the transmitter and the amount of time the RX circuit spends in each of its two operating modes to estimate the power consumption of the lock-and-roll RX circuit. As before, (3.1) applies, yet the simulated open-loop and closed-loop power consumptions of the RX circuit are virtually identical at 5.5 mW because none of the major blocks can be disabled when the loop is opened. Consequently, in the case of the RX (3.1) simplifies to

$$P_{RX} = P_L \frac{T_{startup} + L_{packet}/R_{data}}{T_{packet}} \quad (3.4)$$

where P_L is the power consumption of the loop in both open and closed modes of operation. Essentially, (3.4) just accounts for the duty cycle of the RX to determine the averaged power consumption. With $P_L = 5.5 \text{ mW}$ and assuming a startup time of about $25 \mu\text{s}$ (different loop filter than the TX is assumed here), receiving a 250 bit packet once a second at 5 kb/s results in an averaged power consumption of about $275 \mu\text{W}$. Worth noting is that the calculations here and in section 3.1.2 have assumed that the discussion involves the power consumption of TX and RX circuits that are being duty cycled to lower their average power consumption. Clearly a duty-cycled TX cannot communicate with a duty-cycled RX without some sort of low-power timing/synchronizing circuit which would be difficult to implement based on

separate, non-coherent references. The assumption is that a typical communication system consist of one low-power chip operating from an on-chip power source and using an on-chip antenna with duty-cycled TX and RX circuits to communicate with another chip that operates with a standard off-chip battery and antenna and with TX and RX circuits that are not duty cycled. The two devices need only negotiate TX/RX cycles at that point which is a much simpler problem. Thus, the averaged power consumption numbers calculated in this chapter apply only to the duty-cycled circuits which ideally consume little enough power to be able to operate from an on-chip power source.

From section 3.1.2, one recognizes that lowering the startup time of the TX loop is desirable to lower the overall power consumption. On the other hand, the RX's power consumption is not really affected by the startup time, and a loop filter bandwidth that improves the overall stability of the loop is more likely preferred.

3.4.3 PLL Loop Component Selection

The startup time $T_{startup}$ is the time required for the PLL to frequency lock (to a multiple of) and then to phase lock the VCO to the reference signal. The frequency acquisition time is inversely proportional to the square of the PLL's loop bandwidth (ω_{3dB}) while the phase acquisition time is inversely proportional to ω_{3dB} [51]. Therefore, the overall startup time can be reduced by increasing the loop's bandwidth which decreases both the frequency and phase acquisition times. However, the loop bandwidth cannot be increased without bound as the loop will eventually become unstable. The critical reference frequency $f_{ref,unstable}$ at which instability occurs is given by

$$f_{ref,unstable} \approx \zeta(1 + \sqrt{2})\omega_{3dB} \approx \zeta\omega_n \quad (3.5)$$

where ζ is the loop's damping constant (assumed to be < 1.5) and ω_n is the loop's natural frequency [51]. Allowing for a sufficient safety margin over instability, the reference frequency of the PLL is typically a minimum of 10 times larger than $f_{ref,unstable}$ and therefore

$$f_{ref,min} \approx 10\zeta\omega_n. \quad (3.6)$$

A larger loop bandwidth (and hence a larger natural frequency) requires a proportional increase in the reference frequency in order to provide the same margin over

instability. However, increasing the loop bandwidth results in a decrease of the loop filter's capacitances (C_1 and C_2) which are given by

$$C_1 = \frac{I_{CP}K_{VCO}100\zeta^2}{2\pi f_{VCO}f_{REF}} \quad (3.7)$$

and

$$C_2 = \frac{C_1}{10} \quad (3.8)$$

where I_{CP} is the charge pump current, K_{VCO} is the VCO gain in $[\frac{\text{rad/s}}{\text{V}}]$ and F_{VCO} is the VCO output frequency. If we assume that 300 fF is the smallest capacitor that can be accurately integrated on chip, then the resulting minimum C_1 is 3 pF. With these capacitances, an I_{CP} of 100 μA and a K_{VCO} of 250 MHz/V, solving for the reference frequency in (3.7) results in $f_{ref} \approx 80$ MHz. Recall from section 3.1.1 and 3.4.1 that a reference frequency of 81.25 MHz is being used for the lock-and-roll transceiver.

Although a wide loop bandwidth decreases the startup time of the PLL, the phase noise response of the output signal will not be optimal. As a safeguard, the loop filter on the lock-and-roll TX and RX test circuits is made tunable (via laser trimming) such that the resistor and capacitor values can be changed during testing as needed to yield optimal measured results. As laser trimming is expensive, this practice would only ever be used on a test chip to fine tune the design such that in a production situation the desired loop filter bandwidth would be fixed at the frequency that was determined by trimming and measuring the test chip.

3.4.4 RX VCO's Injection Locking Bandwidth

Recall from section 3.3 that the communication ranges highlighted in Figure 3.6 assumed that the RX circuit could injection lock to an incoming FM signal with a $\Delta f = 500$ kHz at peak-to-peak signal level of 115.8 μV at the antenna terminals. Additionally, recall from section 2.3.1 that the injection locking bandwidth of an oscillator is given by (2.2). The RX oscillator has a free-running differential peak-to-peak swing of 1.0 V and a tank inductor with $Q = 5$. This Q is a result of intentionally de-Q'ing the tank of the VCO using parallel resistance in order to optimize the locking bandwidth for the ± 500 kHz input signal, i.e., the signal switches between 5.1995 GHz

and 5.2005 GHz. On the test chip, the parallel resistance is implemented such that it is laser trimmable in order to optimize the tradeoff between the locking bandwidth and the VCO's power consumption. With a gain of 20 dB from the RX LNA, the injected signal into the VCO will have a peak-to-peak swing of 1.16 mV. From (2.2), the locking bandwidth of the oscillator (in Hertz) with this injected amplitude is $f_L \approx 602$ kHz. Thus, the FM modulated input which toggles between $f_0 - 500$ kHz and $f_0 + 500$ kHz is always within the locking range of the receiving oscillator by design. Figure 3.8 shows the overlay of three separate simulation results conducted with the LNA and VCO circuits in the RX in order to verify that the locking bandwidth is adequate to cover the 5.1995 GHz to 5.2005 GHz range.

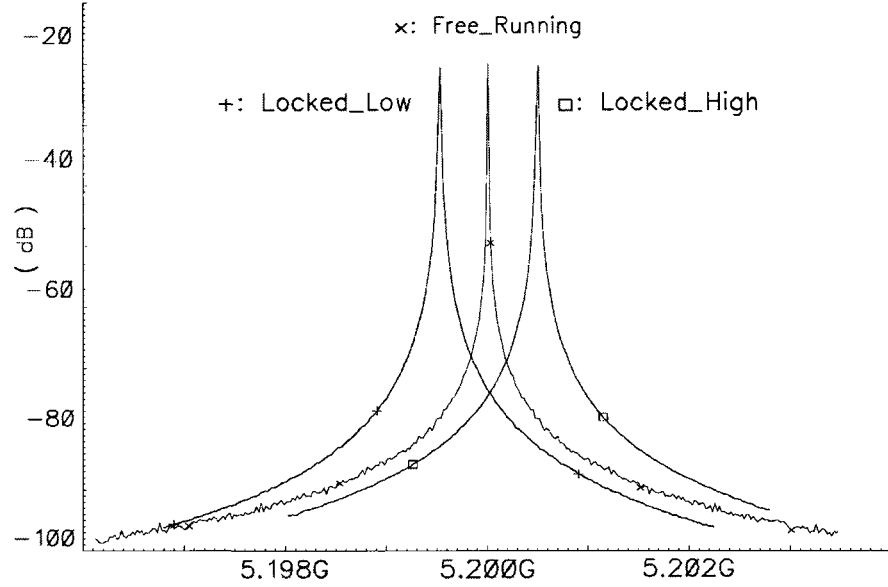


Figure 3.8: Coupled LNA/VCO Locking Bandwidth Verification

3.5 FM Modulation Considerations

The frequency separation, Δf , of the modulated output must be small enough that the VCO in the receiver can injection lock to any instantaneous frequency being transmitted, yet a wide Δf is desirable because it makes for a more distinguishable frequency separation at the input to the PFD in the receiver. Similarly, while a high modulation frequency, f_m , enables higher data rates, f_m must be low enough

with respect to the reference frequency that the phase-frequency detector (PFD) and charge pump (CP) in the receiver have enough time to deduce a 1 or a 0. Here, $\Delta f = 500$ kHz and $f_m = 5$ kHz. The tradeoffs between f_m , Δf and communication range are definitely worthy of careful consideration.

3.6 Tradeoffs Between Power, Range, and Data Rate

In open-loop mode when the RX is demodulating the BFSK input signal there is a frequency shift of the injection-locked VCO signal, corresponding to a transition in the bitstream, and the time required for the receiver loop to demodulate is dependent on the phases of the inputs to the PFD, namely the reference (f_{ref}) and the divider output (f_{div}). In open-loop demodulation, the PFD will behave as a frequency detector/comparator because f_{div} is either higher (representing a logic high) or lower (representing a logic low) than f_{ref} . The solid line in Figure 3.9 illustrates the PFD's behavior seen at the output of the secondary CP for the case when $f_{ref} > f_{div}$. The initial phase difference is assumed to start at point A. The phase difference ($\theta_{ref} - \theta_{div}$) increases with time, passing through points B and C until the phase difference equals 2π (or multiples thereof) where a cycle slip occurs because the phases of f_{ref} and f_{div} are aligned (point D). This pattern repeats without bound towards point Z.




Figure 3.9: Secondary CP Output vs. Phase Difference at the PFD

If there is a transition in the bitstream, then $f_{ref} < f_{div}$ and the time required for the CP to begin reversing the direction of current flow through the loop filter

depends on the phase difference at the PFD inputs. If the transition happens when the phase difference is 0 radians (or multiples of 2π), indicated by point A (or point D), then there is no delay before the CP begins reversing current. If, however, the transition happens when the phase difference is slightly lower than 2π (point C), then the phase difference would begin decreasing, passing through point B and onto point A where the CP finally begins reversing the current. The worst case delay until the current reverses occurs when the phase difference is slightly lower than 2π at the time of a bit change, and the length of the delay is inversely related to the beat frequency between f_{ref} and f_{div} . The period of this beat frequency is also the time between cycle slips. Here we have used a reference frequency of 81.25 MHz in the receiver, a data rate of only 5 kb/s, and a $\Delta f = 500$ kHz. These settings were chosen such that the maximum delay between a bit change (and corresponding frequency change) at the input to the PFD in the receiver and the resulting bit change at the output of the receiver's CP is ≈ 50 % of the bit length. This result might seem overzealous, but note that there is only a small chance that the delay will be that bad for any one bit, and that the data rate can always be decreased in testing if the output signal integrity suffers. The beat period at the input of the PFD is given by

$$T_{BEAT} = \frac{1}{\Delta f/N} = \frac{1}{500 \text{ kHz}/64} = 128 \mu s \quad (3.9)$$

Increasing Δf would clearly decrease this wait time, but as (2.2) suggests, this would require greater received power to keep the receiving VCO injection locked. Recalling (3.3) we see that P_R can be increased with a decrease in range, or increased antenna gain. Thus there are many tradeoffs that can be made. Figure 3.10 summarizes the relationship between the data rate, the locking bandwidth, and the communication range of the system.

3.7 Lock-and-Roll Transceiver Summary

The proposed lock-and-roll transceiver is unique in that operates at 5.2 GHz and is therefore capable of communicating using an on-chip antenna, but by design both the transmitter and the receiver are low-power, so much so that when duty cycled under normal operating conditions, the system can potentially be powered using an on-chip




Figure 3.10: Data Rate, Locking Bandwidth, and Communication Range Tradeoff

power source. The transceiver might very well be the first design to enable the use of an on-chip antenna and an on-chip power source simultaneously.

The TX circuit is based on an integer-N PLL which is unique in many ways. Once locked, the loop can be opened by way of a carefully designed transmission gate, and the control voltage held near constant while a unity-gain buffer sources the leakage current drawn by the VCO's varactors. Once the loop is open, the PFD, CP, and the divider are disabled to lower the power consumption and the VCO is directly modulated to yield a BFSK FM output. Additionally, the inductor in the core of the *LC* VCO serves double duty as the antenna for the transmitter, eliminating the need for a power amplifier (PA) as the output signal radiates directly from the VCO.

The RX circuit is also based on an integer-N PLL that can be opened once the loop is locked. Unlike the TX circuit where the PFD and the divider are disabled to save power, the RX circuit uses the fundamental components of the PLL to demodulate the FM input signal. The on-chip antenna is conjugately matched to the input of an LNA with 20 dB of gain which injection locks the VCO when the loop is open. The PFD and a second CP compare the modulated output of the divider to the reference and produce a rail-to-rail output signal that is low when the incoming signal is lower than 64 times the reference frequency, and high when the input frequency is higher than 64 times the reference.

By design the system has a communication range of 1.75 m at 5 kb/s when one chip using the -22 dBi on-chip antenna communicates with another chip using a 6.7 dBi patch antenna. This communication range can be increased at the expense of the data rate.

The implementation details of the TX circuit represent the work of the author's colleague, Victor Karam, and are the subject of a different thesis [52]. Similarly, the details of the antenna design pertain to the thesis of another colleague, namely Atif Shamim. The details of the RX circuit's design represent the work of the author and are the subject of Chapters 4 through 7 of this thesis.

Chapter 4

Injection-Lockable VCO Design for Low-Power Applications

The standard application for a VCO is to serve as a frequency (and or phase) reference circuit. In heterodyne transceivers that communicate at RF frequencies and convert signals from one frequency to another, the output of an *LC* VCO is commonly applied to the input of a mixer that up-converts or down-converts the message signal to a predetermined frequency. As such, the VCO output must be accurately controlled. The output frequency (and or phase) will determine the frequency (and or phase) of the converted signal, and in most cases the VCO's output amplitude will partially determine the output amplitude of the converted signal due to the finite gain of the mixer. Phase noise, which is related to the *Q* of the circuit, will translate to phase noise on the converted output of the mixer, while low VCO output amplitude may translate to a weak mixer output. Weak output from a transmitter's frequency up-conversion mixer (FUM) can result in lower output power delivered to the antenna due to the finite gain of the PA. Weak output from a receiver's frequency down-conversion mixer (FDM) can result in lower signal-to-noise ratio (SNR) at the input of the ADC that drives the digital baseband. In each case, the end result is usually a negative impact on the link margin of the system, ultimately limiting the communication range and rendering the radio more susceptible to interference. In most cases the design should be current limited (as opposed to voltage or headroom limited) to improve the oscillator's isolation from supply noise and to maintain good linearity. In summary, traditional VCO design strategies often strive to achieve maximum *Q* and high output swing, where the VCO is most usually current limited.

Nontraditional uses for an *LC* VCO, however, can have somewhat opposing requirements. Recall from section 2.3.1 that an oscillator can be viewed as a narrow-band filter with high gain. When the oscillator is used as a precision reference, the bandwidth of the filter should be narrow to decrease phase noise (equivalent to having higher Q), while the gain should be high to guarantee startup and maximize the linear output swing. Maximizing both Q and gain results in a circuit that starts up readily, and is not easily pulled off center by interfering signals. However, for applications that require the oscillator to be easily injection locked over a broad band (as is the case when the oscillator is to lock to a modulated signal as in the lock-and-roll receiver) the Q should be decreased to widen the response of the filter. Additionally, oscillators that have high gain are difficult to injection lock as their tendency to oscillate at their free-running center frequency is strong and must be overcome by additional injected signal power. As a result of these characteristics, injection-lockable oscillator designs must achieve a balance between gain and Q , such that the circuit still oscillates, but can be readily injection locked.

This chapter explains and demonstrates design strategies for achieving injection-lockable oscillators that are simultaneously optimized for locking bandwidth and for power consumption. The fundamental topics of oscillator design are reviewed initially, covering concepts such as the necessary conditions for oscillation, oscillator gain margin, noise considerations, resonant tanks and loaded vs. unloaded circuit Q , to provide a basis for comparison of different design strategies. Many of the fundamental concepts relating to oscillator design are covered in more detail in Appendix A, allowing for the primary focus of this chapter to be design strategies for meeting the specific requirements of the VCO in the lock-and-roll receiver.

Unique strategies that achieve low-power, lockable oscillators are discussed, and the oscillator in the core of the lock-and-roll receiver is presented as a design example, showing simulated and measured results. Key aspects of the layout that were implemented to maintain symmetry and to make the test circuit laser trimmable (fixed capacitance in the load, de- Q 'ing resistors, etc.) are highlighted, along with layout tactics that reduce the effects of parasitics on the 5.2 GHz oscillation frequency.

4.1 A Review of Oscillator Design Fundamentals

Regardless of the application being targeted when designing a VCO, one must consider a number of fundamental criteria to guarantee a successful design. The Barkhausen criteria, which is referenced by nearly all authors that broach the topic of oscillator design [46], defines the necessary conditions for oscillation.

4.1.1 The Barkhausen Criteria and Gain Margin

The Barkhausen criteria (which is derived in section A.1 of Appendix A) states that for sustained oscillations to exist at any particular frequency, the gain around the loop must be unity and the phase must be zero or a positive integer multiple of 2π .

At first glance, this academic and somewhat oversimplified oscillator analysis is often difficult to translate into practical *LC* oscillator design strategies, giving rise to questions such as: “What causes oscillations to grow in the first place at startup?”; “What causes them to saturate at steady state?”; and “What, exactly, is it about the circuit that controls the frequency of oscillation?” Suspending discussion on the third issue momentarily, the issues of oscillation growth and saturation are dictated by the gain of the oscillator. While the Barkhausen criteria outlines that a loop gain of unity is necessary for sustained oscillation, in fact real oscillators have nonlinear gain characteristics that result in small-signal loop gain of much more than unity, which causes oscillations to grow.

Having an initially high loop gain is a desired, and necessary condition for starting the oscillation process, so much so that the extent to which this requirement is satisfied by a particular design has been given its own term – and is regarded as the “gain margin” of the design. For the sake of discussion, one can consider a slightly more realistic view of the oscillator than the simple feedback block diagram (shown in Figure A.1) to quantify gain margin.

Using a simplified model of the popular $-G_m$ oscillator, one can demonstrate that if $G_m \gg 1/R_L$ (where G_m is the net transconductance of the oscillator and R_L is the net impedance at resonance), then the loop gain is greater than unity and the condition promotes oscillations, if not then the loop gain is less than unity and oscillations will decay and eventually die out. This topic is discussed further in section A.2 in Appendix A. As the only initial input to an oscillator at startup is

thermal noise, the loop gain must be greater than unity at some frequency in order for oscillations to develop, and the amount by which the gain exceeds unity (or 0 dB) is called the gain margin of the design.

4.1.2 Current-Limited vs. Voltage-Limited Oscillator Designs

Striving for a loop gain well in excess of unity would appear to violate the Barkhausen criteria which claims that the loop gain must be exactly unity for sustained oscillations, though in reality what happens with integrated LC oscillators is that the gain starts out much higher than unity at startup which causes thermal noise to be amplified (and correlated at a frequency determined by the LC resonance, see section 4.1.4) such that oscillations grow. Eventually, the growing signal swing across the gain devices causes one of two scenarios to occur, either the circuit is current limited (traditionally by a tail current, mirroring device) in which case the gain around the loop will saturate to unity because the gain devices are not afforded the required average current to sustain signal growth, or the circuit is voltage limited and the output signal will start to clip, and the gain devices will fall out of saturation for a period of every cycle where their gain is much lower, and thus the average gain saturates to unity once again. The first case is typically much better than the second from a noise and linearity point of view.

Maximizing Isolation from Noisy Supplies

The output impedance of a CMOS transistor is dependent on the operating point of the device (further discussed in section A.2.1 of Appendix A). Transistors have higher output impedance when operated in saturation than they do in the triode region and for this reason, isolation from noisy supply voltages is maximized when a design is current limited such that the gain devices never fall out of saturation. Consider the simple NMOS $-G_m$ oscillator shown in Figure 4.1.

Looking into the drain of either NMOS device (M1 or M2) in Figure 4.1, the impedance from the output to VSS (a noisy on-chip supply in most cases where the VCO is implemented on the same chip as many other RF blocks such as in a transceiver) will therefore be maximized from a ground noise isolation point of view




Figure 4.1: NMOS, $-G_m$, LC Oscillator

if the devices are kept in saturation, and thus a current-limited design is preferable in terms of achieving maximum isolation from supply noise. The very same argument holds true with regards to PMOS devices with sources tied to VDD. Many truly complementary CMOS VCO designs, such as that used in the lock-and-roll receiver, make use of both PMOS and NMOS devices to maximize the gain margin for a given supply current, and thus striving for a current-limited design maximizes isolation from both the VDD and VSS supplies on chip.

4.1.3 Designing for Minimal MOSFET Noise Contribution

Of additional concern to the integrated oscillator designer is the noise contributed by the transistors themselves. In the case of CMOS transistors, there are three sources of noise that are typically discussed when reviewing the subject, namely thermal noise, shot noise, and flicker noise. A review of each type of noise, including the theoretical source of each noise in a MOSFET device and some common equations used for predicting CMOS transistor noise can be found in section A.3 in Appendix A.

The important conclusion to be noted from section A.3 for the purpose of designing low-noise oscillators is that longer channel devices exhibit less thermal noise while longer and wider devices exhibit less flicker noise. Of course, the gain margin is proportional to the width to length ratio and larger devices have higher gate (fixed) capacitance that will reduce tuning range and so the overall theme of balancing the design tradeoffs is evident once again.

4.1.4 *LC* Resonance and Tank *Q*

Up until this point, the discussions in this chapter have made reference to the “tank” of the *LC* oscillator and to its *Q* without addressing the remaining question that was posed in section 4.1.1: “What is it about the *LC* oscillator that controls its frequency of oscillation?” In fact it is the *LC* resonant circuit itself, also known as the tank circuit, that controls the frequency of oscillation and to a large extent the amplitude of the oscillator output depending on the *Q*, or quality factor, associated with the tank.

The frequency at which the inductive and capacitive admittances cancel sets the free-running frequency of oscillation for the circuit according to

$$\omega_0 = \frac{1}{\sqrt{LC}} \quad (4.1)$$

where ω_0 is known as the resonance frequency of the tank.

Section A.4 in Appendix A examines the concept of *LC* resonance in detail. Additionally, section A.4 reviews the concept of tank *Q*, differentiating between the unloaded and loaded *Q* of the circuit. The unloaded *Q* (sometimes represented as Q_U) is the *Q* of the tank, alone, unloaded by the transconductance of the VCO circuit. The term loaded *Q*, or Q_L , is often used to describe the relationship between the bandwidth of the overall VCO (when the tank is loaded by the transconductor) and the resonant frequency ω_0 . As Q_L is representative of the bandwidth of the complete oscillator circuit, it dictates the phase noise of the circuit which is an important measure by which most VCO designs are traditionally judged.

4.1.5 Phase Noise, Oscillator Pulling, and Injection Locking Bandwidth

The phase noise of an oscillator is a measure of just how clean the oscillator output is from a frequency domain point of view. As the discussion in section 4.1.4 highlights, the frequency of oscillation is dictated by (4.1) (where it should be noted that C in this equation represents the total capacitance seen across the VCO's output terminals, including all parasitic capacitances from routing, transistor gate capacitances, fixed capacitors and varactors) but of course this equation defines the point of maximum tank impedance where the impedance is not exactly zero at all other frequencies. As mentioned in section 4.1.4, at frequencies far above ω_0 the impedance of the capacitor dominates the parallel tank circuit causing the frequency response of the parallel network to roll off towards zero impedance at a rate of about 20 dB/dec. Similarly, the inductor dominates at frequencies much lower than ω_0 and a similar roll off is seen. Close in to ω_0 however, the tank impedance is either slightly capacitive or slightly inductive with a large real component (contributed by R_p) and energy present in the tank at small frequency offsets contributes to phase noise.

In simple terms, the only input signal to the integrated oscillator is white noise (or broadband energy) which is amplified by the relatively broadband response of the transistor gain elements, and filtered by the resonant tank, which even if high- Q , has finite selectivity resulting in phase noise. Many publications provide excellent explanations and insight into the generation of phase noise, [10], [12], [53], where nearly all agree with the simple model representation for the frequency response of an oscillator as shown in Figure 2.4. The generally accepted, albeit simplified (as it doesn't apply at very small offsets from ω_0 as it predicts infinite phase noise), equation for predicting an oscillator's phase noise power at offset $\Delta\omega$ [12], [46] is

$$PN = \left(\frac{|A(s)|_{\omega_0}}{2Q_U \Delta\omega} \right)^2 \frac{FkT}{2P_s} \quad (4.2)$$

where $|A(s)|$ is the oscillator's forward gain (recall Figure A.1, typically close to unity at resonance) and P_s is the power in the oscillator.

RF VCO designs are generally judged on how low their measured phase noise profile is because this, as shown in Figure 2.4, is a measure of how small their loaded bandwidth B_L is and how low their gain is to energy at any frequency away from ω_0 . In

general, a low B_L is a good thing because the energy sloshing around in the oscillator is tightly correlated at ω_0 and interfering signals or noise that are injected into the tank from noisy supplies or elsewhere are attenuated by the dominant oscillation. Yet even the highest Q oscillator can be overwhelmed by an interferer signal at ω_{inj} . At low injected power levels the injected tone appears as a sideband spur within the output signal profile of the VCO, subject to the attenuation of the loaded filter at that particular frequency, with intermodulation harmonics at multiples of $\Delta\omega = |\omega_0 - \omega_{inj}|$. As the power of the injected tone is increased what inevitably happens is that the oscillation frequency is pulled away from ω_0 and towards ω_{inj} as the dominant correlation is being challenged. If the injected power is increased even further what eventually happens is that the transconductance is correlated to ω_{inj} and the gain response of the VCO changes completely, adopting the phase noise profile of the injected signal itself. At this point, the oscillator is said to be “injection locked” to the tone at ω_{inj} . This progression, from weak injection through to oscillator pulling, and eventually oscillator injection locking has been studied at some depth by the author and others, [53], [9], [8], [14], [13], [15], through simulation and measurement. Figure 4.2 through Figure 4.4 show this progression as simulated by the author with a CMOS oscillator that has a free-running oscillation frequency of 9.94 MHz. In each case a tone was injected at 9.82 MHz, but with increasing power levels. Figure 4.2 shows that because the injected tone is only 120 kHz away, even at relatively low injected power levels the oscillator is starting to be pulled away from its 9.94 MHz free-running frequency.




Figure 4.2: 9.94 MHz Oscillator with Weak 9.82 MHz Injection




Figure 4.3: 9.94 MHz Oscillator with Strong 9.82 MHz Injection

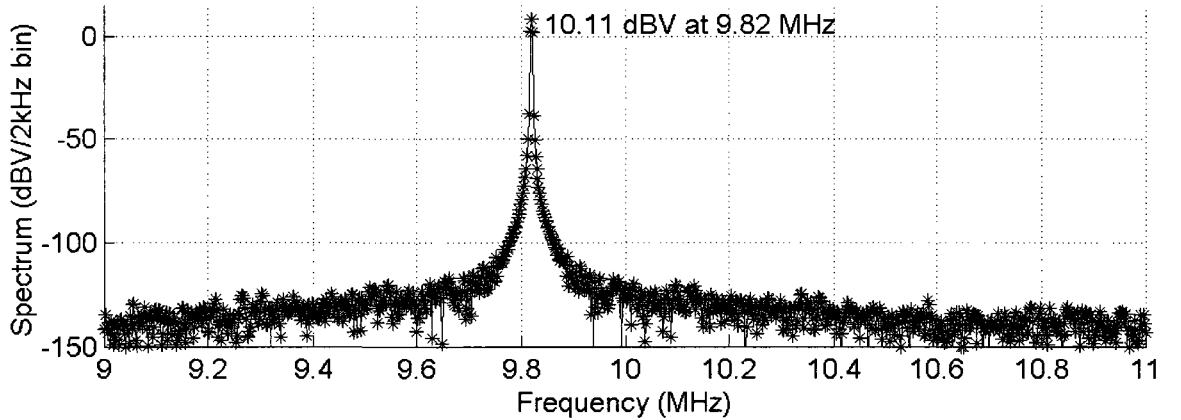


Figure 4.4: 9.94 MHz Oscillator Injection-Locked at 9.82 MHz

An interesting observation can be drawn while studying Figure 4.2 and Figure 4.3 in particular. In fact the intermodulation spurs on the opposite side of ω_0 are higher in amplitude, despite the low-side injection, and the overall spectrum is quite asymmetric. The author studied this phenomenon in depth and came up with a theory supported by a simplified model with simulation results that showed the snap-back effect of the pulled oscillator's frequency transition in time [9]. This analysis, while interesting and worthy of reference in the context of discussing oscillator injection locking, is tangent to the focus of this thesis and will not be discussed further here. However, what is very relevant to the functionality of the lock-and-roll receiver is the mathematical expression (which inevitably results from studying the pulling

and locking effects on oscillators) for the injection-locking bandwidth of the oscillator. There are two classical formulas used to define the injection-locking bandwidth of an oscillator. The first was proposed by Adler [15] in 1946 while the second was proposed by Kurokawa [54] in 1973. The approaches have been compared [55], and the general consensus is that Kurokawa's equation predicts a larger bandwidth than Adler's equation and is more accurate for optical oscillator circulators, while Adler's equation is more pessimistic (predicts narrower locking bandwidth) and is more accurate when considering electrical oscillators. A major difference between the two approaches lies in their consideration of the loaded Q_L . The author has seen good correlation between the bandwidth predicted using Adler's equation and simulated and measured results and so this is the equation that was used to design the VCO in the lock-and-roll receiver. Adler's equation is introduced in section 2.3.1 but is repeated here given how fundamental it is to the design of the VCO in the lock-and-roll receiver. Adler predicts

$$\omega_L \approx \frac{\omega_0}{2Q_U} \frac{V_{inj}}{V_{osc}} \quad (4.3)$$

where ω_L is the single sided locking bandwidth, i.e., the oscillator can be locked from $\omega_0 - \omega_L$ to $\omega_0 + \omega_L$, and (V_{inj}) is the amplitude of the injected signal while (V_{osc}) is the amplitude of the free-running oscillator.

As mentioned, when the oscillator is injection locked it adopts the phase noise profile of the injected tone. This occurs because the oscillator always maintains constant output power, and as the injected power increases, the oscillator gain and loaded Q_L drop [53]. Both the injected tone and the noise see decreased gain, and this directly results in reduced phase noise. Measurements of the phase noise of an integrated, 3.7 GHz CMOS VCO test chip (both free running, and injection locked) shown in Figure 4.5 confirm this behaviour.

The oscillator measured in Figure 4.5 shows a clear 20 dB improvement in phase noise when injection locked relative to its free-running state.

4.1.6 Oscillator Fundamentals Summary

The typical role of an integrated RF VCO (as in a heterodyne type receiver) benefits from a design that has ample gain margin so as to guarantee startup over temperature and process variation, while maximizing signal swing to optimize the output signal

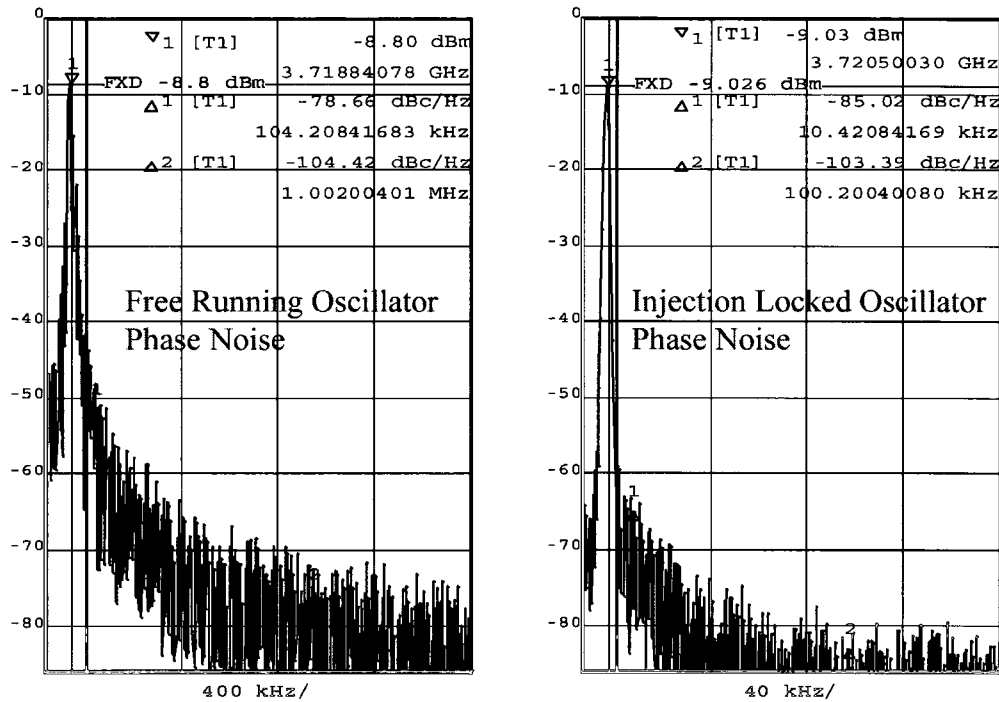


Figure 4.5: Free-Running and Injection-Locked Oscillator Phase Noise

of the mixer circuit that follows. A current-limited design that keeps the transistors in saturation optimizes supply isolation, and making use of a high- Q resonant tank circuit results in a low phase noise profile, and a VCO that resists injection locking. When it comes to designing a VCO that is to consume as little power as possible while guaranteeing a large (but definable) injection-locking bandwidth for use in the lock-and-roll receiver, in fact the considerations are all the same as those encountered when approaching a traditional design, the tradeoffs are just balanced differently.

4.2 VCO Design for the Lock-and-Roll Receiver

Figure 4.6 shows the schematic of the VCO that was designed to meet the requirements of the lock-and-roll receiver. Recall from section 3.3 that the primary requirement of the VCO, beyond the obvious requirement that it start up and oscillate readily at the desired communication frequency of 5.2 GHz, is that it be able to injection lock to a small enough signal to enable the desired communication range of 1.75 m for the targeted application. According to the calculations in section 3.4.4,

meeting this communication range requires that the oscillator be able to injection lock to an injected signal with amplitude of 1.15 mV anywhere between 5.1995 GHz and 5.2005 GHz. As such, Adler's equation was considered early in the design process. Beyond these requirements, current is to be reduced as much as possible given the self-powered applications that are targeted.

4.2.1 Advantages of the Complementary Differential Topology

The topology that was chosen is a cross-coupled complementary (uses NMOS and PMOS devices) $-G_m$ oscillator as shown in Figure 4.6. The circuit is arguably perfectly symmetrical and the differential topology provides excellent rejection to common-mode signals that might be present on the supplies given the complexity of the final chip. Using both PMOS and NMOS devices maximizes the gain margin of the circuit for a given supply current as the current drawn by tail transistor M3, which regulates the bias current to the VCO, will always be drawn through both NMOS and PMOS gain stages – a sort of current reuse.

The drawback of this topology compared to a VCO that has only NMOS or PMOS gain elements is that the maximum output swing that can result, without leaving the current-limited regime, is lower because the signal swing at Out_p and Out_n must not exceed $VDD - VDS_{Sat}$ or drop below $2 * VDS_{Sat}$. A quick glance back at Adler's equation (4.3) shows that as the locking bandwidth is inversely proportional to V_{osc} , a large output swing is not desirable for this application anyway and so this tradeoff is wisely made given the requirements.

One might ask why a current-limited design is important if indeed the chosen topology is relatively immune to supply noise given the differential circuit, but in fact the differential arrangement is really only of benefit to low-frequency noise signals. When the PMOS and NMOS pairs of transistors are switched during oscillation the common mode signal, if low in frequency relative to the oscillation frequency, will appear relatively equally at Out_p and Out_n and the differential result is near zero, but as the noise on VDD or VSS approaches or exceeds the frequency of oscillation the sampling effect of the oscillator results in translations through to Out_p and Out_n

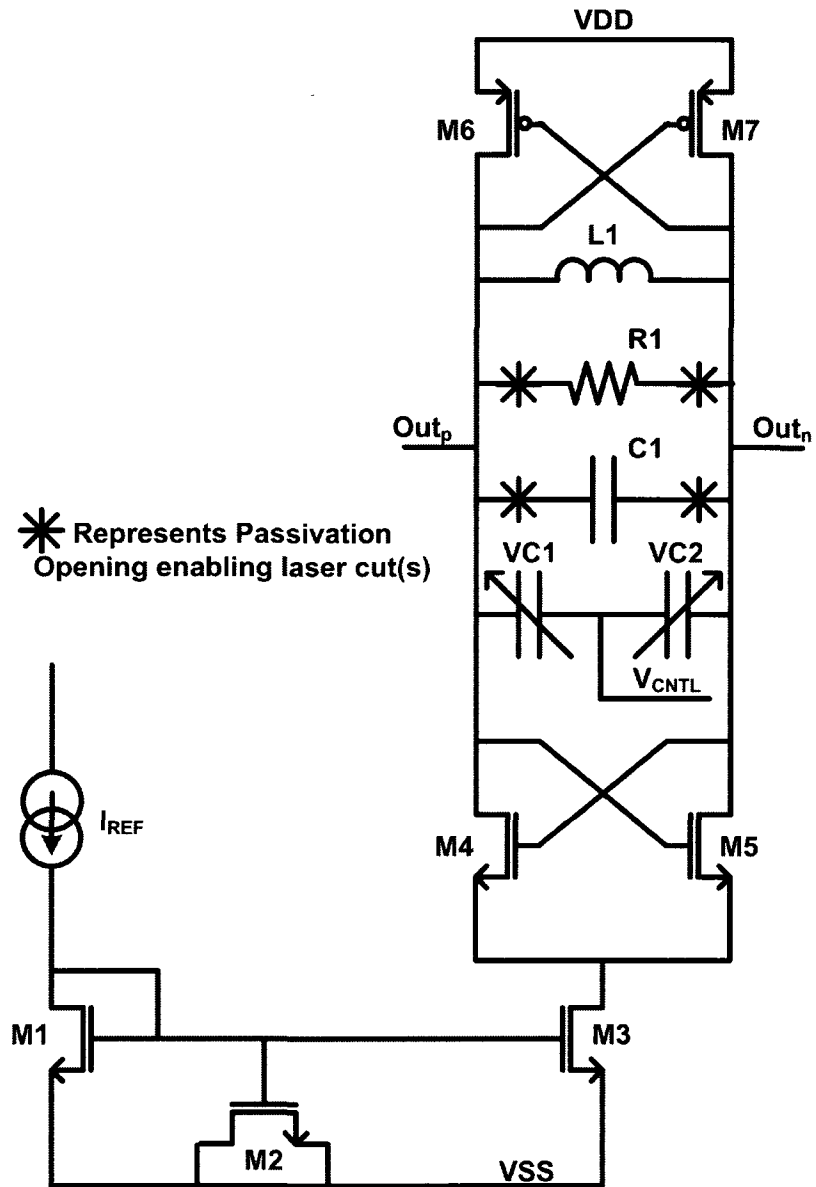


Figure 4.6: Lock-and-Roll Receiver VCO Schematic

that do not cancel differentially, and indeed the added isolation provided by devices operating in saturation (recall section 4.1.2) is of benefit.

4.2.2 Inductor Selection – Minimizing Process Variation Regardless of Q

A common approach to integrated VCO design when working in a technology with reliable inductor models is to simulate all inductor sizes available and to pick the inductor that simulates to have the highest Q ; assuming, of course, that layout size and the actual inductance value are also suitable for the design. This approach usually maximizes tank Q (recall that integrated inductor Q typically dominates the tank Q) which results in the best overall phase noise (recall equation (4.2)). In general, center-tapped (or differential) inductors have lower Q , given the lower level metals (which are typically thinner and higher impedance with more parasitic capacitance as they are physically closer to the substrate and underlying circuitry) and via farms that must be used to allow for routing overlap, than the traditional single ended (two terminal) alternative. The complementary VCO topology is beneficial in this regard because unlike the simpler NMOS circuit shown in Figure 4.1 that requires either a differential center-tapped inductor (lower Q) or two separate inductors (would nearly double the overall layout area for the whole VCO), it can make use of a high- Q single-ended structure. This subtlety being noted, while the lock-and-roll receiver's VCO benefits from the tighter layout enabled by the single-ended inductor that was chosen, the higher inductor Q is actually irrelevant in this case because the circuit was intentionally de- Q 'd to increase the locking-bandwidth. In fact a 2.0 nH single-ended kit inductor was used and the selection was made merely based on process corner and temperature simulations which suggested that this inductor had the smallest variation over all variables. The nominal simulated Q of the inductor alone was about 11.2 given the thick top metal of the process, yet a bank of five resistors (represented as R1 in Figure 4.6) were connected in parallel with the tank (forming a resistance of 540 Ω) intentionally de- Q 'ing the overall tank to about 5, and resulting in a net $R_{p,eff}$ of about 415 Ω .

4.2.3 Transistor Sizing – Trading-Off High Gain Margin and Increased Tunability for Higher Phase Noise

As the discussion in section 4.1.3 concludes, all other considerations aside, to minimize transistor noise contributions transistors should be made long and wide. Heeding this advice to extremes, however, will result in very large gate capacitances which will limit the tunability of the oscillator as they will overpower the variable capacitor (or varactors) in the tank (VC1 and VC2 in Figure 4.6). As the tank Q of the VCO was intentionally lowered (as necessary for achieving the required locking bandwidth), the gain transistors M4, M5, M6, and M7 were kept relatively short (about twice the minimum gate length) and wide in order to achieve the required $g_{m,eff}$ for a gain margin of about 10 dB.

For reference, the complementary topology shown in Figure 4.6 has an equivalent transconductance [56] of approximately

$$g_{m,eff} = \frac{g_{m,n} + g_{m,p}}{2} \quad (4.4)$$

where $g_{m,n}$ and $g_{m,p}$ are the individual transconductances of transistors M4, M5, and M6, M7 respectively. The transistors were increased in width/length ratio until DC operating point simulations suggested that $1/g_{m,eff} \approx 3R_{p,eff}$ for a bias current of about 1.0 mA, (resulting in a gain margin of about 10 dB), where $R_{p,eff}$ is the net parallel resistance of the tank which can be estimated as

$$R_{p,eff} \approx R_{p,ind} || R1 = (Q_{ind}\omega L) || R1 \quad (4.5)$$

Terms $R_{p,ind}$ and Q_{ind} in (4.5) are the equivalent parallel resistance and Q of the inductor alone, respectively. In this case $R_{p,eff} \approx 415 \Omega$ which results in wide, short transistors. If the length of the devices had been made much larger than twice the minimum gate length in favour of reducing noise, and the width/length ratio kept constant to maintain the desired gain margin, the tuning range would have been unacceptably low given the large, fixed gate capacitance. Worth noting is that, as with CMOS inverter design, the PMOS devices are sized at roughly three times the width/length ratio of the NMOS devices such that $g_{m,p} \approx g_{m,n}$. Also worthy of note is that the transistor size and the bias current of 1.0 mA could have been traded off

against each other to result in higher current consumption with the benefit of smaller devices that contribute less fixed capacitance to the tank circuit, and vice versa, all the while maintaining the gain margin of 10 dB. The 1.0 mA bias condition was found to achieve a good balance. Clearly tradeoffs were made in sizing the transistors in the design to benefit gain margin (important given the intentionally low tank Q) and tunability while sacrificing noise performance. In fact the simulated phase noise of the VCO was as bad as -99 dBc/Hz at an offset of 1 MHz from the carrier (a fantastic result given this VCO's intended application!)

4.2.4 Tank Circuit and Tunability

Figure 4.6 shows the VCO circuit complete with the 2.0 nH tank inductor (L1), and the 540 Ω physical resistance (R1) that was added in parallel to lower the overall Q of the tank to about 5. In order to achieve the 5.2 GHz oscillation frequency, equation (4.1) dictates that a total capacitance of $C_{\text{eff}} = 468$ fF is required across the terminals of the tank. This capacitance was accomplished partially using accumulation mode varactors (VC1 and VC2 in Figure 4.6) and partially using fixed capacitors, represented as C1. In fact C1 was accomplished using two 360 fF MIM capacitors that were connected in series with the bottom metal plates connected together forming a more symmetrical structure than using one 180 fF capacitor as indicated in Figure 4.6. The remaining 277 fF was accomplished using varactors, where the noisier bottom plates of the capacitor (well of the accumulation mode devices) were directed towards the V_{CNTL} node to better isolate the VCO output from substrate noise.

One might wonder why C_{eff} was not completely implemented using tunable capacitors in order to maximize the tuning range (by maximizing K_{VCO}) and better guarantee that the VCO can reach 5.2 GHz with process variation, temperature, and layout parasitics. In fact having a very large K_{VCO} is problematic, especially in the case of the lock-and-roll receiver, because it makes the VCO's output frequency excessively sensitive to noise and fluctuations on V_{CNTL} . Recall from the discussion on the overall operation of the lock-and-roll receiver in section 3.4.1 that the loop initially phase locks the VCO and then the loop is opened and V_{CNTL} is held as constant as possible, through careful design (buffer filter, transmission gates with dummies, unique charge pump, etc.), while the VCO is injection locked by the incoming signal.

Despite design efforts to minimize the droop, the loop filter voltage, V_{CNTL} , will drop over time and eventually the VCO can no longer lock to the weak incoming signal. An excessively high K_{VCO} would result in the VCO being pulled away from 5.2 GHz more so for the same droop in V_{CNTL} and the loop would have to be closed again earlier, in order to regain frequency lock to 5.2 GHz. This scenario would result in fewer bits being transmitted for each close-open-close cycle of the loop.

4.2.5 Optimizing for Injection Locking Bandwidth and Low Power

Up until this point in the VCO design discussion, little has been noted about the inclusion of transistor M3 as shown in Figure 4.6, other than to highlight (recall section 4.2.1) that it reduces the maximum potential output swing of the current-limited VCO. Despite that drawback, M3 provides a useful element for controlling the current consumption, the gain margin, and to some extent the injection-locking bandwidth of the oscillator. Transistor M3 forms a current mirror with the diode connected reference transistor M1, and by adjusting the reference current through M1 (which in the case of the test chip is supplied from off chip), the bias current to the core of the VCO can be adjusted which increases or decreases $g_{m,\text{eff}}$. Therefore, if the real oscillator fails to startup due to insignificant gain margin in the presence of real life parasitics and process variations, the gain margin can be increased by increasing the bias current. Similarly, if the injection-locking bandwidth is too small, the bias current can be decreased which will lower $g_{m,\text{eff}}$, decrease the output swing, and increase the locking bandwidth according to (4.3).

Lastly, as the injection-locking bandwidth of the oscillator is of maximum importance for this design, it is unlikely that the VCO will be operated with an output swing that threatens to push M4 and M5 out of saturation. As Adler's equation (4.3) shows, the injection-locking bandwidth is inversely proportional to V_{osc} , and so the amplitude of the free-running oscillator will be kept as small as other factors allow. At nominal bias current of 1.0 mA, the peak output amplitude of the VCO simulates to be about 500 mV, thus Adler's equation predicts a one-sided locking bandwidth of $\omega_L \approx (2\pi(5.2e9))/(2(5)) (1.15e-3)/(500e-3) = 2\pi 1.2e6$ rad/s. Thus, the nominal locking bandwidth of the oscillator calculates to be roughly twice the required bandwidth

that was identified in section 3.4.4 as being necessary for achieving a communication range of 1.75 m. Note that if parallel resistance $R1$ had not been included to lower the overall tank Q from upwards of 11 to around 5, even the nominal locking bandwidth of the oscillator would likely have failed to meet the design specifications.

Transistor M2 in Figure 4.6 is a 10 pF MOS cap decoupling the supply reference node close in to the VCO core.

4.2.6 Designing a VCO with Margin – VCO Layout with Laser Cut Options

The layout (and schematic) of the VCO was carefully implemented to allow for the gain margin, oscillation frequency and tuning range, and the injection-locking bandwidth to be somewhat tunable using laser cuts as necessary. While the default layout was designed to work for the intended application without modifications, the overall functionality of the lock-and-roll receiver weighed heavily on the oscillator's functionality, frequency, and locking bandwidth, thereby justifying the precautionary steps that were taken. Luckily, none of the laser cut options were ever exercised given the good agreement between measurement and results that were simulated using careful parasitic extraction. Figure 4.7 shows the top level layout of the VCO in the lock-and-roll receiver.

The 2.0 nH kit inductor sits at the top of the VCO layout and occupies more than 50% of the layout area due to the necessary 50 μm gap between it and the substrate tiedown ring. The majority of the VCO's core circuitry is arranged about a single axis of symmetry, running left to right, and is physically located as close to the inductor as possible to minimize the routing. Transistors M1 to M7 are located on the right side of the layout, where the tunable varactors (and their dummies) are arranged linearly to the left of core transistors.

Both the fixed capacitors and de-Q'ing resistors are located well below the tunable varactors. The largest passivation opening exposes two top metal routes that connect the top terminals of the two de-Q'ing capacitors to nodes Out_p and Out_n (recall that the bottom plates of these devices were connected together to form a symmetrical structure across the tank circuit). Both top metal routes should be cut at the same time if one were to desire a lower C_{eff} , resulting in a higher frequency




Figure 4.7: Lock-and-Roll Receiver VCO Layout

oscillator with an increased tuning range. Only one of the two routes needs to be cut to remove both capacitors as they are connected in series, but cutting only one route would leave an asymmetrical parasitic connected to one of Out_p or Out_n . Directly below the center of the MIM capacitors is the bank of resistors used to de-Q the oscillator tank. There are five resistors in parallel (with additional dummies), where one resistor is connected at lower level metal but the remaining four can be cut from the circuit, two at a time, by using a laser to cut the top metal routes exposed at the bottom of the layout.

4.2.7 Simulated and Measured Results

Final extracted simulations of the VCO circuit were performed with the LNA (described in Chapter 5) connected to the tank given that the load it presents to the VCO has a slight effect on the oscillation frequency and tuning range. Both the VCO and the LNA layouts were carefully extracted to model all parasitics as much as possible, given the strong effects such parasitics also have on the final oscillation frequency and tuning range. The design kit that was used provides two options (coupled and decoupled) for extracting parasitic capacitances. The first option, coupled, calculates line to line capacitances and adds parasitic capacitors to the netlist between any nodes with overlapping metal in the layout. The second option, decoupled, lumps all parasitic capacitances to a reference node of the designers choosing (in this case the substrate node), such that each node in the circuit has one parasitic capacitance added to the netlist, connected between the node and the substrate. At the time that final simulations were being run prior to tapeout, the author had little first-hand knowledge of the correlation between circuit simulations in this kit using the two different options and measured results. However, the author's experiences using a different design kit with similar extraction options showed that the decoupled option tended to overestimate the parasitic capacitance, predicting a lower oscillation frequency and narrower tuning range than reality. The coupled option, while it predicted less parasitic capacitance, tended to agree better with the measured results and so final simulations were performed using the coupled option for parasitic capacitance extraction.

As the VCO's current consumption was made tunable by design, worth noting is that all measurements were conducted with the VCO bias adjusted to consume

1 mA of average current which aligns with the nominal simulation condition. The measured current consumption roughly followed the 10:1, $I_{DD}:I_{ref}$ relationship that was expected given the mirror ratio between M1 and M3, M4 (recall Figure 4.6, and the oscillator's gain margin proved ample to start up oscillations at the nominal 1 mA bias condition – agreeing well with simulation.

Oscillation Frequency and Tuning Range

Figure 4.8 shows the output spectrum of the VCO with connected LNA (plotted amplitude vs. frequency), as simulated using the extracted netlist over a range of control voltage settings at nominal temperature. From Figure 4.8 one can conclude that at the time of tapeout, the tank circuit was well centered such that a mid-rail control voltage setting (600 mV), resulted in an output frequency of 5.2 GHz.




Figure 4.8: Simulated VCO+LNA Tuning Range using Coupled Extraction

Figure 4.9 shows a comparison between the tuning range that was simulated with the extracted circuit and the one that was measured. In each case the slope of the curve represents the sensitivity of the oscillator, K_{VCO} . Calculated based on the useful

output range of the charge pump (see Chapter 6), the average simulated $K_{VCO, \text{sim}} = 1.1 \text{ GHz/V}$ and the average measured $K_{VCO, \text{meas}} = 0.7 \text{ GHz/V}$. Clearly the extracted simulation, using the coupled option for extracting parasitic capacitance, appears to underestimate the overall capacitance across Out_p and Out_n of the VCO. The effects of having additional parasitic capacitance would be a lower oscillation frequency (due to the larger C_{eff}), recall equation (4.1), and a smaller tuning range (and K_{VCO}) given the lower ratio of tunable capacitance to fixed capacitance within the overall makeup of C_{eff} . Both of these outcomes are clearly observed when comparing the simulated and measured results in Figure 4.9. Another likely explanation for the discrepancy is that while the effect of the LNA load on the VCO tank was included in the simulation, there are two buffers connected to the output of the VCO (see Chapter 6) which drive the divider input and the pad driver circuit which were not included in the simulation. The VCO output is connected to a MOSFET gate at the input of each of these buffers which will contribute additional capacitance to the total C_{eff} of the VCO. Nevertheless, despite the slight discrepancy between simulation and measurement, the results show that the VCO is able to be tuned to the desired frequency of 5.2 GHz with a control voltage that is achievable at the output of the charge pump. Additionally, the lower than expected K_{VCO} actually makes the VCO less susceptible to droop on V_{CNTL} during open-loop operation which improves the amount of time that the loop can be operated in that state, and increases the number of bits that can be transmitted for one closed-open-closed cycle of the loop. Note that the use of the buffers to drive the output pads makes characterizing the output signal swing from the VCO impossible.

Locking Bandwidth

Figure 4.10 shows the result of three separate simulations of the extracted VCO with connected LNA, overlayed one on top of the other. Both the on-chip input match between the LNA and the antenna, and the antenna's lumped element model itself, shown in Figure 3.5, are included in the extracted netlist. The first simulation was conducted with both the VCO and LNA powered up but with no input signal applied across the output terminals of the antenna, representing the free-running state of the oscillator. The output frequency for the control voltage that was used (roughly mid-rail at 600 mV) is 5.1965 GHz. The other two simulations that are shown were




Figure 4.9: Simulated vs. Measured Tuning Range

conducted under the same conditions with the same netlist, but where an input signal with a peak-to-peak signal swing of $115.8 \mu\text{V}$ was applied across the terminals of the antenna model, at frequencies of 5.196 GHz and 5.197 GHz respectively. Note that the injected tones are at $f_0 - 500 \text{ kHz}$ and $f_0 + 500 \text{ kHz}$. Both output spectra show that the VCO is injection locked to the incoming signal, which not only attests to the adequate locking range of the VCO circuit, but to the adequate input match and gain of the LNA circuit in accordance with the link budget analysis presented in section 3.4.4. While not shown, the same simulations were repeated, increasing the separation between f_0 and the injected tone in increments of 50 kHz each time (keeping the frequencies center of bin for the sake of the FFT calculation), to verify the extent of the circuit's locking range. Simulations showed that the oscillator could lock to input signals between $f_0 - 1250 \text{ kHz}$ and $f_0 + 1250 \text{ kHz}$ but not beyond. Therefore, the simulated locking bandwidth of the oscillator is $\omega_{L,sim} \approx 2\pi * 1250e3 \text{ rad/s}$ which is roughly equal to the $\omega_L = 2\pi * 1200e3 \text{ rad/s}$ that is estimated using Adler's equation. Given that the simulated gain of the LNA and the efficiency of the input matching circuit both factor into the simulated result, a slight discrepancy is understandable.

Note that this calculation (and simulation) reflects the 1.0 V peak amplitude that was achieved from the TX VCO which is roughly double what was assumed in the conservative link budget estimate in section 3.3. The simulated locking bandwidth claims a significant margin (factor of two) over the 500 kHz locking bandwidth that is required to guarantee a communication range of 1.75 m for the overall system.




Figure 4.10: Simulated Extracted VCO+LNA Locking Bandwidth Check

Figure 4.11 shows the output spectrum of the measured VCO (with LNA connected), injection locked to a modulated input signal that is being injected at the LNA's differential input port. Prior to enabling the input signal the control voltage was adjusted to $V_{\text{CNTL}} \approx 1$ V, at low resolution bandwidth and low span on the spectrum analyzer, to pretune the real VCO to a free-running frequency of 5.200 GHz (note Figure 4.9). With the modulated input signal enabled and switching between 5.1995 GHz and 5.2005 GHz at a rate of 1 kHz, the VCO circuit clearly injection locks. Unfortunately, little can be proven about the true locking range of the oscillator to input tones at a signal strength of 115.8 μ V peak-to-peak because of the impedance mismatch that is present at the input of the LNA circuit. While

the LNA input match was designed to interface with the integrated antenna, test chip real estate limitations demanded that the LNA inputs be brought to test pads where the input match could be characterized and test signals applied to verify the basic functionality of the lock-and-roll receiver. Thus, the exact voltage swing that exists on chip at the input of the LNA could not be determined as the complex input impedance was being supplied a signal through a bondwire, and traces and cables normalized to a $50\ \Omega$ system. Nevertheless, the result proves that the oscillator can injection lock to the desired input signal frequency.




Figure 4.11: Measured VCO Spectrum, Injection Locked to Modulated RX Input

4.3 VCO Design Summary

Understanding the fundamental concepts that apply to oscillator IC design is key to developing circuits that perform as desired, regardless of the application or the requirements. All oscillators must satisfy the Barkhausen criteria for oscillation and have adequate gain margin to guarantee startup given process variations and operating temperature fluctuations. Most oscillator applications benefit from a design with

low phase noise and an output spectrum that is free from spurs. As such, current-limited designs are popular as they keep the active devices operating in the saturation region and improve supply isolation while a differential topology improves immunity to common mode noise as well. The noise generated by the transistors themselves will contribute to the overall phase noise profile but sizing the devices appropriately can help to reduce thermal and flicker noise, while shot noise can all but be ignored for CMOS designs.

As with any IC design task, the tradeoffs must be considered carefully. While the width/length ratio and bias current of the design will regulate the gain margin of the circuit, the overall transistor size will influence phase noise, and the ratio of the sum of all fixed capacitances (including transistor gates and parasitics) relative to variable capacitance in the LC tank will dictate the tunability of the design (and K_{VCO}). The oscillation frequency is dictated by $\omega_0 = 1/\sqrt{LC}$ and the Q of the inductor will dominate and set the unloaded Q_U of the oscillator in most cases. High- Q VCO designs have lower phase noise profiles and are less susceptible to pulling and oscillator injection locking in the presence of a strong interferer. All oscillators can be injection locked at any offset frequency if enough injected power is applied, and Adler's equation has been shown to be relatively accurate for predicting the locking bandwidth of an electrical oscillator.

In unique scenarios, such as for the lock-and-roll receiver, the VCO is used in a nontraditional role serving as a band-pass filter (BPF) with gain. The requirements of the lock-and-roll receiver demanded a VCO design with a two-sided locking bandwidth in excess of 1 MHz and as such the design tradeoffs were optimized differently. Rather than choosing the kit inductor that simulated to have the highest Q , the inductor with the tightest tolerance over process and temperature was selected and the Q was intentionally lowered by adding parallel resistors to the tank such that Adler's equation predicted a sufficient locking bandwidth. The transistor sizes were optimized for gain margin and low current consumption, given the low tank Q and the self-powered application, at the expense of overall phase noise performance. Extracted simulations with the VCO and LNA were used to fine tune the oscillation frequency by adjusting the fixed capacitance in the tank, while an adequate ratio of fixed to variable capacitance was maintained to assure tunability. The layout was implemented with passivation openings and top level metal routing that facilitated easy trimming of

the fixed capacitors and the de-Q'ing resistors in the tank, if necessary, to achieve adequate gain margin and tuning range, while these fall-back options were never exercised. Simulations of the locking bandwidth confirmed the LNA/VCO pair's ability to amplify and injection lock to the predicted output of the integrated antenna, validating the 1.75 m communication range that was estimated with Friis' equation.

Measured results confirm that the oscillator's tuning range is adequate and sufficient to achieve the 5.2 GHz communication frequency at a control voltage that is compatible with the charge pump circuit. The measured K_{VCO} is lower than the simulated result which used coupled (rather than decoupled) parasitic extraction to estimate the parasitics, and neglected to account for the additional loading the buffer circuits in the test chip placed on the VCO output.

Due to limitations of the test chip design, the exact signal amplitude at the differential LNA input terminals cannot be determined (in favour of being able to characterize the LNA input match) and this limits characterization of the locking bandwidth of the VCO for a $115.8 \mu V$ input. Yet measurements show that the VCO can be injection locked to the modulated input frequency for which it was designed to track, and the circuit can therefore be used to validate the novel concept of the lock-and-roll receiver.

Chapter 5

Injection-Locking Circuit Design

The lock-and-roll receiver topology that is presented at a system level in Chapter 3 is unique in that it makes use of an integer-N PLL that is required to operate, at least some of the time, in a nontraditional way. In order to receive a stream of data bits, the PLL initially operates as a traditional closed-loop feedback system, frequency locking (and eventually phase locking) the VCO to 64 times the 81.25 MHz reference signal, yielding a 5.2 GHz VCO output signal. This step essentially pre-tunes the VCO's gain response to be centered on the 5.2 GHz modulated signal being transmitted by the lock-and-roll transmitter. Once the VCO is locked, the loop is then opened, and the VCO is injection locked to the incoming signal and the remaining loop components work to demodulate the bitstream. The VCO design is explained in detail in Chapter 4, where the approach that was taken to simultaneously optimize the locking bandwidth and the power consumption is discussed. The resulting VCO design proved to have ample tuning range (for the purpose of pre-tuning in closed-loop mode), adequate gain margin, and according to Adler's equation and simulation results, it can be easily injection locked to a 5.1995 to 5.2005 GHz input signal with an amplitude of 115 μV applied to the antenna. The role of the LNA circuit is to injection lock the oscillator, serving as an interface between the integrated loop antenna and the oscillator. As such, the LNA must be adequately impedance matched to the inductor, and provide enough gain to the input signal that the VCO can be injection locked. Additionally, the LNA's output must be connected to the tank of the VCO, such that the VCO can be injection locked by the LNA's output signal, but without disrupting the tuning range of the VCO, its center frequency, its gain margin or its locking bandwidth.

This chapter compares alternatives for coupling signal into the core of an LC VCO in order to injection lock it without substantially changing the loading of the VCO's tank circuit itself. Both voltage coupling and current-steering approaches are analyzed. As a practical example, the design of the LNA circuit in the lock-and-roll receiver is explained along with its simulated and measured results. A unique requirement of the LNA circuit in the lock-and-roll RX is that it be conjugately matched to the on-chip antenna that has a low impedance of $7.1 + j66.0 \, \Omega$, and so the design of the low- Q matching circuit is also presented in this chapter.

5.1 Coupling Voltage vs. Steering Current

There are essentially two approaches that can be taken to electrically inject the desired signal into the core of the VCO for the purpose of injection locking. The injection circuit can be designed as a separate amplifier with its own supply and load yielding an output voltage that can be coupled into the core of the VCO, or a current-steering approach can be adopted. Figure 5.1 and Figure 5.2 show the two approaches, respectively, implemented in simplified forms.




Figure 5.1: Locking Circuit Schematic with Coupled Output Voltage

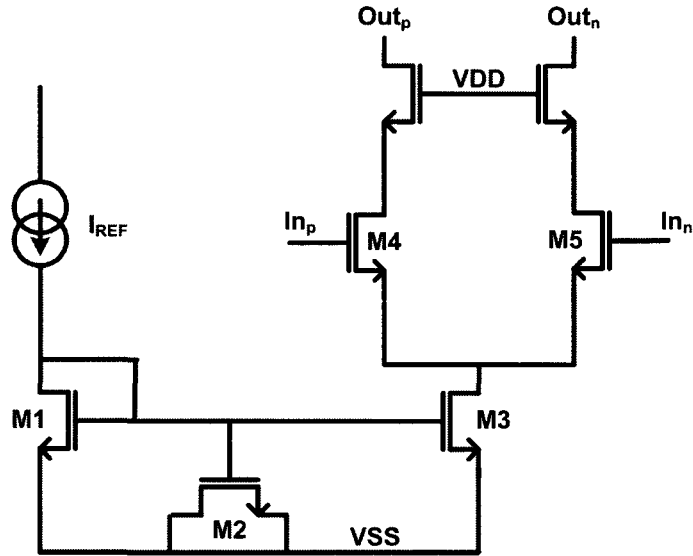


Figure 5.2: Locking Circuit Schematic using Current-Steering Approach

Prior to highlighting the differences between the two approaches and their advantages, some discussion on their similarities and the interchangeability of some of the aspects of the designs shown in Figure 5.1 and Figure 5.2 is warranted.

5.1.1 Differential vs. Pseudo-Differential, Cascode Topologies and Tail Currents

The amplifier circuit shown in Figure 5.1 is a pseudo-differential common-source amplifier with resistive load. The amplifier is said to be only “pseudo” differential because even though the gates are biased at the same voltage and the input signal is applied there differentially, the sources of the NMOS transistors are connected to ground with no shared tail current. The schematic in Figure 5.2, in comparison, is a true differential circuit with a shared tail current, where increasing the gate voltage on one side with respect to the other will shift the balance and ultimately reduce the drain current on the opposing side – a scenario not true of the pseudo-differential circuit. Differential circuits are relatively immune to both AC and DC common mode input noise while pseudo-differential circuits provide only AC immunity in that regard, when considering a differential output. The benefit of the pseudo-differential design is that in low-power designs where the supply voltage is low, there is more

headroom available to the gain devices as there is no tail current (mirror device) that requires $V_{DS_{\text{sat}}}$ of headroom to operate. Both the current-steering and coupled-output-voltage alternatives to designing an injection locking circuit could be pseudo-differential or truly differential circuits, biased with either a common tail current or gate voltage reference.

The circuit shown in Figure 5.2 is a differential cascode gain stage with NMOS devices, typically operated in saturation, connected above the gain devices. This approach improves the output impedance seen looking back from the VCO which is clearly advantageous if one is trying to maintain a certain gain margin and tank Q in the VCO. The disadvantage of the cascode approach, as with the tail current, is reduced operating headroom. Again, both the current-steering and coupled-output-voltage approaches could be implemented with a cascode gain stage.

In both cases, the designs could be optimized rather equally for noise figure, and input match, yet there are inherently some advantages of using one alternative over the other given certain system level requirements.

5.1.2 Optimizing for Efficiency with Hard Switched Inputs

The major benefit of the current-steering approach is that it is generally more power efficient than the coupling voltage approach because there is no separate load to the supply from which to draw current. All AC current that is drawn through the differential pair is sloshed through the tank of the VCO and directly contributes to the injection locking effort. This approach works best if the differential pair is switched very hard, ideally with a rail-to-rail input voltage which maximizes the ratio of AC to DC current in the stage. A possible drawback of this approach, however, is that the DC current drawn by the stage is also pulled from the VCO circuit. Consider the VCO schematic for the lock-and-roll receiver shown in Figure 4.6. The DC current drawn by the locking circuit in Figure 5.2, if used with the lock-and-roll VCO, will be drawn through the VCO's PMOS devices M6 and M7 from the supply rail. This in turn will affect the balance of $g_{m,p}$ and $g_{m,n}$ when the locking circuit is enabled compared to when it is not. Additionally, the DC drop across M6 and M7 in the VCO will reduce headroom to the locking circuit making it very difficult to implement a cascode design, and even more difficult to use a cascode design and a tail current

concurrently in low-power applications with a low supply voltage. Lastly, without any additional gain stage implemented in front of the locking circuit and interfacing with the antenna, the differential pair certainly won't be hard switched because the signal amplitude coming off the antenna (and the input match) will rarely be large enough to do so, resulting in substantially more DC than AC current in the gain stage (i.e. class A operation) and exacerbating the disruption of the VCO.

5.1.3 Minimizing Disruption on the VCO Core

The coupled voltage approach shown in Figure 5.1, while clearly drawing current from a separate load to the supply, provides a cleaner divide between the VCO and LNA circuits. While not all of the AC current will be transferred to the VCO because of the additional load, the ratio of the LNA's load relative to the load imposed by the VCO connected at the output can be optimized to maximize the efficiency of the circuit. In fact, as the AC gain of the classical common-source amplifier circuit is well known to be approximately equal to $A_v = g_m Z_{L,tot}$ [57], where $Z_{L,tot}$ is the parallel combination of the LNA's load (including the complex output impedance of the gain devices themselves at ω_{inj}) and the VCO's overall parallel tank impedance ($R_{p,eff}$) at resonance, maximizing the LNA's own load impedance improves both the LNA's gain and the efficiency of the locking effort. Note that the injected tone at ω_{inj} was assumed here to be very close to the center frequency of the oscillator (ω_0), allowing the simplification that the load imposed by the VCO circuit is $R_{p,eff}$. In fact ideally the resistive loads shown in Figure 5.1 would be replaced altogether with high- Q RF choke inductors such that the overall loading of the LNA would be dominated by $R_{p,eff}$ from the VCO, with only the small amount of AC current that is recirculated through the complex output impedance of the LNA's gain devices themselves not contributing to the injection locking goal.

An additional benefit of the coupled voltage approach is the ease with which a tuned filter could be implemented at the independent load to provide improved immunity to interfering input signals and noise. In applications such as the lock-and-roll receiver where the modulated input signal is very narrow-band, implementing an LC filter using only on-chip components that would have enough selectivity to

be useful would be a challenge, yet the opportunity exists with the coupled voltage topology.

5.1.4 Designing for Low Voltage Supply and Weak Input Swing

As section 5.1.2 explains, when the locking circuit can be driven hard with a near rail-to-rail input signal (approaching class-B operation), the current-steering approach is often preferred, as was the case when implemented by DeVries and his colleagues in their sub-sampled RF receiver [58], [59], [60]. However, when the locking circuit is to interface directly with the antenna and handle input signals of very small amplitude, the coupled voltage approach's appeal improves. Coupling voltage allows for the LNA and VCO circuit designs to be optimized individually and for this reason the approach was implemented for the locking circuit in the lock-and-roll receiver. At the expense of slightly increased current consumption, the transistor sizing and bias current can be optimized for gain, input match, and noise considerations more easily without worry of disrupting the behaviour of the VCO circuit. Additionally, as the VCO circuit is shown in Chapter 4 to have been heavily de- Q 'ed in favour of increasing the locking bandwidth, the load presented by the VCO to the LNA circuit very much dominates the overall loading of the LNA such that little AC current is lost to the separate load which is introduced by choosing this topology.

5.2 The Lock-and-Roll Receiver LNA

5.2.1 Circuit Topology and Design

The schematic for the lock-and-roll receiver's locking circuit is shown in Figure 5.3. Comparing the circuit with the simplified voltage coupling topology shown in Figure 5.1 one can observe that the circuit uses active PMOS load devices (M5 and M6) in place of a resistive load. The use of the active loads allows for a high-impedance load at 5.2 GHz (maximizing the signal transfer to the VCO for locking), while providing much less DC drop across the load that is connected to the supply than what would result from simple resistor loads, maintaining headroom for the gain devices M3 and

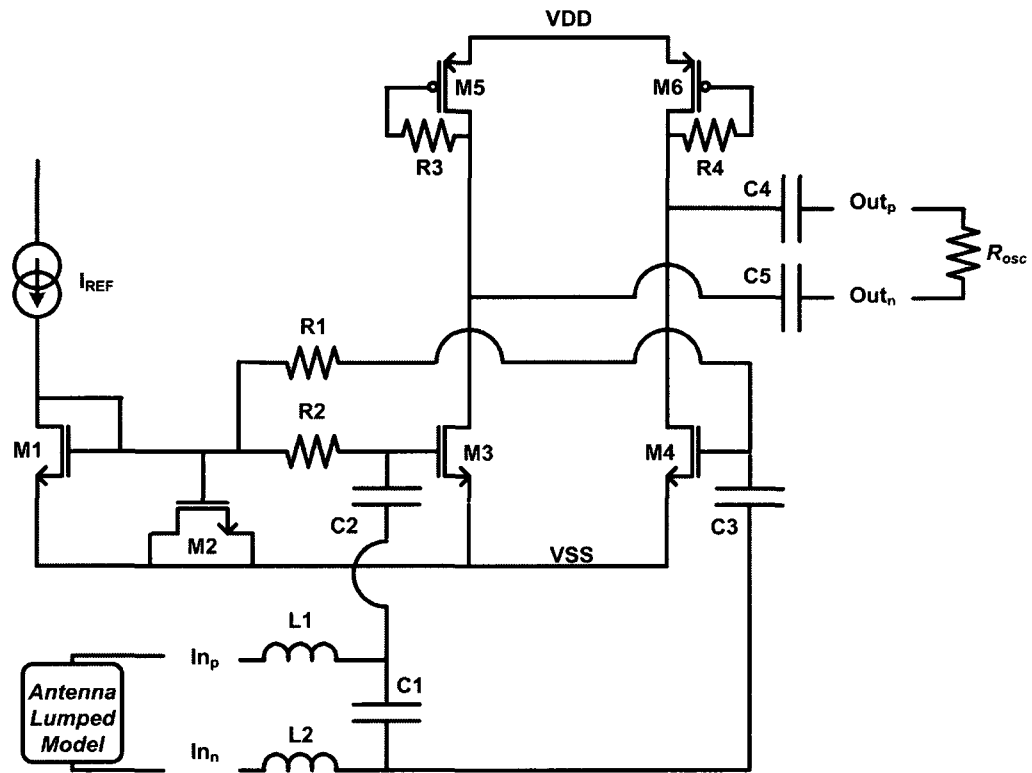


Figure 5.3: Locking Circuit Schematic for the Lock-and-Roll Receiver

M4. Ideally, large inductive chokes would be used rather than M5 and M6 to maximize headroom and resulting in the VCO being the only AC load on the circuit (other than the output impedance of M3 and M4), but implementing chokes on chip is impractical and M5 and M6 proved to provide a suitable compromise between using resistors and inductors. As the parasitic gate-to-source capacitance of M5 and M6 is a relatively small impedance at 5.2 GHz, resistors R3 and R4 keep the parasitic from essentially shorting out the output (and VCO tank) to VDD. As the DC gate current is nil, R3 and R4 do not affect the DC response of the diode connected devices. Similar to the VCO topology presented in Chapter 4, the locking circuit uses a gate bias voltage that is generated with a reference current and transistor M1, while M2 serves as a decoupling capacitor. Resistors R1 and R2 are large and isolate the AC input signal from the reference current, and like R3 and R4 they do not affect the DC bias condition as there is no DC MOSFET gate current. Capacitors C2 and C3 are large AC coupling capacitors that do not factor into the input match but do isolate the DC bias of the amplifier from the AC input at the antenna. Capacitors C4 and C5 are carefully sized to balance the amount of signal that is coupled into the tank of the VCO with the impedance that is presented to the VCO's tank circuit. Both the output impedance of the LNA and its voltage gain are quoted after these capacitors that also isolate the two circuits from a DC perspective. Inductors L1 and L2, along with capacitor C1, form the matching circuit that interfaces with the antenna. The resistor R_{osc} is included for simulation purposes only and was sized at $415\ \Omega$ (recall $R_{p,eff}$ from Chapter 4) to represent the loading affect of the VCO on the LNA circuit. During final verification simulations, the actual VCO was used in place of R_{osc} to verify the LNA gain and the injection locking bandwidth as summarized in section 4.2.7. To design the input match and to simulate the circuit gain accurately, the equivalent model of the antenna (recall Figure 3.5 and Table 3.1) was also included in most simulations, connected as shown in Figure 5.3.

5.2.2 Trading-Off Noise, Current, Transistor Size and Output Impedance

So far in this chapter the terms LNA and locking circuit have been used rather interchangeably. In fact the locking circuit in the lock-and-roll receiver is somewhat

different than a traditional LNA, primarily with regards to the way it was optimized and the traditional design tradeoffs balanced. In a typical heterodyne type receiver, the LNA circuit is the first gain element in the receiver and often interfaces directly to the antenna, or to a filter between the antenna and the LNA. The locking circuit in the lock-and-roll receiver is similar in this regard. The major concerns of the typical LNA designer are to implement a circuit that has very low noise figure, because as Friis' [61] equation for the distributed noise figure in radio receivers (not to be confused with Friis' equation for link budget from section 3.3) suggests, the noise figure of the first gain element in the receiver chain dominates the overall noise figure of the receiver. Friis' equation can be summarized as

$$F = F_1 + \frac{F_2 - 1}{G_1} + \frac{F_3 - 1}{G_1 G_2} + \dots \quad (5.1)$$

where F is the overall noise factor of the system (i.e., the receiver), and F_n and G_n are the noise factor and linear gain, respectively, of the n 'th stage in the chain. The noise figure (NF) is simply the noise factor (F) expressed in decibels (i.e., $NF = 10\log(F)$).

Friis' equation also highlights that the noise factor of each subsequent block in the chain is divided down by the product of the gains of the blocks that precede it, and so high gain is also desirable for the LNA circuit – lessening the effect of the noise of all subsequent blocks on the overall noise figure of the chain.

As the LNA could be faced with a very large input signal if the physical separation of the RX and TX is small (resulting in low path losses, recall equation (3.3)) the linearity of the amplifier is often carefully scrutinized. Metrics that are often used when judging an overall LNA design are therefore NF, third order intercept point (IP3), 1-dB compression point (P1dB) and of course current consumption. If the amplifier incorporates a tuned load, the 3-dB bandwidth might also be carefully scrutinized. Given these considerations, a very common recipe [62] for LNA design follows:

1. Determine the current density for the gain transistors, given the design process being used, for which the lowest NF can be achieved and bias the devices at this density regardless of their size.

2. Size the gain devices such that the real part of the driving impedance for which the lowest NF is achieved is equivalent to the real part of the actual complex driving impedance - typically this is $50\ \Omega$.
3. Add inductive source degeneration, which is well known [46] to increase the real component of the input impedance seen at the gate, until the real part of the input impedance matches that of the driving impedance.
4. Add a reactive component, typically an inductor is required, in series with the gate to conjugately match the input of the amplifier to the driving load.

The result of this 4 step process is a design that is simultaneously matched for noise performance and power transfer from the antenna or pre-filter. In fact this formula was very loosely followed to optimize the design of the lock-and-roll receiver's amplifier. The difficulty in following the recipe exactly for the lock-and-roll receiver's LNA arises when one considers the $7.1 + j66.0\ \Omega$ impedance of the antenna (recall section 3.2.1), which roughly resembles $630\ \Omega$ in parallel with a 2 nH inductor at 5.2 GHz. The size of M3 and M4 would have been prohibitively large if step 2 was followed exactly, resulting in large parasitic capacitances that would have weighed heavily on the balance of the VCO tank. Additionally, the lock-and-roll test chip was designed with a budget of roughly 1 nH of downbond inductance (estimated from 3 downbonds in parallel, 3.0 mm long, assuming 1 nH/mm) and so the addition of extra degenerative inductance (according to step 3) was to be avoided if possible given the limited die area allocated to the project.

5.2.3 Simulated LNA Performance

Through simulation, a delicate balance was achieved whereby the transistors were sized to achieve adequate NF and gain, for a relatively low bias current, with an output impedance that did not greatly impact on the VCO's performance when capacitively coupled to the tank circuit, and most importantly, facilitated a reliable, low- Q on-chip match to the antenna. Table 5.1 summarizes the final simulated performance of the lock-and-roll receiver's amplifier versus the common metrics used to evaluate LNA circuits. Naturally, all simulations are summarized after extracting layout parasitics, and with the lumped antenna model and VCO loading in place.

Table 5.1: Lock-and-Roll Receiver LNA Extracted Performance

Metric	Nominal Simulated Performance at 5.2 GHz
NF_{\min}	1.3 dB
NF	4.7 dB
A_v	21 dB
S11	-13.4 dB
P1dB (output referred)	0 dBm
IP3 (output referred)	9 dBm
Current Consumption	1.2 mA
Output Impedance	$45.3 + j202.6 \Omega$

5.2.4 LNA Output Impedance and the Effect on the VCO

The amplifier's output impedance of $45.3 + j202.6 \Omega$ is actually one of the more important aspects of the design for the lock-and-roll receiver. At 5.2 GHz, this impedance can be represented equivalently by a 952Ω resistor in parallel with a 143 fF capacitor. The effect that this impedance has on the performance of the VCO was easily negated by decreasing the size of the 360 fF fixed capacitors in the tank (recall section 4.2.4) to accommodate the extra capacitance, while the equivalent 952Ω simply lowered the equivalent $R_{p,\text{eff}}$ of the tank from 415Ω to 290Ω . With the fixed de- Q 'ing resistors having been laid out to facilitate laser surgery and with the adjustable tail current in the VCO allowing for some control of the gain margin in case the oscillator fails to startup at the default bias condition, the fixed resistors in the core were left as per the original design and layout.

5.2.5 Low- Q On-Chip Input Match Design

One of the most critical aspects of the lock-and-roll receiver's LNA design is the on-chip input match that interfaces with the large integrated loop Antenna. Without a conjugate match to the antenna, equation (3.3) and the 1.75 m communication range it predicted in section 3.3 are not valid. Recall that the analysis assumed a conjugate match to the antenna and 20 dB of LNA gain in order to assure an injected signal amplitude at the VCO that yielded an adequate locking bandwidth as predicted using Adler's equation (2.2). An important element of the match is that it be relatively

low- Q given the fact that it is to be implemented using on-chip components which tend to have much looser tolerances than their off-chip counterparts. A high- Q match, implemented on-chip, would be very sensitive to the typical variation of on-chip inductors and capacitors – resulting in an S11 that is not reliable from part to part, or equivalently, low die yield.

In general, the Q of a matching network can be estimated [63] as

$$Q_{match} = \sqrt{\frac{R_{p1}}{R_{p2}}} - 1 \quad (5.2)$$

where R_{p1} and R_{p2} are the largest and smallest equivalent parallel resistances presented to either side of the match, respectively. Given the process variation of on-chip passive elements today, a matching network Q of 1 to 3 is typically safe. Beyond a Q of 3 the yield will be low if an S11 of better than -10 dB is desired. For example, one can generally expect to achieve an S11 of better than -10 dB at RF frequencies, reliably over process and temperature fluctuations, when matching 50 Ω to 250 Ω , which requires a matching network $Q_{match} = 2$, using today's modern CMOS processes. On the other hand, one will be hard pressed to match 50 Ω to 850 Ω , requiring $Q_{match} = 4$, using on-chip components if an S11 of better than -10 dB is to be achieved with high yield over process and temperature. The higher the communication frequency the more difficult it becomes to guarantee S11 for all conditions.

Given the multiple downbonds of the lock-and-roll receiver test chip and the predicted inductive degeneration of roughly 1 nH (recall section 5.2.2) that they contribute to the LNA circuit, transistors M3 and M4 in Figure 5.3 were sized such that a matching network Q of lower than 3 could be achieved without the use of additional degeneration inductors on chip (to increase the real part of the impedance looking into M3 and M4), which would consume valuable die area. Recall from section 3.2.1 that the impedance of the antenna at 5.2 GHz is $Z_{ant} = 7.1 + j66.0 \Omega$. As such, the target differential impedance looking into the terminals of the LNA circuit (including the match) in order to achieve a conjugate match is $Z_{in} = 7.1 - j66.0 \Omega$. The simulated differential input impedance looking directly into coupling capacitors C1 and C2 (see Figure 5.3) of the final LNA design was $Z_1 = 194.2 - j567.6 \Omega$ at 5.2 GHz, which resembles roughly a 1.85 k Ω resistor in parallel with a 50.0 fF capacitor. Thus, the Q of the required matching circuit can be estimated from (5.2)

as $Q_{match} \approx 1.4$. This result should facilitate a reliable on-chip match over process and temperature fluctuations. After matching, the input impedance of the LNA simulated, under nominal conditions, to be $Z_{in} = 8.5 - j52.4 \Omega$. Figure 5.4 shows the key impedances plotted on a Smith chart along with the impedance translations caused by each of the matching components.




Figure 5.4: LNA to Antenna Match Translation on a Smith Chart

5.2.6 Match Variation over Process and Temperature

While the use of series inductors in the input match has the drawback of consuming substantial die area, the benefit is that inductors have a relatively tight tolerance over process and temperature as their inductance is mostly dictated by their shape (thus

the consistency of the lithography is a dominant factor). In fact the CMOS process that was used for the lock-and-roll receiver test chip had inductors that varied in inductance by less than 2% over process and temperature (3σ limit). The inductor Q , however, varied by about 10% because it is largely dominated by the resistances of the metals that make up the coil which are more sensitive to doping levels. Variations in substrate resistivity and oxide capacitance also affect inductor Q . That said, the Q variation among wafers from the same lot, where doping and oxide thicknesses should be similar, is likely much smaller. Even less variation would be expected among dice from the same wafer.

The variation of integrated capacitors is much worse than that of inductors. MIM capacitors such as the one used in the input match of the lock-and-roll receiver are sensitive to variations in the oxide thickness that separates the two (or three) plates. The models of the CMOS process used for the lock-and-roll receiver test chip suggest that the MIM capacitors have a 3σ variation of about 15% over process and temperature (where temperature only contributes about 2%).

An excellent reference on the factors that influence the accuracy and matching of on-chip components, both active and passive, is Alan Hasting's *The Art of Analog Layout* [64].

The input match of the lock-and-roll receiver was simulated over worst case corners and temperature. The fluctuation in S11 is shown in Figure 5.5, where the variation in the input match is nearly completely dominated by the worst case 15% variation in MIM capacitance. The worst case S11 of the match simulated to be slightly worse than -8 dB and the target of -10 dB over all corners could not quite be guaranteed despite the low- Q match. While the result falls short of achieving the S11 goal under all conditions, in a production scenario, detailed measurements of the overall affect on the receiver's performance with parts manufactured using process skews or selected from a binning process would be necessary to truly identify a performance problem that would affect the yield. Excess margin on the locking bandwidth of the VCO and the LNA gain could potentially accommodate any shortcomings of the match. Additionally, the match could potentially be redesigned to be less sensitive to process and temperature with the tradeoff of having higher nominal S11.




Figure 5.5: S11 Simulation with Extracted LNA over Process and Temperature

5.2.7 LNA Circuit Layout Including Input Match

The LNA circuit was carefully laid out, including the input match, to be as symmetrical as possible given the pseudo-differential topology. Figure 5.6 shows the layout which is nearly perfectly symmetrical about an axis running vertically through the center of the LNA core and between the two inductors in the input match. The core of the LNA was laid out using common centroid [64] layout techniques and dummy cells for M3, M4, M5, and M6, R3 and R4 (see Figure 5.3) to best preserve symmetry over process gradients. Unlike the VCO layout (recall Figure 4.7), there are no laser tunable options in the LNA layout. Where a low gain margin or locking bandwidth in the VCO would have been disastrous (if uncorrectable) when it comes to proving the functionality of the novel lock-and-roll receiver using the test chip (a primary goal of the thesis), an error in the tuning of the input match was deemed to be an acceptable risk (not warranting laser trimmable elements). The overall functionality of the receiver can still be verified with a mistuned input match by applying more input power given that the LNA inputs are connected to pads for measurement and




Figure 5.6: LNA Layout Including Input Match

characterization. Probe de-embedding structures were added to the test chip, as is explained in Chapter 7, to allow for the input pads that are connected to the LNA match to be calibrated out of the measurements. Due to the fact that the LNA's output is so closely tied to the VCO tank circuit for injection locking, the output signal of the LNA could not be measured directly and the number of measurements that could be done to verify the performance of the LNA itself were somewhat limited.

5.2.8 Measured Results

As mentioned, the output terminals of the LNA could not be physically accessed for measurement purposes which somewhat limited the type of verification that could be performed against simulated results. While the gain and linearity of the amplifier could not be verified against the results summarized in Table 5.1, the current consumption and the input match were measured. Additionally, recall that the measurements of the VCO's locking performance, oscillation frequency, current consumption and tuning range (which are presented in section 4.2.7) were made by applying an input signal to the input terminals of the LNA circuit's input match. These measurements therefore confirm the negligible effect that the LNA's output impedance has on the VCO's performance. Unfortunately, as the input match is not designed for a $50\ \Omega$ system, the true input signal that is delivered to the terminals of the match is unknown during testing and careful characterization of the LNA's gain (and consequently the VCO's locking bandwidth for the intended worst case received signal strength) is impossible. The tradeoff, however, is that the input match can be measured to validate the conjugate match to the antenna.

Like the VCO, the LNA's current consumption was made tunable by design. All measurements were conducted with the LNA bias adjusted to consume 1.2 mA of average current which aligns with the nominal simulation condition. The measured current consumption roughly followed the $10:1, I_{DD}:I_{ref}$ relationship that was expected given the mirror ratio between M1 and M3, M4 (recall Figure 5.3).

To truly measure a differential input impedance a 4-port network analyzer (NA) is necessary. A 4-port NA is extremely expensive, and no such hardware was available at Carleton University, so a 2-port network analyzer measurement was made and a free software tool, *AppCAD* [65] available from Agilent Technologies, was used

to mathematically translate the results to yield an estimated differential S11 measurement. The tool proved quite useful in this regard. The input impedance of the LNA was characterized at the nominal bias condition ($I_{DD} = 1.2$ mA), using RF probes, after calibrating the VNA and de-embedding the bondpads using the on-chip de-embedding structures that were included on the test chip (discussed in Chapter 7). Figure 5.7 shows the differential input impedance measurement, $Z_{in,meas}$ at 5.2 GHz, that results from using AppCAD to analyze the 2-port measurement. The result is plotted along with the simulated differential input impedance ($Z_{in,sim}$) on a Smith chart, normalized to $50\ \Omega$, that shows both the lines of constant resistance and constant conductance on the same plot for the sake of discussion. A possible explanation for the deviation between simulation and measurement is also depicted visually on Figure 5.7.

The measured result is actually quite close to the simulated result which is encouraging given the added translation step that was involved. Recognizing that the parallel capacitance in the match translates the impedance along lines of constant conductance on the Smith chart, while the series inductors translate the impedance along lines of constant resistance, one can hypothesize about what may be contributing to the difference in simulation versus measurement. If the shunt capacitor in the match was actually 20% smaller than the intended value, and the series inductors were actually 20% smaller than intended, the translation that would be expected is that shown in orange on Figure 5.7. Recognizing, however, that the 3σ variations on capacitors and inductors in this process are 15% and 2% respectively (recall section 5.2.6), such a scenario is surely unlikely. More likely is that the impedance difference results from the summation of many factors and is a distributed effect. The modeled versus actual Q of the matching components, the input impedance of the transistor devices themselves, the accuracy of the de-embedding structures and the calibration procedure for the VNA will all factor into the impedance that is measured, along with the accuracy of the mathematical translation of the 2-port data using AppCAD. That being noted, the measured and simulated impedances show relatively good correlation and suggest that with some fine tuning the input match could likely be adjusted to provide an excellent conjugate match to the on-chip antenna.




Figure 5.7: Measured vs. Simulated Differential Input Impedance

5.3 Injection-Locking Circuits Summary

The injection locking bandwidth of an oscillator is directly proportional to the strength of the injected tone relative to the free running tone of the oscillator as indicated by Adler's equation. As such, the role of the locking circuit, to amplify the injected tone to the necessary signal strength and to inject it into the oscillator core without disrupting the delicate impedance balance of the tank circuit, is an important one.

The locking circuit can be designed to steer current through the core of the VCO circuit or to couple voltage across the tank. The current-steering approach works well when the input signal is strong enough to hard switch the transistors in the locking circuit, thereby driving the circuit towards class-B operation which maximizes the AC portion of the current being drawn from the VCO. The drawback with the current-steering approach is that, while there is no separate load for the circuit and all the AC current in the circuit contributes to injection locking, the DC current must also be drawn from the VCO which can disrupt the balance between $g_{m,p}$ and $g_{m,n}$ if a complementary CMOS topology is used for the VCO. The coupled voltage approach uses a separate load to the supply, thereby decoupling the DC bias of the amplifier from that of the VCO, with the drawback of the AC signal at the output being split between the VCO tank and the separate load at the amplifier.

Regardless of the locking circuit topology that is chosen, both differential and pseudo-differential circuits can be used, with or without cascode devices. True differential circuits have a single tail current and reject both AC and DC noise at the input, whereas pseudo-differential circuits reject only AC noise but allow for more headroom to the gain devices – an advantage when designers face the low supply voltages that inherently accompany low-power, CMOS technologies. Cascode topologies help to increase the output impedance of the amplifier, thereby reducing the disruption on the VCO tank circuit that is connected to the output, yet much like the tail current of a true differential circuit they often starve the gain devices of precious headroom in low supply applications.

The coupled voltage approach was chosen for the circuit in the lock-and-roll receiver primarily because of the clean divide it provided between the LNA and VCO designs from a DC perspective, given the expected class-A operation with small input signal swing, and also because it maximized the headroom available for the gain

devices given the low-voltage supply. The design is a pseudo-differential topology without cascode devices, further maximizing headroom to the gain devices. The LNA in the lock-and-roll receiver is unique in that it was optimized not merely for NF, but to eliminate the need for further inductive degeneration (beyond the down bonds) to enable a low- Q integrated match to the integrated antenna (with complex impedance), thereby keeping the layout area that is required to a minimum.

While the exact gain, linearity and NF of the amplifier cannot be measured because the output terminals do not connect to pads that interface with the outside world (as doing so would be to the detriment of the VCO's tank circuit impedance), the input match is connected directly to differential input pads that can be probed to characterize S_{11} . The measured impedance seen at the LNA's input match generally agrees with simulation and suggests that with some fine tuning a completely integrated match to the integrated antenna is feasible.

Chapter 6

PLL Component Designs that Enable Open-Loop Operation

The lock-and-roll receiver that is outlined in Chapter 3 is unique in that it is completely integrated and is designed to function in both the closed-loop and open-loop modes of operation. Chapters 4 and 5 outline the design and implementation of the core RF circuits, namely the VCO and LNA, where the traditional tradeoffs are balanced, at times unconventionally, to meet the unique demands (mainly the injection locking bandwidth) the system places on those blocks. This chapter addresses the design and implementation of the components in the lock-and-roll receiver loop that specifically enable the open-loop mode of operation. As this mode is unique to the lock-and-roll topology, the concept of designing a loop that is optimized to operate in this mode does not fall under the umbrella of topics that are traditionally discussed in the context of designing PLL and frequency synthesizer circuits.

Recall from Chapter 3 that the loop is initially closed and phase locked to the reference signal, where the control voltage settles out to the DC level that pre-tunes the VCO output frequency to be centered on the carrier frequency used by the transmitter with which the receiver is communicating. With the pre-tuning step accomplished, the loop is then opened and the oscillator is injection locked to the incoming modulated signal while the control voltage is held near constant. With the VCO having been pre-tuned to the center of the modulated signal of the transmitter, the modulated signal falls within the injection-locking bandwidth of the receiver's VCO (which can be predicted using Adler's equation, recall section 4.1.5) and the VCO therefore locks to the signal readily and the remaining loop components are used to demodulate its output. If the control voltage does not stay constant however,

the VCO's center frequency will drift (taking its locking bandwidth along with it) and the receiver will fail when the incoming signal no longer falls inside of the VCO's injection-locking band. Note here that the term "locking band" is used to refer to the absolute band of frequencies to which the VCO can be injection locked by a given input power at any particular instance in time, whereby the locking bandwidth refers to the width of the locking band that is centered at the oscillator's free-running frequency, which is adjusted via V_{CNTL} . The injection-locking band drifting beyond the modulated input signal's bandwidth is one of two mechanisms for failure which are introduced as a result of control voltage droop. Given the Δf of the modulated input signal, if the center frequency of the RX VCO drifts by more than Δf from its nominal condition, the PFD will no longer see the divider's output as switching between output frequencies that are faster or slower than the reference, and demodulation of the data will fail as a result. The first failure mechanism results from a loss of synchronization between the transmitter and the receiver, while the second mechanism is a result of the receiver's architecture and its sensitivity to Δf . In reality the second mechanism is likely to dominate most of the time as similarities between the TX and RX loop designs will tend to cause them to drift in the same direction and presumably at similar rates, assuming similar operating temperatures, etc.

There are two primary concerns that threaten to disturb the control voltage in the open-loop mode of operation, where both are exacerbated by implementing the topology on a single IC. The first concern is charge injection at the time the loop is opened, where charge redistribution at the moment the loop switches states can accumulate on the loop filter capacitors and change the DC control voltage that is held (and thus the center frequency and locking band of the oscillator). Charge injection is minimized for the lock-and-roll receiver loop by optimizing the design and layout of the charge pump circuit to this end, and by adopting much the same transmission gate topology, for the loop filter switch, that was pioneered by those designing switched capacitor circuits that faced a similar dilemma. The second concern is control voltage decay over time, which is somewhat unavoidable. As the control voltage is determined by the charge held on the loop filter capacitors, any leakage paths that enable charge to escape from the loop filter capacitors contribute to control voltage droop with time, and eventually to the receiver failing, or rather requiring the loop to be closed again and the control voltage refreshed. As leakage from the loop filter can never be

eliminated altogether, the design strategy focuses on minimizing the leakage as much as possible such that the time taken for the VCO's locking band to drift beyond the incoming modulated signal, or for the VCO's center frequency to drift by more than Δf , is sufficient to demodulate a reasonable quantity of data (recall section 3.1.1 and section 3.4.1). The control voltage droop is minimized primarily through the use of a transmission gate to further isolate the finite output impedance of the disabled charge pump from the loop filter, and by the use of a unity-gain buffer circuit that isolates the loop filter from the leakage imposed by the VCO's varactor diodes. Figure 6.1 shows the lock-and-roll receiver block level diagram and highlights the main loop components that are optimized to enable the open-loop operating mode, and which are analyzed and discussed in this chapter. While the PFD circuit does not really facilitate the open-loop operating mode, understanding the design of the PFD is fundamental to understanding the charge pump design and so the PFD design is presented here. The divider circuit, however, is analyzed in Chapter 7.

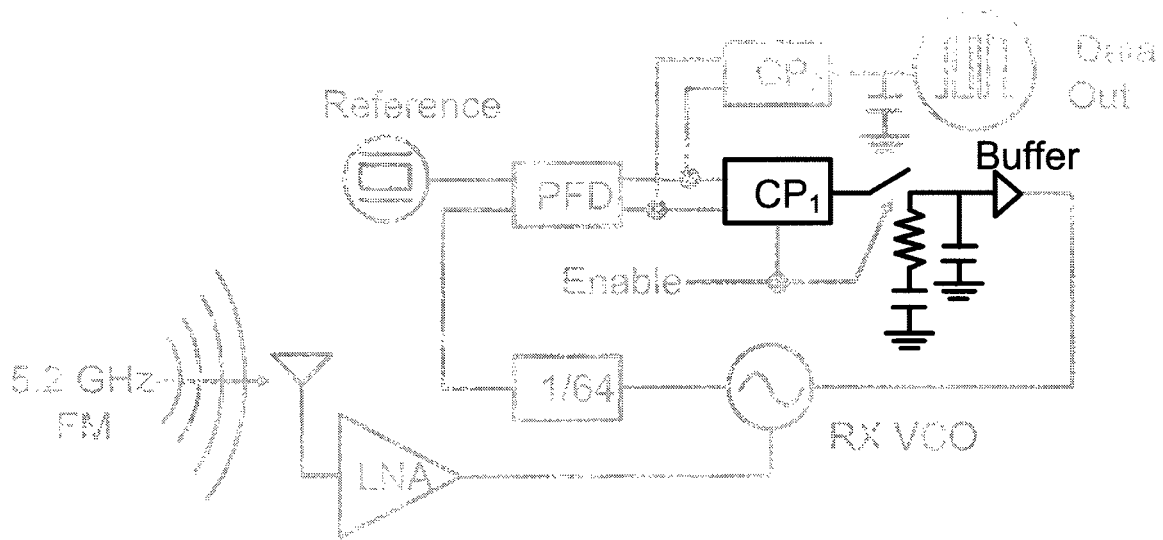


Figure 6.1: Lock-and-Roll RX Components Enabling Open-Loop Mode

The author notes that as the lock-and-roll receiver design was part of a larger project (the lock-and-roll transceiver) that included the lock-and-roll transmitter circuit implemented by colleague Victor Karam [52], most of the blocks discussed in this particular chapter are common to both the RX and TX test chips. The PFD, the tunable loop filter, and the divider circuit (discussed in Chapter 7) were designed and

implemented by Karam while the charge pump, the loop-filter buffer and switch are the author's designs. In the end both test chips used all of these blocks successfully, and the careful design of all of these blocks is the result of much collaboration between the two designers.

6.1 Highly Adjustable Loop Filter Design

At the core of the lock-and-roll receiver loop is the classic second-order loop filter itself which is comprised, in simplified terms, of a series RC circuit ($R1||C1$) connected in parallel with a capacitor ($C2$). When the loop is closed the loop filter dictates much of the loop's transient response, affecting the settling time, noise filtering, and the loop bandwidth (recall section 3.4.3). When the loop is open, however, the sole role of the integrated loop filter is to preserve, for as long as possible, the control voltage that was achieved during the closed-loop pre-tuning mode, thus maintaining the desired center frequency of the oscillator and the correct injection-locking band. If the center frequency of the RX VCO drifts too far away from that of the TX VCO the locking band will not overlap with the modulated transmit signal and the communication system fails. Failure also occurs if the VCO's center frequency drifts by more than Δf . As such, leakage from the control voltage node must be minimized.

The capacitors in the integrated loop filter are all implemented using triple layer MIM capacitors that are available in the technology used for the lock-and-roll RX test chip. At first glance the fact that both capacitors $C1$ and $C2$ connect to VSS suggests that they might be good candidates for implementation using MOS caps, especially given the large integrated capacitances that are required for a narrow loop bandwidth. While MOS caps typically have a higher capacitance density than MIM caps, thus decreasing precious die area when used in applications such as this, they have lower Q and are generally much lossier than MIM capacitors which are formed at higher level metals with less associated parasitics. MOS caps are typically made by shorting the drain and source terminals of a MOSFET device together to form one terminal of the capacitor, while the gate forms the second terminal of the capacitor and the gate oxide serves as the dielectric. The drain and source diffusions, and their associated parasitics to the substrate, must inevitably connect to one terminal of the capacitor, and thus MOS capacitors are rarely used when one

terminal of the capacitor does not connect to a relatively benign node such as on-chip ground (i.e., VSS). Beyond the disadvantages imposed by their leakage, MOS capacitors, by the very nature of the fact that they are derived from a MOSFET, have capacitance values that vary with gate voltage. Operated in the accumulation mode or strong inversion mode their capacitance is roughly fixed, but as the gate voltage approaches the threshold voltage for the device the capacitance decreases substantially until strong inversion is achieved [64]. As the control voltage of the lock-and-roll receiver (and transmitter) will vary substantially, leakage notwithstanding, MOS caps are a poor choice for the integrated loop filter regardless of the die area that might be saved by using them. Luckily, the dual layer MIM capacitors that are available in the design kit that was used offer nearly the same capacitance density as their MOS cap counterparts, and thus little additional die area was sacrificed at the expense of achieving lower leakage and stable capacitance over voltage. All resistors used in the loop filter were implemented using poly resistors over triple wells in the substrate, rather than diffusion type resistors, again chosen to minimize parasitics and associated leakage paths.

Given the somewhat risky nature of designing an all new, high-frequency PLL test chip in a design kit with which the author (and his colleagues) had no previous experience, the loop filter for the lock-and-roll RX test chip was designed to be highly adjustable using laser fuse options. Figure 6.2 shows the schematic for the completely integrated loop filter, including arrangement of the laser fuse options.

6.1.1 Achieving a Balance Between Fast Acquisition, Increased Stability, and Leakage Robustness

Recall from section 3.4.3 that selection of the loop filter component sizes has a strong effect on the loop's damping constant and the loop bandwidth, which in turn affect the settling time of the loop. The phase acquisition time is inversely proportional to the loop bandwidth while the frequency acquisition time is inversely proportional to the square of the loop bandwidth. Thus a larger loop bandwidth makes for a faster settling loop. Unfortunately, as with most RFIC design issues, the tradeoffs require much thought and delicate balancing. If the loop bandwidth is made excessively large in favour of quick settling, the loop will become unstable, recall equation (3.5),

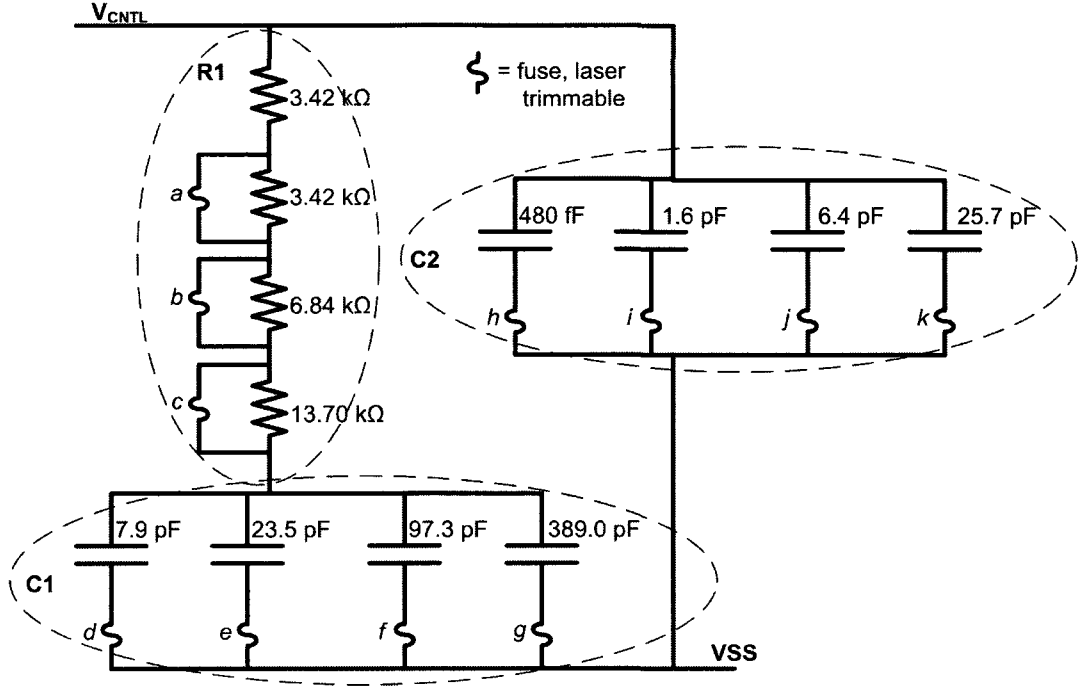


Figure 6.2: Lock-and-Roll RX Tunable Loop Filter Schematic

and as a large loop bandwidth requires the use of small loop filter capacitors, while such a scenario is good for conserving die area it makes for a loop filter voltage which is very susceptible to leakage in open-loop operation. Recognizing that the fundamental equation defining the properties of a capacitor is $C = Q/V$, where Q and V are the charge stored (in Coulombs), and the voltage across (in volts), respectively, a capacitor with capacitance C (in Farads), one can clearly conclude that small capacitors store less charge to maintain the same voltage as larger capacitors, and are thus quicker to discharge when faced with a leakage current (I), where $I = \partial Q / \partial t$.

A free behavioural-model-based simulator for PLL design, called “PLL” [66], previously available from Eagleware Corporation (now Agilent), can be used to estimate the second-order loop filter components necessary to achieve a required loop bandwidth given the charge pump current and targeted loop damping constant, ζ , as inputs. In the case of the lock-and-roll receiver, the nominal charge pump current is $I_{CP} = 100 \mu A$, and $\zeta \approx 1$. The software tool essentially solves, simultaneously, the

mathematical expressions that govern the response of this type of loop which can be approximated [51], as

$$\omega_n = \sqrt{\frac{I_{CP}K_{VCO}}{2\pi C_1}} \quad (6.1)$$

where ω_n is commonly referred to as the natural frequency of the loop,

$$\zeta = \frac{R_1}{2} \sqrt{\frac{I_{CP}K_{VCO}C_1}{2\pi}} \quad (6.2)$$

where ζ is the damping constant of the loop, and

$$\omega_{3dB} = \omega_n \sqrt{1 + 2\zeta^2 + \sqrt{4\zeta^4 + 4\zeta^2 + 2}} \quad (6.3)$$

where ω_{3dB} is the 3 dB loop bandwidth which can be approximated, for $\zeta < 1.5$, as

$$\omega_{3dB} \approx (1 + \zeta\sqrt{2})\omega_n. \quad (6.4)$$

Table 6.1 outlines the estimated loop bandwidths than can be achieved by exercising various combinations of laser fuse options within the integrated loop filter circuit shown in Figure 6.2. By no means is Table 6.1 an exhaustive list of the possibilities, but it shows the flexibility and wide range of bandwidths that can be achieved. The default setting, expected to yield a loop bandwidth of $\omega_{3dB} \approx 3.42$ kHz, will be the slowest of the various options when it comes to loop settling, yet it will make for the most stable loop and most importantly, it will likely be the most robust filter to control voltage droop when operated in the open-loop state given the large capacitors and the charges they accumulate.

Table 6.1: Lock-and-Roll RX Loop Filter Laser Options

R1	C1	C2	Loop Bandwidth	Fuses to be Exercised
3.42 k Ω	518.0 pF	34.2 pF	$2 \pi * 215$ kHz	default, no fuses cut
6.84 k Ω	129.0 pF	8.5 pF	$2 \pi * 425$ kHz	<i>a, k, g</i>
13.70 k Ω	33.2 pF	2.1 pF	$2 \pi * 850$ kHz	<i>a, k, g, b, f, j</i>
27.40 k Ω	7.9 pF	480 fF	$2 \pi * 1750$ kHz	<i>a, k, g, b, f, j, c, e, i</i>

One can quickly (and roughly) estimate the effects of leakage current on the performance of the receiver using the default loop filter arrangement. Assuming that the default loop filter can be approximated as a single 550 pF integrated capacitor from the point of view of tolerating leakage current (such an assumption is arguable given the RC decay profile inherent of R1C1, yet with the leakage current being drawn from an unknown node a more precise analysis is difficult), an estimate of the leakage current limit required to allow the target of 250 bits received (recall section 3.4.1), per closed/opened loop cycle, can be calculated. Recall from section 4.2.7 that the measured $K_{VCO,meas} \approx 0.7$ GHz/V and that from section 4.2.5, the estimated one-sided locking bandwidth (which cannot be precisely measured) was roughly 1.2 MHz. Thus, for the center frequency of the VCO to drift by enough that the injection-locking band no longer covers the modulated input signal with $\Delta f = 500$ kHz, the required droop on V_{CNTL} is about 1 mV. However, when V_{CNTL} droops by merely 715 μ V, with $K_{VCO,meas} \approx 0.7$ GHz/V the center frequency of the VCO has moved by 500 kHz and the PFD no longer recognizes the divider output as being faster and slower than the reference with every bit change, and thus 715 μ V of droop causes the RX to fail when $\Delta f = 500$ kHz. At a communication rate of 5 kb/s, transmitting 250 bits takes approximately 50 ms. Assuming the closed-loop mode pre-tunes $V_{CNTL} \approx 600$ mV, the charge stored on the 550 pF capacitance can be calculated as $Q = (600e - 3) * (550e - 12) \approx 330e - 12$ Coulombs, and for a droop in voltage of 715 μ V, a leakage current would have to rob the capacitor of $(715e - 6) * (550e - 12) \approx 395e - 15$ Coulombs. To do so in 50 ms would require a leakage current of $(395e - 15)/(50e - 3) \approx 8$ pA. While the calculated leakage limit would appear to be very small, 8 pA is actually much larger than the leakage in a MIM capacitor, and still larger than the leakage expected in a MOS capacitor at certain bias levels. Nevertheless, the result attests to the importance of minimizing the sources of leakage on the loop filter and to using the tightest loop bandwidth that can be afforded given the tradeoff with settling time.

6.1.2 Loop Filter Layout with Laser Trim Tunability

Figure 6.3 shows the layout of the tunable loop filter that occupies roughly 15% of the overall die area for the lock-and-roll receiver test chip due to the large capacitances

that were implemented. The passivation openings that were left to facilitate easier laser cutting of the fuse options are clearly visible. The note “pwr” that appears over every passivation opening merely results from a design rule check (DRC) work-around that was implemented, as the kit required the passivation openings to be labelled as a “pwr” or “gnd” node at the top level in order to be ignored from the automated routine that otherwise looked for adequate electro-static discharge (ESD) circuitry connected to every top metal route that was exposed to the outside world.




Figure 6.3: Lock-and-Roll RX Tunable Loop Filter Layout

Clearly the loop filter for the lock-and-roll receiver is highly tunable if required. While having the laser trim options is helpful from a test chip perspective, a production chip making use of the topology would only implement one filter option (after deciding on the most optimal design through extensive testing), and the layout could be compacted and further optimized. Despite the large loop filter capacitors, the estimated 8 pA leakage limit (for receiving 250 bits per cycle) attests to the importance of further measures for protecting the loop filter’s control voltage from low-impedance paths in the open-loop state.

6.2 Loop Filter Switch Design

Even with the relatively high-impedance off state of the charge pump design (which is discussed in section 6.4), the loop filter requires additional isolation from the output node of the charge pump during open-loop operation in order to achieve satisfactorily low charge leakage. Figure 6.4 shows the schematic of the loop filter switch that was implemented. At the core of the switch are transistors M1 and M2 which are complementary and form a traditional CMOS transmission gate. Both are sized for low R_{on} impedance, though they are not made excessively large as doing so would increase the risk of charge injection at the time of switching.

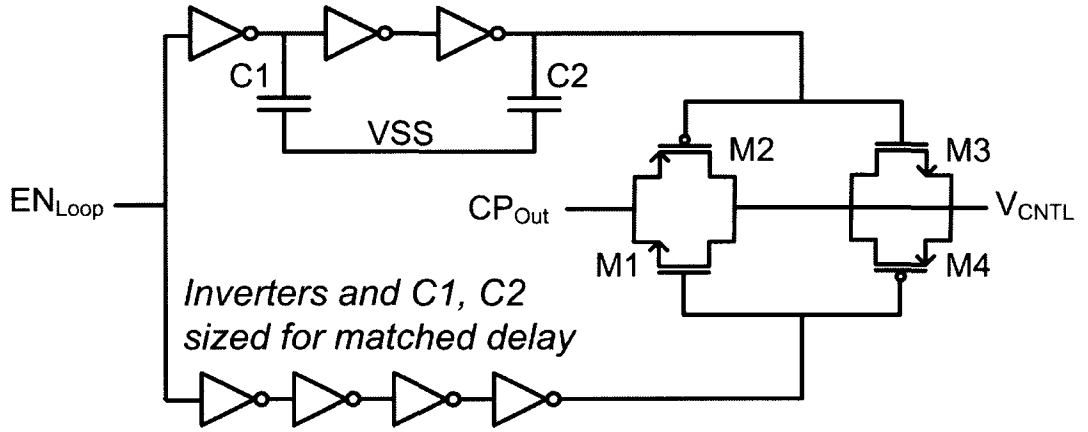


Figure 6.4: Lock-and-Roll RX Loop Switch Schematic

6.2.1 Charge Injection and Mitigating the Effects with Dummies

The phenomenon of charge injection refers to the redistribution of channel charge to the drain and source nodes of a MOSFET at the instance the device is switched from the “on” to “off” state, or equivalently when V_{GS} is switched from $V_{GS} > V_{TH}$ (which introduces a channel below the gate), to $V_{GS} < V_{TH}$ (where the channel disappears as charge redistributes). Given the importance of the subject with regards to the performance and accuracy of switched capacitor circuits, there are many good texts that explain the issue, its effects, and methods of mitigating them [67], [68], [69]. The process is roughly depicted in Figure 6.5 for an NMOS device, where one can see

that having $V_{GS} > V_{TH}$ attracts negative charge under the gate to form the channel. When the NMOS is switched off such that $V_{GS} < V_{TH}$, the charge that was previously

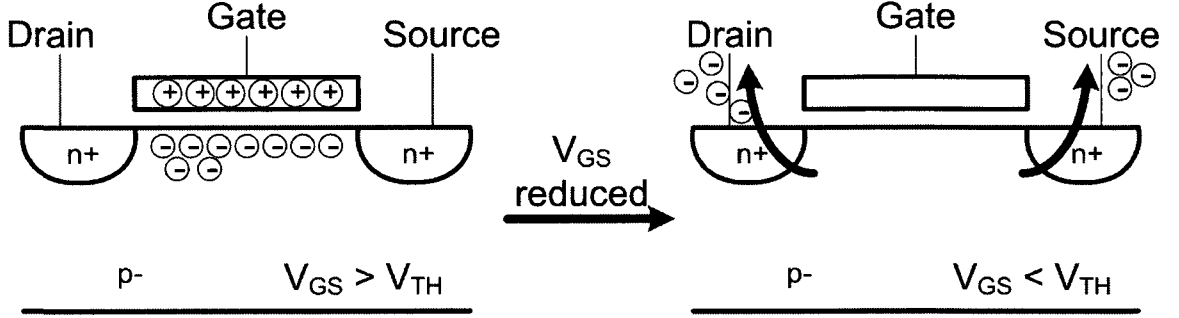


Figure 6.5: Channel Charge Injection Effect in NMOS

built up under the gate dissipates and largely exits the device through the drain and source connections. Depending on the shape of the channel (which is largely dictated by V_{DS}) the charge may or may not exit the device evenly from the drain and source, typically a 50%:50% split is assumed on the basis of low R_{on} and low V_{DS}), and assuming more or less similar impedances seen by the drain and source. The same scenario exists for PMOS devices, where the polarity of the charges in Figure 6.5 are reversed. In the case of the lock-and-roll receiver loop, the drains of M1 and M2 in the loop filter switch (see Figure 6.4) are connected directly to the loop filter. When the switch is opened, roughly half of the charge that was previously collected under the gates of M1 and M2, if no further steps were taken, would accumulate on the loop filter capacitors and instantly change the control voltage that is held on the filter, and consequently the center frequency and the locking band of the VCO would be altered from their precisely tuned states. The common approach to mitigating the effects of charge injection is the use of a “dummy” transmission gate that is connected as shown in Figure 6.4, whereby transistors M3 and M4 make up the dummy switch. Note that the drain and source terminals of M3 and M4 are shorted together, such that the switch is essentially always closed and therefore does not affect the behaviour of the overall circuit. Note also that the gates of M3 and M4 are connected to alternate polarities of the enable signal (compared to M1 and M2). Transistor M3 is sized to be exactly half the width of transistor M1, and transistor M4 is sized to be exactly half the width of transistor M2. The concept assumes that with half the gate area of the devices used in the real switch, and with the gate polarities reversed, the half

of the channel charge on M1 and M2 that exits the drain terminals upon closure of the switch should be completely absorbed by M3 and M4 where a channel is being induced at exactly the same instance in time. M3 should absorb the charge from M1, and M4 the charge from M2, eliminating (or at least greatly reducing) the amount of charge that is added or subtracted from the loop filter that is connected to the output of the switch. Lastly, note from Figure 6.4 that a chain of inverter cells is used to balance the timing of the control signals that appear on the gates of M1 and M2, where capacitors C1 and C2 help to accomplish the required delay by slowing down the slew rates of the inverters that drive them. The signals at the gates of M1 and M2 aren't necessarily balanced to achieve a perfect mid-rail crossing, but are rather adjusted such that given the sizes of M1 and M2 (which were sized to minimize and balance $R_{on,p}$ and $R_{on,n}$), their turn-on and turn-off times are as close as possible.

Noting that a leakage current on the order of 8 pA or more would reduce the number of bits that can be received at 5 kb/s to less than 250 bits/cycle, the use of the loop filter switch in combination with the high-impedance charge pump and loop-filter buffer circuit is clearly warranted. While the PFD circuit is not directly involved with opening the loop and preserving V_{CNTL} , understanding its design facilitates a better understanding of that of the charge pump circuit.

6.3 PFD Circuit Design and Behaviour

The phase-frequency detector circuit used in the lock-and-roll receiver loop is a simple and traditional three-state digital CMOS implementation [51]. The schematic is shown in Figure 6.6, where the complete circuit is comprised of two flip flops and an AND gate. The flip flops have their D inputs tied to the positive VDD supply and are clocked by the rising edges of the output of the reference signal (Ref input) and the divider (Div input), respectively. The rise of the Up output signal results from a rising edge of the Ref input signal, insinuating that the charge pump should “pump up” V_{CNTL} such that the VCO will increase in frequency and better match the phase of the reference. Conversely, a rise of the Div input signal triggers the Down output to rise, insinuating that the charge pump should “draw down” V_{CNTL} such that the VCO will decrease in frequency and match the phase of the reference. In both cases, the flip flops are reset, causing Up and Down outputs to fall, when both Up and

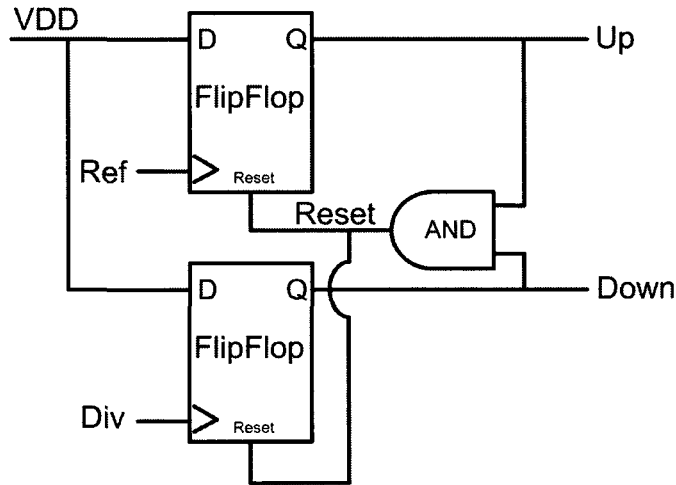


Figure 6.6: Lock-and-Roll RX Tristate PFD Schematic

Down are driven high simultaneously (and only momentarily given that resetting the flip flops resets their outputs low). Prior to achieving frequency lock, if the reference signal is much higher in frequency than the output of the divider, the Up output will be high most of the time with the Down output only pulsing high momentarily as the flip flops are reset. The reverse scenario is true when the reference frequency is much lower than that of the divider output.

6.3.1 Trading-Off Lower Acquisition Time for Reduced Loop Filter Leakage

Assuming, for sake of discussion, that the loop starts out with the reference frequency higher than that of the divider output, as frequency acquisition is approached, the duty cycle of the Up signal will be reduced, thereby slowing down the rate of acquisition. Additionally, the frequency of occurrence of cycle slips, whereby the duty cycle of the Up signal is reset to nearly 0% and must recover over time, increases, further slowing the acquisition rate. For this reason a five-state PFD is often preferred over

the three-state PFD, where if there are two rising edges of the Ref input between rising edges of Div, the PFD kicks into a “turbo Up” state that activates a second set of Up and Down outputs that are connected to a second charge pump, resulting in faster acquisition. The second charge pump current is often higher than that of the primary charge pump, further increasing the rate of acquisition. A complementary “turbo Dn” mode makes up the fifth state of such a PFD machine. While one could certainly argue the benefits of increasing the acquisition rate of the lock-and-roll RX loop in terms of lowering the average current consumption of the topology, recall equation (3.4), the second charge pump that would be required would inevitably increase the leakage from V_{CNTL} during open-loop operation, and so the three-state PFD approach was adopted with a single charge pump being connected to V_{CNTL} .

6.3.2 PFD Dead-zone

Worthy of note is that both the three-state and five-state PFD topologies, implemented with CMOS digital logic cells, suffer from what is commonly referred to as a “dead-zone”. The term dead-zone refers to the fact that when the Ref and Div inputs are very close in phase, the PFD is essentially unable to further influence the charge pump and V_{CNTL} , and so a static phase error will always remain. The dead-zone arises, primarily, from the finite delay on the rising edge of the flip flop output signal, which usually has a lower slew rate than that of the AND gate that drives the Reset signal. For small phase offsets between Ref and Div, the flip flop output may not have risen beyond the threshold voltage of the charge pump input before the AND gate has driven Reset high and the flip flop output is reset. The result is no change in the output current from the charge pump, and the loop does not correct for the small phase offset between Ref and Div. Careful design can optimize the slew rates of the digital cells to minimize the dead-zone but some finite phase offset is to be expected. As precise phase matching is not required by the lock-and-roll receiver loop, only frequency tuning of the RX VCO to that of the TX VCO for the purpose of overlapping the injection-locking band with the modulated input signal, little time was spent refining the slew rates of the digital cells although to a first order, the dead-zone was minimized through careful combined simulation of the PFD with the charge pump.

6.4 High-Impedance Charge Pump Design

The charge pump for the lock-and-roll receiver loop is designed so that it can be disabled and put into a high-impedance off state in order to minimize leakage from the loop filter during open-loop operation. The design is used in duplicate on the receiver test chip, as shown in Figure 6.1, whereby the primary charge pump (CP_1) is used in closed-loop operation to pre-tune the VCO, and the secondary charge pump (CP_2) is used in open-loop operation to demodulate the received data. Figure 6.7 shows the schematic for the charge pump circuit.

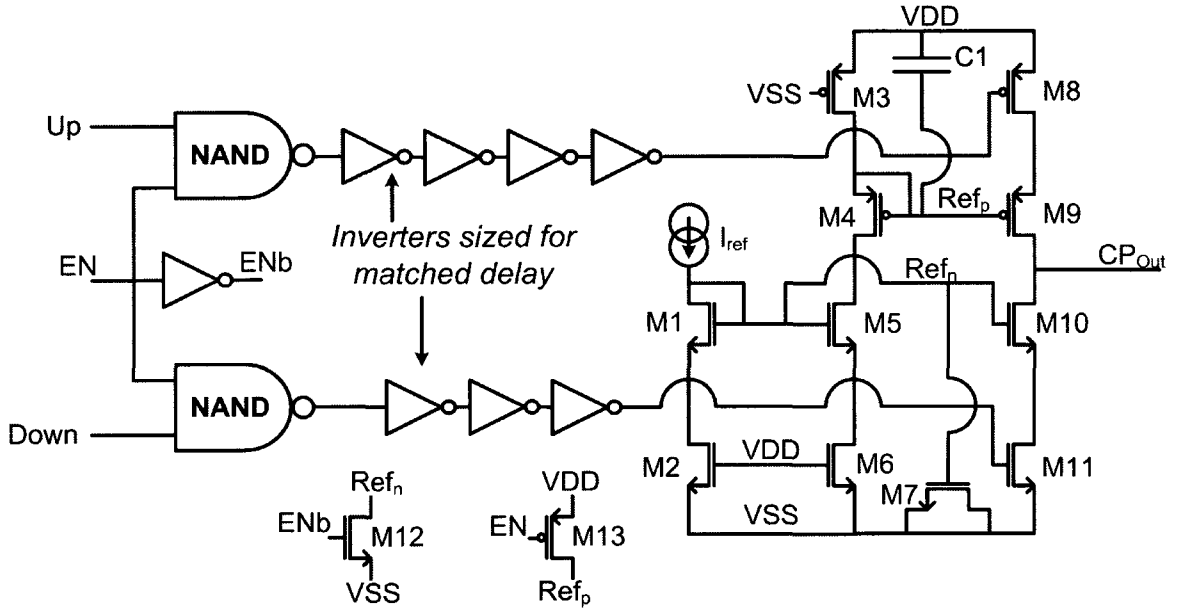


Figure 6.7: Lock-and-Roll RX High-Impedance Charge Pump Schematic

6.4.1 Matching Up/Down Pump Profiles, Simulated Output Current Response

When CP_1 is enabled during closed-loop mode, the charge pump translates the *Up* and *Down* input signals into current being supplied or drawn from the loop filter that is connected, through the loop filter switch, to the charge pump output. Transistors M5 and M10 form current mirrors with diode connected M1 that passes the reference current that sets the charge pump output current, I_{CP} , according to the ratio of the

current mirrors. Both charge pumps in the lock-and-roll RX use a 1:1:1 mirror ratio such that $I_{CP} \approx I_{ref}$. Transistors M2 and M6 are sized the same as M11, such that the V_{DS} drop across all three is similar, providing better current matching between the three branches of the current mirror. Where M11 is required to switch M10 on during a high pulse of the Down input, by drawing down the source of M10 such that $V_{GS,10} \approx V_{GS,1}$ and I_{ref} is drawn from output CP_{Out} , M2 and M6 are included to help match the source voltages of M1, M5, and M10. Transistor M3 plays a similar role, matching the voltage drop across M8, which switches on M9 during a pulse of the Up input, thereby causing I_{ref} to be supplied to CP_{Out} . Transistor M9 forms a current mirror with diode connected M4 which translates the NMOS-based reference current to a PMOS-based one. Nodes Ref_n and Ref_p are sensitive to switching noise that results primarily from the parasitic gate-to-source capacitances of M10 and M9, respectively, and as such the nodes are decoupled to the supplies by means of the MOS cap formed by M7 and MIM cap C1. Note that Ref_n is decoupled to VSS given that (ignoring M2, M6, and M11) it is the gate voltage on M10 relative to VSS that sets the Down output current, whereby the gate voltage of M9, relative to VDD, is what largely dictates the Up output current. Transistors M9 and M10 are carefully sized to minimize output noise and to have matched slew rates, and the sizing of M1, M4, and M5 follows. Transistors M8 and M11 are sized to have matching (and fast) slew rates, where the two chains of inverters that drive their gates from the NAND gates at the input are carefully sized along with M8 and M11 to balance the turn-on times of M8 and M11.

Simulating the circuit with both Up and Down input signals pulsed simultaneously is very useful for design purposes, whereby a well-balanced design will see nearly no current supplied to or drawn from CP_{Out} under such conditions. If the turn-on times of M8 and M11 are well matched, along with mirror devices M9 and M10, all current through the output stage will transfer directly from VDD to VSS. The cascoded nature of M10 and M11, along with M8 and M9, provides a high output impedance when the charge pump is enabled, which like a cascode current mirror design, helps to maintain constant output current under varying output voltage scenarios. Such behaviour is highly sought after for a charge pump design where the output voltage will vary largely depending on the state of the loop, and maintaining a

relatively constant output current equates to steadied loop dynamics at various control voltages given the dependence of ω_n , ζ , and ω_{3dB} on I_{CP} (recall section 6.1.1 and equations (6.1), (6.2), and (6.4)). While having a high output impedance contributes to a more constant output current under varying output voltage conditions, achieving a good balance between the PMOS and NMOS sections of the circuit will equate to a balanced output current profile. Figure 6.8 shows the simulated output current profile of the charge pump circuit. The Pump Up curve in Figure 6.8 is achieved by




Figure 6.8: Simulated Charge Pump Output Current vs. V_{CTL}

forcing the Up input of the circuit to VDD, while forcing the Down input to VSS, and simulating the output current that is supplied to the output node, with the DC output voltage swept from VSS to VDD. Conversely, the Pump Down curve is the resulting current drawn from the output when the Down input is connected to VDD with the Up input connected to VSS while the output voltage is swept. Ideally the curves would be as flat as possible, indicating high output impedance, and would show the same current. That being said, the simulated result shows good balance between Up and Down currents, with a near perfect balance at the mid-rail voltage for the process (600 mV), and with a maximum Up/Down current mismatch of roughly 15% inside of $200 \text{ mV} < V_{CTL} < 1.00 \text{ V}$. The slope of the curves suggest an output impedance of roughly $35 \text{ k}\Omega$ when operated in either the Up or Down modes. Below a 200 mV output voltage, transistors M10 and M11 fall out of saturation and

the output current decreases. The same is true of transistors M8 and M9 for output voltages beyond 1.00 V.

When the charge pump is disabled (CP1 is disabled in open-loop mode and CP2 is disabled in closed-loop mode), transistors M12 and M13 pull nodes Ref_n and Ref_p to VSS and VDD respectively, disabling M9 and M10. M8 and M10 are disabled regardless of the state of the Up or Down input signals by nature of the input NAND gates that are tied to the enable (En) signal, and the resulting output impedance of the disabled charge pump is high.

Note in Figure 6.1 that the connections between the PFD and CP2 are reversed with respect to the connections between the PFD and CP1. This phase reversal results in the final output bitstream matching the polarity of the input bitstream to the transmitter, given the compensating nature of the PFD outputs when taken in the context of closed-loop operation.

6.4.2 Minimizing Charge Injection and Leakage Through Design and Layout

Worthy of discussion is that the cascode charge pump could just as easily have been designed with the switches placed towards the center of the circuit, in essence swapping the positions of M8 with M9, and M10 with M11. Doing so would have pretty well eliminated the need for transistors M2, M3, and M6 as the current mirrors would have been firmly referenced to VSS and VDD respectively. The change from the implementation that was adopted seems to greatly simplify the schematic and the task of matching Up and Down currents. The drawback with this approach, however, is that placing the switches (M8 and M11) closer to the output node greatly increases the charge injection they contribute (when switched) to node CP_{Out} , based on the same principle presented in section 6.2.1. Such charge injection leads to clock spurs on the VCO output during closed-loop operation (not of great concern for the lock-and-roll receiver loop but definitely of concern for most PLL and synthesizer reference circuits), and possibly affects the charge stored on the loop filter at the moment the loop is switched from closed to open-loop operating modes. By placing the switches towards the outside of the circuit (i.e., closer to VDD and VSS), the switches see a low-impedance path (to VDD or VSS) at their source and the charge is readily

transferred there when the switches are turned off. Figure 6.9 shows the layout of the charge pump circuit used in the lock-and-roll receiver. The layout maps well visually




Figure 6.9: Lock-and-Roll RX High-Impedance Charge Pump Layout

with the schematic, whereby the signals travel from left to right. While not clearly visible in Figure 6.9, transistors M1, M5, and M10 are carefully laid out close together, using the same orientation, and with dummy cells in order to improve their matching. Common centroid layout techniques [64] were not used in order to keep the routing simple, thus reducing clock feedthrough, and the devices could not be interdigitated due to the fact that they do not share the same drain or source connections. By using an even number of device fingers to form each transistor, the dummy finger cells on each side of the layout for M10 can be connected between the source terminal of M10 and VSS. While the mention of such layout details might seem trivial, even the finite

off impedance of a dummy finger cell (with gate connected to VSS) would otherwise contribute to increased leakage from the loop filter during open-loop operation, if indeed it were connected to the drain of M10. Transistor M9 and its dummies are laid out in much the same way.

6.4.3 Measured Output Current Response

Figure 6.10 shows the output current profile of the charge pump that was measured in the lab for comparison purposes with the simulated results shown in Figure 6.8. To generate the Pump Up curve, the VCO's supply was disabled such that the divider's output never toggles. As the PFD, divider, and charge pump share the same on-chip supply the divider could not be disabled. To generate the Pump Down curve, the reference signal was disabled with the VCO enabled. The measured results suggest




Figure 6.10: Measured Charge Pump Output Current vs. V_{CTL}

a lower output impedance than simulated, roughly $20 \text{ k}\Omega$ instead of $35 \text{ k}\Omega$, and therefore higher current mismatch as the output voltage deviates from the nominal mid-rail point of 600 mV. That said, the Up and Down currents are nearly perfectly matched at $V_{\text{CTL}} = 600 \text{ mV}$, and track the $100 \text{ }\mu\text{A}$ reference current well given the 1:1:1 mirror ratios in the core of the circuit. While the high-impedance off state of the charge pump and the loop filter switch both help to reduce the leakage from the loop filter during open-loop operation, perhaps the most important tool in the

fight against leakage is the loop-filter buffer circuit that was implemented on the lock-and-roll receiver test chip.

6.5 Unity-Gain Loop-Filter Buffer Design

With the overall design of the lock-and-roll receiver loop more or less complete, including the high-impedance charge pump, the integrated loop filter, which makes exclusive use of high-quality MIM capacitors and poly resistors, and the unique loop filter switch, transient simulations of the open-loop behaviour suggested that further steps were necessary to reduce the leakage current drawn from the loop filter. Careful analysis identified that the accumulation-mode varactors [70] used in the tank of the VCO circuit to give the block its required frequency tunability, provided a substantial leakage path during open-loop operation. Much like MOS caps, which are similarly structured devices, accumulation-mode varactors are well known to be lossy despite their excellent range of capacitance tunability. Recall from section 4.2.4 that the lossy well connections (anodes) of the accumulation-mode varactor diodes were connected to V_{CNTL} in order to better isolate the VCO's tank circuit from the parasitics associated with that side of the diodes. The design of a unity-gain, amplifier buffer circuit to be inserted between the loop filter and the VCO's V_{CNTL} input signal followed. Figure 6.11 shows the schematic for the loop-filter buffer circuit that was designed, in essence, to disconnect the varactor leakage from the integrated loop filter. Referring to Figure 6.11, transistors M1 and M2 form the common-source amplifier with active load, whereby M1 and M2 are sized to provide unity gain. Note that the introduction of a common-source, inverting amplifier to the loop means that the Up and Down connections between the PFD and the charge pump must now be reversed in order to preserve the correct loop polarity. As the test chip was designed such that the loop-filter buffer could be used optionally, a mux circuit was implemented (not shown in Figure 6.1) between the PFD and CP1 such that connections between these blocks are configurable. The implementation of the MUX circuit and its use is explained in Chapter 7. Transistor pairs M3 and M4, M5 and M6, and M7 and M8 form CMOS transmission gates that allow the buffer circuit to be connected in series between the loop filter and the VCO, or shorted out altogether if its use is not desired. Given the latter option, note that the supply for the loop-filter buffer is the enable signal




Figure 6.11: Lock-and-Roll RX Unity-Gain Loop-Filter Buffer Schematic

for the buffer itself. Finally, note that unlike the transmission gate that opens and closes the loop (recall section 6.2), the transmission gates in the loop-filter buffer are intended to remain open or closed during both the open-loop and closed-loop modes of operation, depending on whether the loop-filter buffer is desired or not. Therefore, no dummy cells are used with the transmission gates in the circuit as charge injection should not be a concern.

The ultimate test of the effectiveness of all the design strategies that were employed to minimize the leakage current drawn from the loop filter is the simulated and measured VCO drift rates when the loop is operated in the open-loop mode.

6.6 Simulated vs. Measured VCO Drift Rates

Table 6.2 compares the open-loop performance of the lock-and-roll receiver test chip, simulated both with and without the loop-filter buffer enabled in the loop, and measured with the loop-filter buffer enabled. The second column of data represents the open-loop time margin, or in other words how long the loop can remain open before

the VCO drift that results causes the RX to fail according to one of the two mechanisms outlined at the beginning of this chapter, recall that at a data rate of 5 kb/s, 50 ms is required to transmit the target of 250 bits per cycle assuming $K_{VCO} = 700$ MHz/V (see section 3.4.1). The third column in Table 6.2 summarizes the resulting VCO drift rate, the fourth column summarizes the resulting rate of droop on the loop filter voltage, and the fifth column highlights K_{VCO} . The sixth column of data is an extrapolated result, estimating the leakage current that is likely experienced by the loop filter circuit using the same assumptions and math that are outlined in the estimated limit from section 6.1.1.

Table 6.2: Open-Loop Performance, Simulation vs. Measurement

Result	OL Margin	VCO Drift	V_{CNTL} Droop	K_{VCO}	Leakage
Sim Buffer	51 ms	10 Hz/ μ s	9 μ V/ms	1.1 GHz/V	5 pA
Sim no Buffer	240 μ s	2.1 kHz/ μ s	1.9 mV/ms	1.1 GHz/V	1.0 nA
Measured	22 ms	23 Hz/ μ s	33 μ V/ms	700 MHz/V	18 pA

Table 6.2 clearly shows that the simulated results suggest the loop-filter buffer is crucial to achieving adequate open-loop performance, where by disabling it the leakage is so severe that only one bit could be received at a data rate of 5 kb/s before the the RX would fail. The simulated droop rate of V_{CNTL} is reduced by about 200 times when the loop-filter buffer is enabled compared to when it is disabled. The nominal simulated result, with the loop-filter buffer enabled, is capable of demodulating 250 bits at 5 kb/s before the loop needs to be closed again and V_{CNTL} refreshed. The third row of data, which highlights the nominal measured performance, resulted as follows.

Recall that as the LNA inputs are connected directly to differential input pads for the purpose of characterizing the input match (which does not match the LNA to 50 Ω), the exact input amplitude of the modulated signal applied to the LNA cannot be determined. As such, the precise locking bandwidth of the VCO for a given input signal strength cannot be determined from measurement. Given the 1.2 MHz locking bandwidth that was predicted for the design using Adler's equation (recall section 4.2.5) for a 1.75 m separation between the RX and a TX using a patch antenna, the

input signal strength from an FM signal generator was adjusted in the lab such that the measured locking bandwidth (with the FM input applied) for the RX VCO was 1.2 MHz. To accomplish this step, the VCO was operated in the open-loop mode with V_{CNTL} driven from an external DC source to precisely pre-tune the VCO center frequency to 5.200 GHz (at zero span on the spectrum analyzer). With the input signal amplitude adjusted appropriately, the external DC source was removed from V_{CNTL} and the RX loop enable signal was driven with a square wave input, causing the RX to operate in the open-loop and closed-loop modes with a 50% duty cycle. Figure 6.12 shows the waveforms that resulted on the oscilloscope in the lab, whereby the yellow waveform at the top of the screen is the loop enable signal, the green signal at the bottom is the bitstream applied to the signal generator that is creating the FM input signal, and the blue signal in the center is the output of the lock-and-roll receiver. Ignoring the noise on the output signal, which is explained in detail in




Figure 6.12: Lock-and-Roll RX Measured Drift and Point of Failure

Chapter 7, the output bitstream from the lock-and-roll RX matches well with the input bits to the modulator for the first 22 ms after the loop is opened. Beyond 22 ms the output of the RX appears random and the VCO's center frequency has likely

drifted beyond 500 kHz, where the FM input uses $\Delta f = 500$ kHz. From the 22 ms window that was measured, along with the measured K_{VCO} from section 4.2.7, the drift rate of the VCO, the droop rate of V_{CNTL} , and the leakage current are calculated.

With the loop-filter buffer disabled during measurement, demodulation proved impossible unless the input signal amplitude was made excessively large to force injection locking given the higher leakage.

6.7 Suspected Leakage from Tie-Down Diodes

The measured result suggests that while the lock-and-roll RX architecture is successful at demodulating the FM input signal, only 110 bits can likely be received at 5 kb/s (instead of the target of 250 bits) before the loop must be closed again and V_{CNTL} refreshed. Likely there is more leakage from the loop filter than what was simulated. Worthy of note is that the simulated result uses the (pre-layout) schematic netlist, as simulating merely the initial closed-loop pre-tuning behaviour takes upwards of 1 week using the full extracted netlist (post-layout). The effects of parasitics, therefore, are not represented by the simulations.

This difference noted, there is likely another source of error. The design kit that was used for the lock-and-roll receiver test chip has a unique electrical requirement that the gates of every CMOS device and every MIM capacitor plate be connected to the substrate with a “tie-down” diode. Unlike regular antenna rule checks [64] where the ratio of connecting metal area to that of the gate must not exceed a specified limit which can usually be met by using higher metal “jumper” wires in layout to isolate long routes from the gate at the metal level in question, the design kit that was used mandates connection to the substrate at metal 1 or lower for CMOS gates, and imposes similar demands on all MIM capacitor plates. As a result of this rule, the very large area of the MIM capacitors used in the loop filter layout (recall Figure 6.3) necessitates that the V_{CNTL} node be connected to the substrate through a large PN diode having dimensions of roughly $50\ \mu\text{m}$ by $300\ \mu\text{m}$. To further complicate matters, the diodes are not extracted by the extraction tool as circuit elements (at the time of design and layout) and therefore they do not get added to the extracted netlist. The result is that even though running a complete simulation of the extracted PLL in

open-loop mode is prohibitively complicated, simplified simulations of the extracted loop filter alone would not show the effect of the tie-down diode as a source of leakage.

Nevertheless, despite the increased leakage that the measured result suggests, Figure 6.12 and Table 6.2 attest to the fact that the measures taken to isolate the loop filter voltage from numerous sources of leakage were fruitful, as the lock-and-roll receiver can successfully demodulate significant data before needing the loop to be closed and V_{CNTL} refreshed.

6.8 PLL Components that Enable Open-Loop Operation Summary

The unique topology of the lock-and-roll receiver demands a PLL loop with components that are optimized to operate in both the conventional closed-loop mode, as well as the unconventional open-loop mode. Given that the closed-loop mode of operation is merely used to pre-tune the frequency of the RX VCO, thereby aligning the center frequency and the injection-locking band with the modulated output from the transmitter, the closed-loop requirements are actually relaxed compared to those of most synthesizer or PLL reference circuits. Rather, it is the open-loop mode of operation, and the near constant V_{CNTL} that it requires, which places tough restrictions on the overall design.

In open-loop mode, and leakage currents that draw charge from the integrated capacitors in the loop filter will cause V_{CNTL} to droop over time. The effect of V_{CNTL} droop is that the center frequency and the injection-locking band of the VCO will change over time. Eventually, the drifting VCO will cause the receiver to fail due to one of two mechanisms. If the center frequency of the VCO drifts by more than Δf of the modulated input signal, the PFD circuit will no longer recognize the divider output as being indicative of a 1 or a 0 (data bit logic levels). Secondly, if the injection-locking band of the VCO drifts beyond the modulated input signal's bandwidth, the VCO will no longer be injection locked, and again, the divider output will no longer be representative of the data. As such, the loop design must be carefully optimized to minimize leakage from the integrated loop filter and the related V_{CNTL}

droop that follows. The lock-and-roll receiver implements several unique circuits and design strategies in the name of reducing loop filter leakage.

The loop filter design itself is highly configurable to yield various optional loop bandwidths with the use of a laser to cut top metal fuses, post fabrication. Such flexibility allows the test chip to be optimized for startup or noise requirements, while the default arrangement uses the largest selection of capacitors in order to maximize the charge that is held on the filter during open-loop mode, lessening the effect of leakage on the receiver's performance. The filter is completely implemented using high-quality MIM caps and poly resistors over triple wells, rather than MOS caps and diffusion resistors which would contribute to more leakage.

The transmission gate used to open and close the loop borrows from the philosophy of switched capacitor designers, whereby a half-sized dummy switch is used to absorb most of the charge injected towards the loop filter at the moment the switch is opened.

By using a three-state PFD and one charge pump in the loop, rather than a five-state PFD and two charge pumps which would reduce startup time and the current consumption of the system, control voltage droop is minimized yet again.

The charge pump used in the loop is designed such that it can be disabled and doing so presents a high output impedance to the transmission gate and the filter. The placement of the switches in the charge pump was chosen to minimize charge injection, regardless of the added complexity this change brought to the design, and even the connection of the layout dummies close to the output node were implemented carefully in order to maximize the output impedance of the design.

The unity-gain loop-filter buffer is a common-source amplifier with an active load and buffers the loop filter circuit from the leakage drawn by the accumulation mode varactors in the VCO. The buffer can be enabled or disabled for making comparative performance measurements, whereby the simulated benefit of enabling the buffer is a decrease in the open-loop V_{CNTL} droop rate of about 200 times. When disabled, the buffer circuit is bypassed with a transmission gate that shorts its outputs and a mux circuit between the PFD and charge pump enables the control signals to be reversed such that the loop's polarity is maintained.

All in all, while the simulated open-loop performance of the lock-and-roll receiver suggests it is capable of demodulating 250 bits of data at a data rate of 5 kb/s

before the loop needs to be closed and V_{CNTL} refreshed, the measured result suggests that the true leakage from the integrated filter is likely three times what was simulated, and that only 110 bits can be demodulated at that data rate. The increase in leakage is suspected to have been caused by the last minute placement of a large tie-down (PN) diode connecting V_{CNTL} to the substrate, as mandated at the top level by the unique electrical requirements of the design kit that was used for the test chip. As the diode has no schematic representation and is not extracted by the design kit and added to the post-layout netlist, the leakage current that is introduced by the diode was not estimated through simulation. Despite the differences between the simulated and measured open-loop performances, both results attest to the ability of the lock-and-roll receiver topology to demodulate an FM input signal, whereby numerous bits can be discerned at moderate data rates, affirming the solution's adequacy for the short-range, low-power applications for which it is targeted.

Chapter 7

The Lock-and-Roll Receiver Test Chip

This chapter focuses on the lock-and-roll receiver test chip that was implemented for demonstration of the unique receiver proposed in Chapter 3. The test chip puts into practice the IC design strategies and techniques for achieving ultra-low power consuming and fully integrated short range communications systems (manufactured in inexpensive bulk CMOS), as described in Chapters 4 through 6. While the design and layout of the individual core circuits on the test chip are covered in the previous chapters, the overall test chip topology is presented here, examining the project from a broader perspective. The design of circuits such as the divider and the Up/Down pulse mux, which are fundamental to the operation of the receiver loop but are only mentioned briefly in the preceding chapters, are covered here in more detail. Additionally, circuits and structures that were implemented simply for the purpose of testing the lock-and-roll receiver, such as the output buffers (which drive 50 Ω interfaces in some cases), and the probe de-embedding structures (for characterizing the input match) are discussed in this chapter. Figure 7.1 shows a block level diagram of the overall lock-and-roll receiver test chip that was implemented, including all auxiliary and supplemental circuits required to facilitate testing. Figure 7.2 shows a microphotograph of the complete die that was manufactured and tested. Scratch marks on some pads, caused by landing the 8-pin probes on two of the rows of pads on the die during testing, are visible in the photograph.

In addition to analyzing the circuits on the lock-and-roll receiver test chip that are not discussed in the previous chapters, this chapter reviews the test methodology that was adopted to achieve the measured results, and reviews the difficulties that

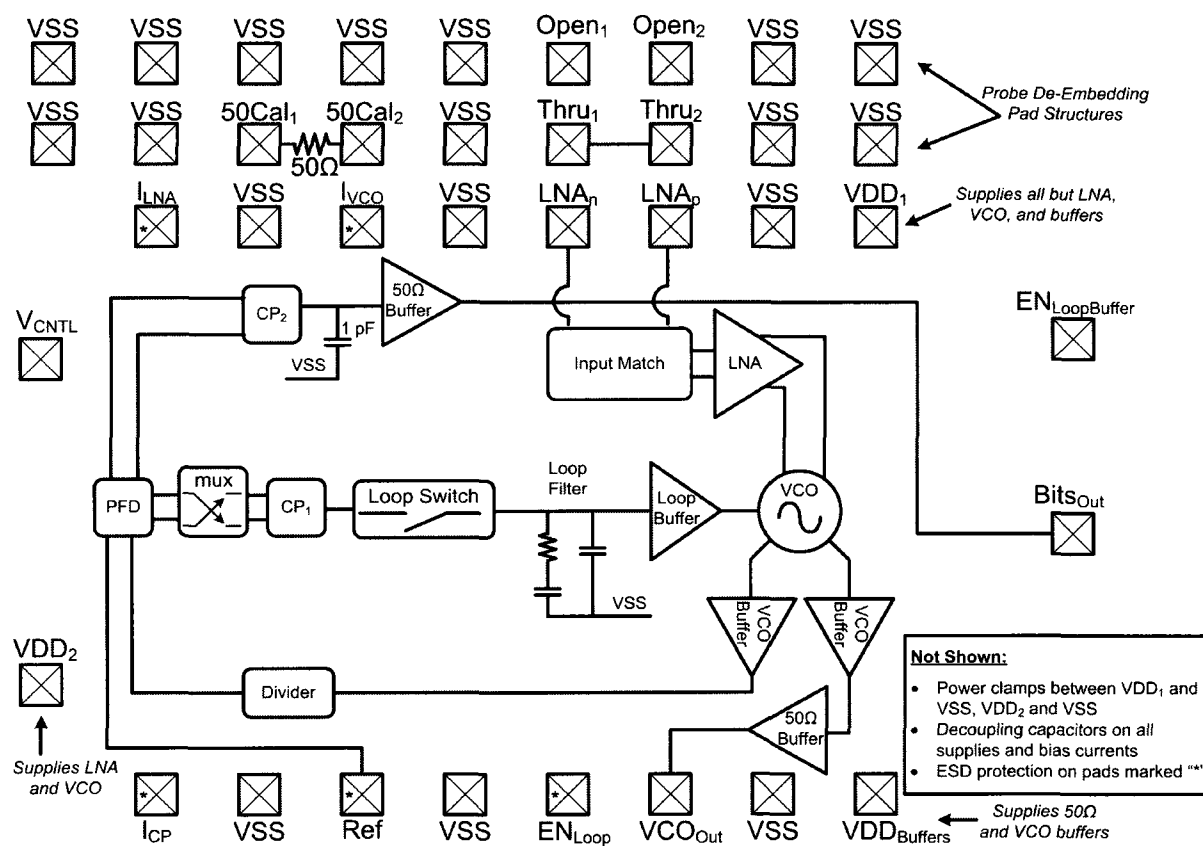


Figure 7.1: Lock-and-Roll RX Test Chip Block Diagram




Figure 7.2: Lock-and-Roll RX Die Microphotograph

were encountered during testing and the work-arounds implemented to maximize the testability of the die.

Finally, the measured output bitstreams of the lock-and-roll receiver test chip, and the limitations of the experimental die that they demonstrate, are presented and explained. While the results confirm that the receiver successfully demodulates the FM input signal, thereby validating the proposed receiver, the noise on the measured waveforms highlight improvements that could be made on subsequent test chips – which are outlined here.

Recognizing that key requirements of the applications targeted by the lock-and-roll receiver are that it be completely integrated and inexpensive to manufacture, thereby mandating ultra-low power consumption, the overall power consumption breakdown as measured with the lock-and-roll receiver test chip is reviewed and compared with other published results, highlighting the strengths and weaknesses of the overall topology.

7.1 High-Speed, Low-Power Divider Design with TSPC Input Stages

The divider used on both the lock-and-roll RX and TX test chips is of the same design, as implemented by the author’s research colleague Victor Karam [52]. When designing PLL and synthesizer circuits having high output frequencies on the order of a few gigahertz or more, the frequency divider circuit inevitably ranks high when considering the power consumption breakdown of the overall die due to the fact that its first few input stages are clocked at or near the high output frequency of the VCO. Historically, CMOS logic cells were unable to switch fast enough for use in the initial divider stages, but with the advancements made in modern IC technologies, standard CMOS cells can now be used, in many cases, with switching speeds of up to 2 or 3 GHz. Beyond these frequencies, current-mode logic (CML) cells are typically used [51] because of the faster switching times they provide, though at the expense of increased current consumption compared to CMOS logic, and requiring complementary (i.e., differential) inputs. An additional benefit of CML circuits, however, is that they typically operate with less input amplitude than CMOS cells, which reduces the

drive-strength requirements on the VCO or any added buffering circuits. Synthesizer circuits often require programmable divide ratios which in turn require configurable divider stages. Often the initial stages are not configurable, however, dividing down the high-frequency input signal with simple and efficient circuits, where subsequent stages of the divider are made programmable to provide the necessary divide ratio flexibility.

The divider in the lock-and-roll receiver is a fixed divide by 64 circuit made up of six cascaded divide by two circuits. The first three stages use a topology known as true single phase clocking (TSPC) which has been shown [71] to offer the performance of CML logic with the benefit of requiring only one clocking phase, and thereby reducing the current consumption of the divider. As the TSPC stages require more input signal amplitude than CML logic, each TSPC stage's input is buffered with a standard CMOS inverter. The last three stages of the design, where the clocking frequency is much slower, are implemented with simple CMOS logic whereby a common flip-flop with “Qb” output connected to “D” input accomplishes the divide by 2 task. Figure 7.3 shows the schematic for the divider topology in the lock-and-roll receiver.

An important factor to consider in the design of high-frequency PLL or synthesizer circuits, is that the divider circuit be capable of switching at the highest output frequency that the VCO is capable of generating. Even if the PLL is not expected to operate with a nominal $V_{\text{CNTL}} = V_{\text{DD}}$, startup conditions may very well result in loop conditions that maximize the VCO output frequency, and if the divider is incapable of switching under such conditions, the loop will be unable to tune the VCO output to the nominal frequency. Operating temperature and process corners will also influence the highest output frequency of the VCO.

The divider that was implemented on the lock-and-roll RX and TX test chips proved successful at dividing down input frequencies beyond 6.3 GHz during testing, as documented by the author's research colleague Victor Karam [52], which is well above the measured output frequency of the RX VCO at all V_{CNTL} settings between VSS and VDD (see Figure 4.9).

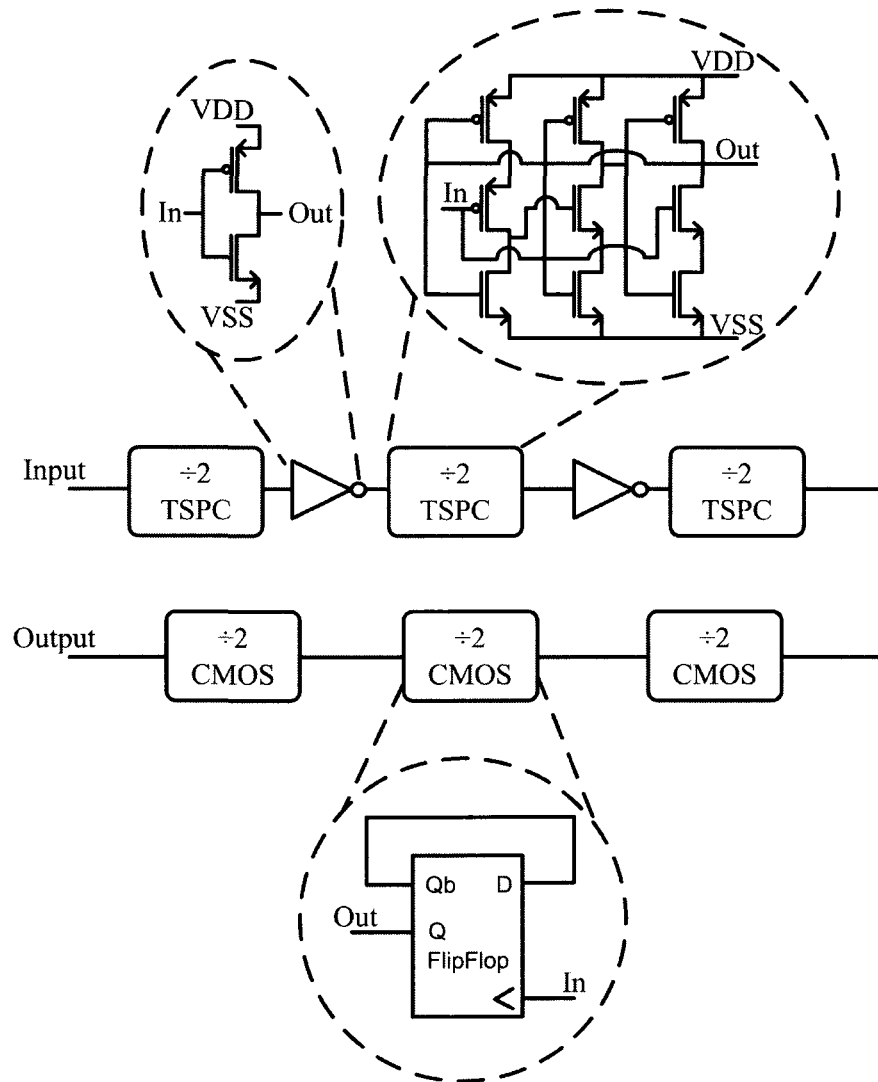


Figure 7.3: Lock-and-Roll RX Divider Schematic

7.2 Digital Up/Down Pulse Mux Design and Use

Recall from sections 6.5 and 6.6 that the addition of an inverting loop-filter buffer circuit between the integrated loop filter circuit and the accumulation mode varactor diodes in the VCO proved essential, in both simulation and measurement, to enabling successful open-loop operation. Without the loop-filter buffer enabled, the VCO drift rate when the loop is opened makes the demodulation of input data all but impossible at a 5 kb/s data rate. To allow for comparative measurements to be taken, both with and without the loop-filter buffer enabled, the addition of a mux circuit is necessary between the PFD outputs and the charge pump inputs such that the correct polarity of the loop can be maintained both with and without the inverting buffer enabled. The schematic for the mux circuit is shown in Figure 7.4, where transistor pairs M1 and M2, M3 and M4, M5 and M6, and M7 and M8 form CMOS transmission gates that accomplish the muxing task. When the “En” control signal (which also

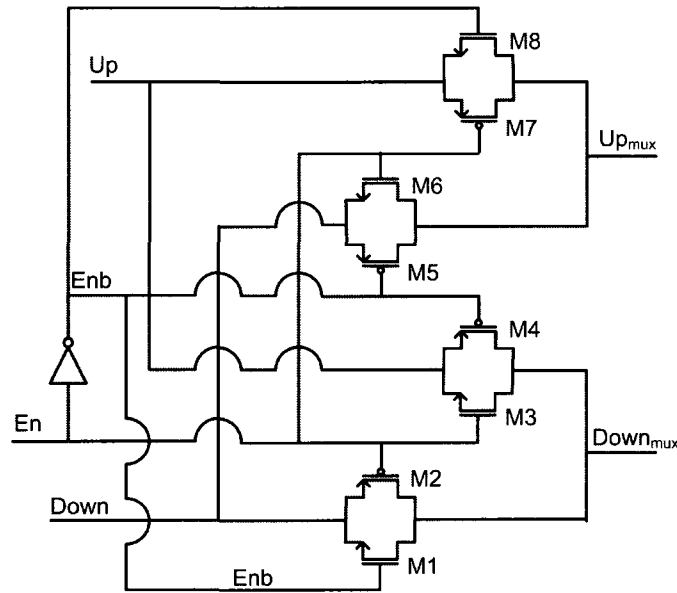


Figure 7.4: Lock-and-Roll RX Up/Down Mux Schematic

enables/disables the loop-filter buffer) is low, the loop-filter buffer is disabled and there is no need to reverse the connections between the PFD and CP_1 . The “Up”

input signal is transferred to the “Up_{mux}” output signal, and the “Down” input signal is transferred to the “Down_{mux}” output. Conversely, when the En signal is high, calling for the mux to switch the connections, the Up input signal is transferred to the Down_{mux} output, and the Down input signal is transferred to the Up_{mux} output.

As the output of the secondary charge pump, CP₂ that is only used in open-loop mode to demodulate the incoming data bitstream, is connected to a buffer circuit (as shown in Figure 7.1) that drives the output pad whereby the buffer itself is an inverting amplifier, the inputs to CP₂ are actually connected straight to the PFD outputs, Up to Up and Down to Down. The reader may recall that the block level diagrams for the receiver (Figure 3.7 and Figure 6.1) show that the connections to CP₂ are reversed, but in fact this is not the case. Figure 3.7 and Figure 6.1 are drawn in this way to communicate the general phase of the signal paths through the TX and RX communication system, whereby implementation details such as the inverting 50 Ω pad buffer circuit are left out of these diagrams in the interest of simplicity and clarity.

7.3 Buffer Circuits that Enable Testability

There are essentially two different buffer circuits implemented on the lock-and-roll RX (and TX) test die that enable the topology to be tested in the standard 50 Ω lab environment. All buffers are supplied off chip from a separate VDD pin (labeled VDD_{buffers} in Figure 7.1) such that the current consumption of these blocks, which are merely implemented for testing purposes, can be measured separately from the current consumption of the core circuit blocks during testing.

7.3.1 VCO Output Buffer Design

The VCO output buffer is a simple, single-ended circuit as shown in Figure 7.5, which buffers the VCO tank from the input capacitance of the divider while simultaneously providing the drive strength required to toggle the first TSPC divider stage. Maintaining a symmetrical load on the VCO, a second VCO buffer circuit is used to connect the complementary (to the output terminal that drives the single ended




Figure 7.5: Lock-and-Roll RX VCO Buffer Schematic

divider) output of the differential VCO to the buffer circuit that drives the $50\ \Omega$ spectrum analyzer used for monitoring the VCO output during testing (recall Figure 7.1). The circuit is a self-biased inverter, whereby the output is connected to the input via large resistor R1 that biases the gates of M1 and M2 at the trip point of the inverter. The topology consumes appreciable current when enabled, even with no input signal, but switches readily with minimal input signal swing thereby complementing the small output swing of the VCO, which (recall section 4.2.5) was reduced to improve the injection-locking bandwidth. Capacitor C1 serves to block the mid-rail DC bias voltage on the gates of M1 and M2 from that of the VCO tank circuit, while AC coupling the VCO output to the buffer. C1 is relatively small (roughly 150 fF) and M1 and M2 are also small (but sized for symmetrical output slew rates), such that the load imposed on the VCO tank circuit is negligible.

7.3.2 Buffering to Interface with the $50\ \Omega$ Domain

The buffer circuit that drives the $50\ \Omega$ load imposed by the spectrum analyzer is shown in Figure 7.6. The buffer is comprised of three common-source amplifier stages cascaded together in series (with resistive degeneration used in the first two stages), whereby the NMOS devices and resistors were sized to increase the drive strength by a factor of 3, approximately, relative to the preceding stage. The output stage is quite large and is able to drive a $25\ \Omega$ load nearly rail-to-rail at 5.2 GHz. The output




Figure 7.6: Lock-and-Roll RX 50 Ω Output Buffer Schematic

of the buffer is AC coupled to the output pad by large capacitor C1, and an on-chip 50 Ω resistor provides a rough impedance match to the 50 Ω spectrum analyzer.

7.3.3 Output Bitstream Buffer Design

The output signal of the secondary charge pump CP₂ (which integrates Up and Down currents across an on-chip 1 pF MIM capacitor), is connected to the input of a buffer circuit that drives the “Bits_{Out}” pad as shown in Figure 7.1. The circuit that was used for this purpose is identical to that shown in Figure 7.6 but with C1 removed given the low-frequency output signal which is essentially rail-to-rail. The Bits_{Out} pad was measured using a high-impedance probe connected to a spectrum analyzer, thus loading the buffer with approximately 50 Ω given the on-chip load.

7.4 Measurement Methodology

The original plan for testing all individual blocks on the test chip and the overall lock-and-roll receiver’s performance was to probe the die using one of the RF probe stations available at Carleton University. Given this intention, the pad arrangement

that is shown in Figure 7.1 and Figure 7.2 enables all of the necessary connections given the three sided, horseshoe shaped stage of the probe station. Landing two 8-pin RF probes from opposite sides of the die (top and bottom given the orientations shown in Figure 7.1 and Figure 7.2), and three (or four) single DC probes, all of the key signals can be accessed. The 8-pin probes have a P-G-S-G-S-S-G-P arrangement, whereby “P” represents a power connection (for supplying a DC current or bias voltage through a Bayonet Neill-Concelman (BNC) type connector), “G” represents a ground connection, and “S” represents a signal connection (intended for RF connections using a Sub-Miniature type ‘A’ (SMA) connector), where all grounds (including outer connections of all SMAs and BNCs) are shorted together to a common node. Note from Figure 7.1 that the two functional (not for de-embedding purposes) rows of 8 pads are mirror images of each other (in terms of P-G-S-G-S-S-G-P orientation) as a pair of 8-pin RF probes is comprised of two probes which are indeed mirror images of each other.

Three separate VDD domains exist on the test chip, whereby the VCO and the LNA are powered from VDD_1 , all buffers are powered from VDD_{buffers} , and all remaining blocks are powered from VDD_2 . Separating the supply of the sensitive analog blocks (namely the VCO and the LNA) from that of the noisier digital blocks (such as the divider) is generally good design practice, and allows for some added flexibility. For instance, the separate supplies allow the charge pump output current to be characterized with the VCO disabled while the PFD constantly sends “Pump-Up” signals (as highlighted in section 6.4.3), and allows for the current consumption of the buffer circuits (which are included simply for the purpose of testing) to be considered separately from that of the core circuits. A typical synthesizer or PLL application, whereby an important requirement of the design is to minimize the clock feed-through spurs at the VCO output, also benefits from having separate supplies for digital and analog sections of the chip. At the very least, “star connecting” the supplies, whereby the domains are kept separate across the chip and shorted together right at the VDD pad itself, reduces noise coupling between sections of the chip. On the lock-and-roll receiver test chip the LNA and VCO supplies are star connected at the VDD_1 pad, and the supplies to the other blocks are similarly star connected at the VDD_2 pad. The same is true of the VSS routes to each block, whereby all VSS pads are connected together around the perimeter of the chip, with the ground ring

carefully broken to form a large “C” shape, thereby reducing the chance of a complete loop acting like an on-chip antenna and coupling in outside noise signals.

In order to characterize the input impedance of the differential LNA circuit (with integrated input match), as summarized in section 5.2.8, the network analyzer must first be calibrated to normalize the measurements. Using a set of precision open, short, through, and $50\ \Omega$ loads to calibrate the instrument to the end of the $50\ \Omega$ SMA cables that are connected to it is useful (indeed it is the first step that was performed in testing) but not sufficient to enable accurate characterization of the LNA input. A set of on-chip calibration structures should be used to de-embed the probe and the pads on the die (including the parasitics to the neighbouring pads) from the 5.2 GHz measurement.

7.4.1 Probe De-embedding Structures, Arrangement and Use

Note from Figure 7.1 that the differential LNA input is connected to the adjacent S-S (signal-signal) pads, located three and four pads in from the top right hand side of the die (ignoring the two rows of de-embedding pads momentarily). The top two rows of pads, which are actually 9 pads wide, form a rather unique set of de-embedding structures whereby all four of the required calibration standards can be measured with only two rows of pads – thereby making efficient use of the test chip die area. Figure 7.7 highlights the procedure for calibrating the network analyzer using the two rows of calibration pads.

Lowering the 8-pin probe initially in the red position (as outlined in Figure 7.7) results in an “open” calibration, whereby the pads touching the S-S probe tips are floating and not connected to anything on-chip. Raising the probe and lowering it next in the green position provides the “through” calibration, whereby the two adjacent pads touching the S-S probe tips are well shorted together with a wide, thick top metal route, and connected to nothing else on the die. Raising the probe yet again and rotating the die a complete 180° allows the probe to next be lowered in the blue position, whereby the S-S probe tips are connected to grounded pads which are well connected to the ground ring on the chip, providing the necessary “ground” (or “short”) calibration. Finally, by raising the probe one last time and lowering it in the orange position, an on-chip $50\ \Omega$ load is connected across the adjacent S-S probe

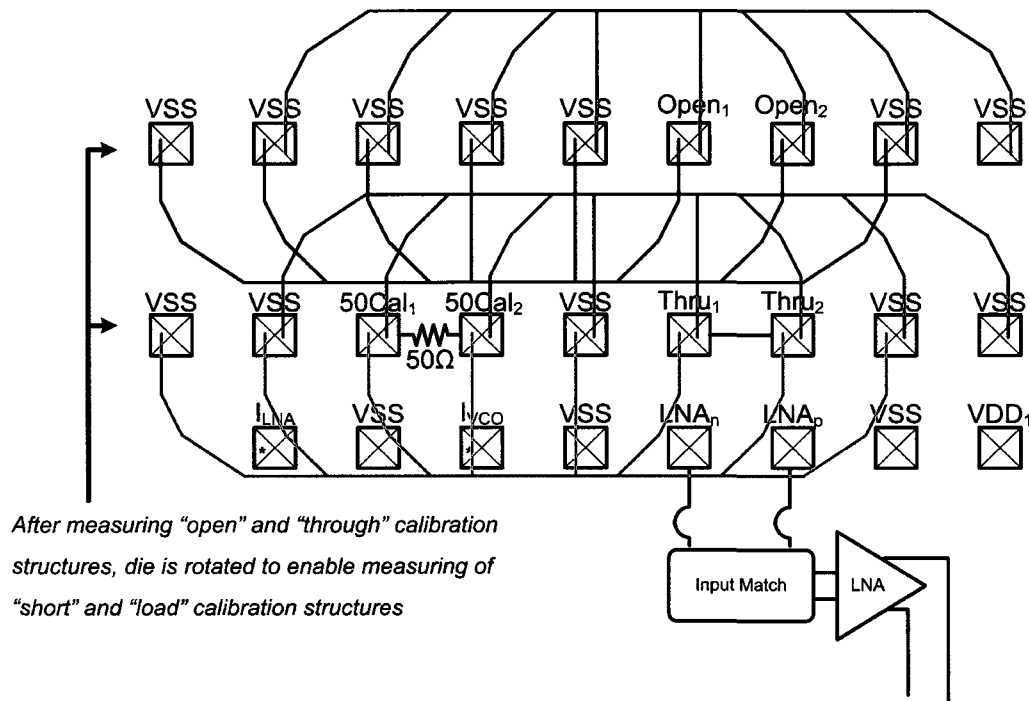


Figure 7.7: Receiver Test Chip Probe De-Embedding Options

tips, yielding the final “load” calibration. The 50 Ω calibration load is implemented on the test chip using the back-end of line (BEOL) top metal type of resistor which offers the tightest integrated resistor tolerance available, 6% at room temperature, with the design kit that was used. While the 50 Ω calibration will not be perfect, using it will likely introduce only a small error to the measured results while allowing for capacitance and parasitics associated with the probe and the pads to be de-embedded for the most part, yielding a far more accurate measurement at 5.2 GHz than what would be possible with the use of the precision SMA calibration standards alone – which only calibrate out the SMA cables themselves from the measurement.

Lastly, note that the pad connections were arranged such that landing a single 8-pin probe to connect to the input of the LNA also provides a means of supplying VDD₁ (and VSS) and the two bias current references needed by the LNA and the VCO. The intention of this arrangement is that only one probe needs to be landed in order to facilitate testing of the LNA and VCO in combination, enabling the match to be characterized with the VCO load connected to the LNA output.

While the method of probing the test chip for testing purposes proved successful (and essential) for characterizing the input impedance of the LNA with integrated matching circuit (recall section 5.2.8), the approach quickly became cumbersome for achieving reliable connections to all the pads required to test the full functionality of the receiver day after day using the shared resource that is the electronics lab at Carleton University. Not only was it difficult to get a solid connection simultaneously with all 16 tips of the two 8-pin probes, landing the DC probes along side the RF probes on the same stage proved difficult to do without having the probes interfere with each other. Achieving good connections to all pads typically took upwards of two hours, and time on the probe station had to be shared with other students and other projects. Additionally, the RF probes are very delicate and are often damaged by first-time users, resulting in expensive repairs that take weeks to be performed with long round trip shipping times, etc., resulting in testing delays which jeopardized the authors plans to meet paper submission deadlines for numerous IEEE conferences. Given the many drawbacks associated with probing a test chip of this complexity, the author set out to find an alternative whereby the chip could be glued down to a printed circuit board (PCB) and bonded out for reliable and repeatable testing purposes.

7.4.2 Enabling Chip-on-Board System Level Testing

Luckily the printed circuit board designed and used by the author for a previous 5.2 GHz research project [9] at Carleton University offered nearly enough flexibility to be reused in its original form. Modifying a few of the spare boards left over from the previous project for which they were intended allowed for the boards to be recycled for the purposes of testing the lock-and-roll RX and TX, while preventing the cost and time penalty to the author and his research colleagues that would otherwise have resulted from having to design a new PCB and have it manufactured. Figure 7.8 shows the modified PCB, whereby the RX die was glued to the ground paddle with conductive epoxy, and the author used an old wedge-bond to wedge-bond bonding machine at Carleton University to bond out all of the pads required for testing.

With respect to the orientation shown in Figure 7.2, the die was rotated 90° counterclockwise and mounted on the PCB. The two SMA connections (with $50\ \Omega$

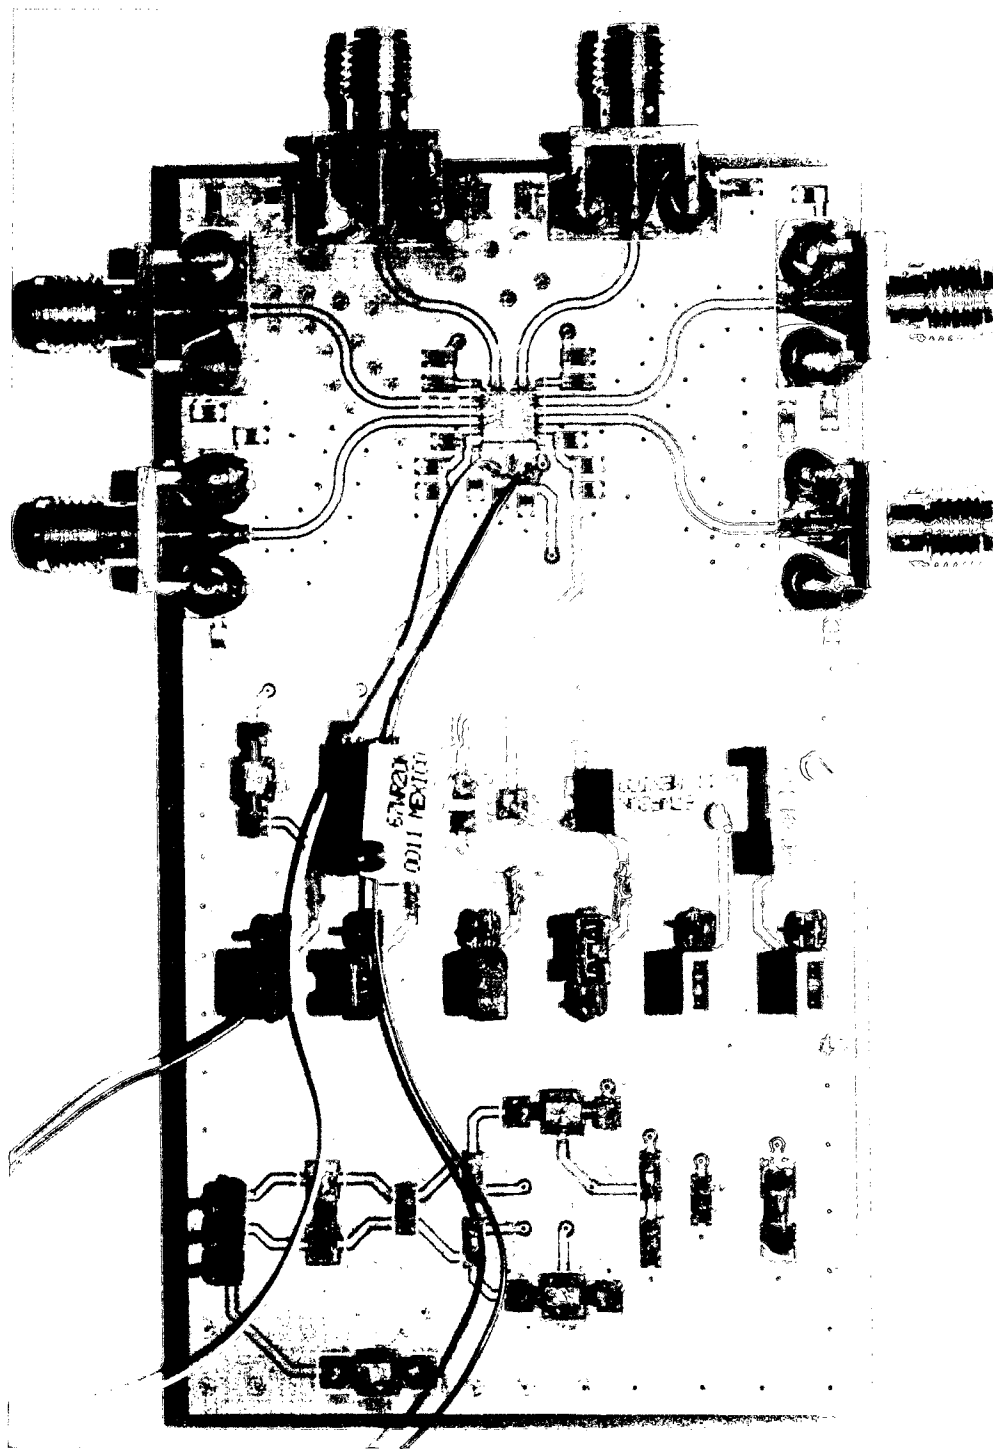


Figure 7.8: Receiver Test Chip Bonded to PCB

routes at 5.2 GHz) on the left side of the board are connected to the differential LNA input pads, the two SMAs on the top of the PCB are connected to $EN_{\text{LoopBuffer}}$ and Bits_{Out} , and the two SMAs on the right of the PCB are connected to the VCO output and the reference input signal for the PLL. The bottom third of the PCB is dedicated to supply filtering at low frequencies where the positive external supply, its negative return, and chassis ground are connected to the terminals of pin header block J1. The middle third of the PCB has six separate pin header blocks and three potentiometers that allow for independent bias control to the core blocks on the test chip. The two leftmost pin headers in Figure 7.8 are used to enable VDD_1 and VDD_{Buffers} independently, while the third pin header from the right allows for the EN_{Loop} signal to be connected to the positive external supply or board ground depending on the placement of the jumper on that header. The three potentiometers allow for precise trimming of the bias current references for the charge pumps, the LNA and the VCO. Below the die paddle itself, two independent islands of top metal on the PCB were isolated and cleaned of solder resist with a sharp knife. The V_{CNTL} and VDD_2 pads were bonded out to these locations where the white and blue wires in Figure 7.8 were soldered to yield access to the final pads that require a connection for RX level testing. The blue wire is connected to VDD_2 and was often shorted to the positive supply along with VDD_1 at the leftmost pin header, while the white wire is connected to V_{CNTL} and was left floating in most cases, except when conducting output current characterization of the charge pump (recall section 6.4.3) and tuning range characterization of the VCO (recall section 4.2.7).

Details of the PCB design, the low-frequency supply decoupling circuit, and the bias control are covered in much detail in the author's M.A.Sc. thesis [9]. Note that the off-chip decoupling capacitor footprints located all around the perimeter of the die paddle are not populated as there is much high-frequency supply decoupling on the lock-and-roll receiver die itself which strives to minimize the need for external components.

With the lock-and-roll receiver test chip firmly mounted on a PCB, having reliable connections to all the required pads in place with excellent bias control options, performing overall performance testing on the complete receiver chain is greatly simplified.

7.5 Measured Receiver Output

The measured output bitstream from the lock-and-roll receiver is shown in section 6.6, where the noise on the signal is not explained, but where the number of consecutive bits that can be successfully demodulated in a single closed-loop, open-loop, closed-loop cycle of the receiver is used to estimate the leakage from the integrated loop filter. Figure 7.9 shows the output of the receiver when the transmitted signal has $\Delta f = 500$ kHz, and a data rate of 1 kb/s for the purpose of showing clear demodulation results. The EN_{Loop} signal is being driven by a square wave signal having a 50%




Figure 7.9: Measured Receiver Output Signal with 1 kb/s Data Rate

duty cycle and a frequency of roughly 20 Hz such that the loop is closed and the VCO re-centered before the receiver fails due to VCO drift given the amplitude of the input signal. Figure 7.9 shows only the open-loop mode of operation when the EN_{Loop} signal is low. The yellow signal on the oscilloscope is the output signal from

the die measured using a high-impedance scope probe via the Bits_{Out} pad, while the blue signal is the pseudo-random bit sequence used to modulate the RF carrier of the signal generator that is driving the LNA input. The result clearly shows that the output of the receiver is a faithful reproduction of the transmitter data, though there is rail-to-rail (nominal VDD for the test chip is 1.2 V) noise on the output that appears at regular intervals. The noise on the output results from the 1 pF integrating capacitor that was used on the output of the secondary charge pump CP₂ being too small given the 100 μ A charge pump current.

7.5.1 Integrating Capacitor Size and Noise on the Final Output

Figure 7.10 shows a zoomed in view of the rail-to-rail noise on the RX output pad, whereby the frequency of the noise is measured to be roughly 7.8 kHz. With the




Figure 7.10: Measured Output Signal Noise Period

output of the VCO switching between 5.2005 GHz and 5.1995 GHz, having been injection locked to the modulated input signal with $\Delta f = 500$ kHz, the output of the

divider block switches between 81.2578125 MHz and 81.2421875 MHz respectively, i.e., 64 times slower than the output of the VCO. Recall that the reference signal for the loop has a frequency of 81.25 MHz, and thus the beat frequency between the reference signal and the output of the divider is always 7.8125 kHz regardless of whether the bit being demodulated is a 1 or a 0. The result agrees perfectly with the beat period calculated in section 3.6 using equation (3.9) for the given modulation. Given the slight frequency difference between the “Ref” and “Div” inputs to the PFD under these conditions, a cycle slip will occur every 128 μ s.

Consider, for sake of discussion, the behaviour of the loop when the data being transmitted is a 1 as shown on the right side of Figure 7.10. The divider’s output signal is 7.8125 kHz faster than that of the reference, and the Down output signal from the PFD will be high much more frequently than that of the Up output. In fact the Up output will only pulse high very briefly (i.e., the duty cycle of Up is very low) when the flip-flops in the PFD are being reset (recall the 3-state PFD behaviour discussed in section 6.3), while the duty cycle of the Down signal will be comparatively large on average, progressively increasing as time passes since the last cycle slip event until the duty cycle is nearly 100% just prior to the next cycle slip event which returns the duty cycle to nearly 0%, and the cycle repeats. Figure 7.11 highlights the progression visually, where a cycle slip occurs where indicated, and the frequency difference between Ref and Div have been exaggerated for clarity. The




Figure 7.11: The Cycle Slip Phenomenon

result is that the voltage will be drawn down across the 1 pF integrating capacitor at the output of CP₂ (recall Figure 7.1) more and more as the cycle slip approaches, and then very little immediately following the cycle slip until the duty cycle improves. Recall that as the 50 Ω output buffer is an inverting amplifier, drawing down the voltage across the integrating capacitor translates to driving the Bits_{Out} pad high.

The logical next question is then, “Why is the noise indicating that upon cycle slipping the output voltage is actually reversed momentarily, if in fact the duty cycle of the Up signal isn’t any larger following a cycle slip?”. In fact the answer to this question is likely best explained by referring back to Figure 6.10 which shows that when the charge pump output voltage is near either VDD or VSS, there is a large mismatch between the Up and Down output currents. Again considering the right section of Figure 7.10, the Bits_{Out} pad is high the majority of the time, indicating that the output voltage of CP₂ is low due to the Down output from the PFD being high, on average, much more so than the Up output. Whenever the Down signal is reset, for a small window of time both the Up and Down outputs are high at the same time. Figure 6.10 shows that when the output voltage is mid-rail, the Up and Down currents are well balanced and the result is that little current will be sourced to or drawn from the integrating capacitor. When the voltage is near VSS however, the Up current is much larger than the down current (as the output voltage across M10 and M11 in Figure 6.7 is $< 2 \cdot V_{DS_{sat}}$), and the result is that for that small window of time I_{CP} will effectively be sourced to the capacitor. When the reset event occurs far away in time from a cycle slip the effect is negligible as the duty cycle of the Down signal is large and quickly removes the charge on the capacitor and the voltage does not rise above the trip point of the 50 Ω output buffer. After a cycle slip however, the duty cycle of the Down signal has been reset to near 0% and the effect dominates the net charge on the capacitor such that the trip point of the 50 Ω output buffer is reached, toggling the final output to the Bits_{Out} pad. As the voltage across the integrating capacitor increases however, the imbalance between Down and Up currents decreases, and recall that the duty cycle of the Down signal is also increasing with time. Eventually the Down signal dominates once again and the voltage on the capacitor is lowered below the trip point of the 50 Ω output buffer, toggling the final output, and maintaining this state until the next cycle slip event.

As can be expected (and is reinforced by the results shown in Figure 7.9 and Figure 7.10), the phenomenon also occurs when the divider's output frequency is slower than that of the reference, with the Up and Down signals playing opposite roles from the scenario depicted in Figure 7.11.

If the size of the 1 pF integrating capacitor was increased, the noise would likely be eliminated as making the capacitor larger would make it much more difficult for the Up current to affect the voltage on the capacitor during the small window of time where it dominates over the Down current, as the larger capacitor would sink more charge before the voltage at the input of the 50 Ω output buffer reached the threshold necessary to toggle the output.

While the noise on the final output signal of the receiver can be attributed to a minor flaw in the design of the test chip (the integrating capacitor should have been made larger), the measured results have been fruitful in that they demonstrate successful demodulation of the transmitted data, and allude to the upper bound on the number of bits that can be received due to VCO drift as analyzed in section 6.6. Additionally, the results are suggestive of the maximum data rate with which the receiver is compatible.

7.5.2 Data Rate Limitations Revisited

Recall the discussion in section 3.6 which outlines the inevitable tradeoff between the data rate, the communication range, and the power consumption of the topology. The analysis highlights that with a reference frequency of 81.25 MHz and $\Delta f = 500$ kHz, the maximum data rate that can be used is roughly 5 kb/s based on the fact that the worst case phase difference at the time of a bit change could result in the receiver taking nearly 50% of the width of the bit to recognize the transition and to follow suit. Recall that increasing Δf would reduce this delay, thereby increasing the maximum data rate, but the injection-locking bandwidth would need to be increased in order to accommodate a larger Δf and the communication range would suffer in turn. In fact the results depicted in Figure 7.12 roughly confirm the result of a worst case phase difference at the time of a bit change. In Figure 7.12 the data rate is only 1 kb/s and $\Delta f = 500$ kHz, and the oscilloscope shows a zoomed in view of the worst case delay between input and output bit transitions as could best be discriminated by




Figure 7.12: Measured Transition Delay with $\Delta f = 500$ kHz

the author. Using the vertical cursors to measure the delay shows that the receiver takes roughly $100 \mu\text{s}$ to recognize the change and to toggle its output accordingly. As the width of a bit when the data rate is 5 kb/s is $200 \mu\text{s}$, the result shows a worst case delay that aligns well with the theoretical limit that was calculated in section 3.6.

The transition delay, which depends on the phase difference between Ref and Div at the time of a bit change, can also be seen in simulation, though conducting long transient simulations with the complete RX circuit is painfully slow. Figure 7.13 shows the results of one such transient simulation, whereby the data rate of the modulated input signal was 65 kb/s (which would no doubt cause the RX to miss some bits altogether at some point) just to force the demodulation of a few bits during the $100 \mu\text{s}$ simulation. The results show that the loop acquires lock in roughly $7 \mu\text{s}$, showing good damping (little overshoot and ringing on V_{CNTL}) and the traditional evidence of cycle slips (leading to charge sharing between C1 and C2 of the loop filter and the associated shark's teeth profile of V_{CNTL} during the acquisition stage [51]). At the time indicated by M2 the loop is opened and bit changes occur at M1 and M3.

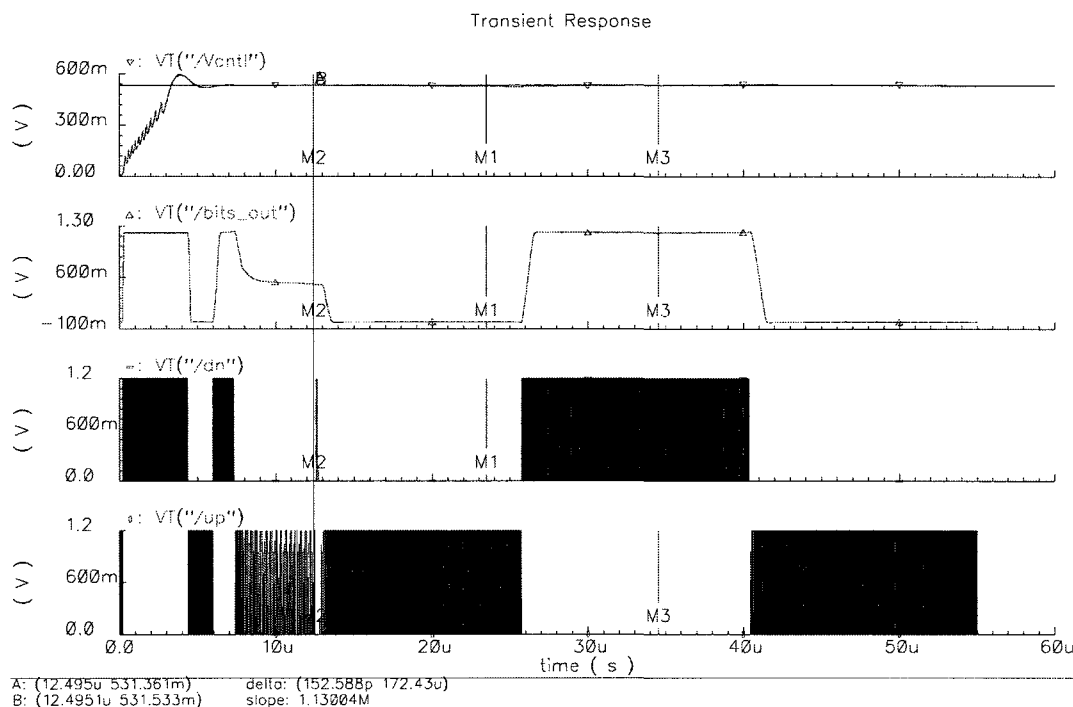


Figure 7.13: Simulated RX Demodulation Showing Delay

In fact CP_2 was enabled for the duration of the simulation merely by accident. The results show that the response of the RX to the first bit change is delayed by about $2 \mu s$, and the delay after the second bit change is about $7 \mu s$.

An interesting observation is that neither the Up or Down outputs from the PFD appear to toggle with every reset event which is largely expected to contribute to the output noise that is visible on the measured RX output. Nor does the simulated output show the rail-to-rail output noise characteristic of the measured output. Likely the schematic for the simulation, which did not include layout parasitics due to the prohibitively long simulation time that would result, underestimates the reset delay of the PFD. Additionally, note that the Up and Down current mismatch of the simulated charge pump (see Figure 6.8) is better than that of the measured charge pump response (see Figure 6.10) which is also believed to contribute to the noise on the measured output (recall section 7.5.1).

Note also that due to the use of the mux and the loop-filter buffer circuit in the simulation, the loop's response to Up and Down pulses from the PFD is opposite

to convention. The “A” and “B” markers in the simulation are merely noting the ripple on the control voltage that was introduced at the moment the loop switch is opened, which is minimized by the presence of the dummy transmission gate in the loop switch (recall section 6.2.1).

While the measured output of the RX test chip attests to its suitability for demodulating low data rate signals for short-range communication applications, the suitability of the approach for use in RFID and medical sensor type applications clearly hinges on the overall power consumption of the topology and how it compares to state-of-the-art receivers with similar abilities (and limitations) in terms of integration, cost, and data transmission.

7.6 Measured Receiver Power Consumption Breakdown

Table 7.1 shows a breakdown of the measured power consumption of the lock-and-roll receiver. The overall power consumption is approximately 5.7 mW regardless of

Table 7.1: Lock-and-Roll Receiver Measured Power Breakdown

Receiver Element	Power Consumption
LNA	1.5 mW
VCO	1.2 mW
Divider	2.2 mW
Remaining RX Elements	0.8 mW
Total	5.7 mW
Duty-Cycled Total	285 μW

whether the loop is operating in the open-loop or closed-loop modes, and the duty-cycled power consumption assumes that the RX is communicating with a TX that is sending 250 bits of data once per second at a data rate of 5 kb/s (see section 3.4.2). Recall that the measured drift (see section 6.6 and Figure 6.12) suggests that the RX can only demodulate approximately 110 to 120 bits of data at 5 kb/s before needing to have the loop closed and the VCO refreshed, but recognizing that both the simulated

(and measured) startup time is on the order of $8\ \mu\text{s}$ (see Figure 7.13), closing the loop to refresh the VCO and breaking the bitstream into two 125 bit lengths (or three 85 bit lengths) does not really affect the overall power consumption of the RX, given that the startup (or settling) time of the closed loop is so small relative to the period of a bit at 5 kb/s. The measured power consumption is about 4% higher than what is predicted in section 3.4.2, but still falls well within the predicted sourcing capabilities of a $3\ \text{mm}^2$ ultracapacitor fabricated on a 2 mm by 2 mm CMOS die as discussed in section 2.5.2, whereby published results suggest such an integrated supply should be capable of supporting a circuit drawing 5 mA of current for sixty 50 ms bursts between chargings.

While the power consumption of the design appears to be sufficiently low for use in the targeted applications, many other factors such as the level of integration, communication frequency, cost, communication range, and the data rate must be considered when evaluating the lock-and-roll RX in comparison with other published topologies.

7.7 Comparing the Lock-and-Roll RX to State-of-the-Art Alternatives

Making a truly fair comparison between the lock-and-roll RX and the published alternatives is a difficult task. The difficulty arises, in large part, due the fact that short-range communications are generally less regulated than other forms of communication (in terms of power levels, modulation scheme, bandwidth, frequency, data rate, etc.) and there are few standards. Taking the argument to extremes, where circuit designers addressing telecommunications applications have to meet strict guidelines to ensure inter-operability and compliance in terms of spectral masks, emissions, etc., the lack of regulations steering the designer of circuits for ultra-low power, short-range communication systems might have him or her feeling like they are operating in the “wild west”! The priorities of the application for which any communication system is designed must be clearly outlined beforehand. The discussion ultimately comes full circle and highlights the importance in understanding the tradeoffs depicted by

Figure 2.1, when it comes to designing short-range, low-power, highly integrated and inexpensive communication systems.

In general, circuits that communicate at low frequencies will consume less power, perhaps enabling the use of on-chip power sources as a result, but likely eliminating the potential for an integrated antenna. As the communication frequency is increased, novel (yet simple) architectures are required to maintain low-power operation, yet the possibility of an on-chip antenna approaches reality. Higher performance technologies such as GaAs, silicon germanium (SiGe), silicon-on-insulator (SoI) and exotic technologies such as silicon-on-sapphire or silicon-on-diamond substrates may help to alleviate some tradeoffs, but using them comes at the expense of increased die cost. The overall cost of a solution is of principal concern when designing for RFID tag (or similar) applications that often require cheap and disposable solutions. Yet in some cases, as for many medical applications, the cost of a solution may be much less important than the overall size of the topology, or the materials that make up the solution when it comes to the presence (or lack) of toxic or hazardous chemicals and elements.

There is no figure of merit that can be used to normalize all published low-power, short-range designs and offer a fair comparison under all circumstances. That claim noted, to propose and implement a unique alternative like the lock-and-roll receiver and to not compare it with the performance of other topologies would seem short-sighted. As such, comparisons are drawn in the following paragraphs, yet the author notes that he is unaware of any published result which has been designed to operate with an integrated antenna and power source, in a bulk CMOS process, for the purpose of minimizing the physical size and the cost of the solution.

The receivers outlined in [72] and [73] both operate with communication frequencies of 433 MHz and do so with low power consumption. In [72] the receiver consumes 1.4 mW and is capable of communicating at 20 kb/s. Unfortunately [72] merely proposes the concept of the receiver and the performance is but speculative as a result, though the authors of [72] believe all blocks could be implemented in 0.5 μm bulk CMOS, requiring only an off-chip matching circuit and an off-chip antenna given the low communication frequency. The measured performance suggested in [73] is similar, whereby the “WiseNET” receiver mentioned consumes 1.8 mW and communicates at 24 kb/s, though the topology requires an expensive off-chip SAW

filter, and an off-chip antenna. Noteworthy is that the authors of [73] claim that at 1.8 mW the receiver consumes 30 times less power than comparable solutions, where the topology makes use of an injection-locked Colpitts oscillator divider, and like the lock-and-roll receiver, the chosen modulation scheme was binary FSK. Thus [72] and [73] consume less power than the lock-and-roll receiver and are capable of communicating at a higher data rate if required by the application, but they are undoubtedly more expensive to implement and while they could be powered using an on-chip energy source they are not compatible with an on-chip antenna. As such, the physical size of the solution would be larger than that of the lock-and-roll transceiver. Neither [72] nor [73] mention the communication range of the topologies.

The receivers in [74], [75], [76], and [77] all operate, along with many other interferers, near 900 MHz where an off-chip antenna is required in all cases. The design in [74] communicates at 40 kb/s and is implemented in 0.18 μm bulk CMOS, but it consumes 29 mW and requires an off-chip transmit/receive switch, a filter, an antenna and a crystal to function. The communication range is not mentioned. The topology in [75], implemented in 0.25 μm CMOS, claims FSK communication at a range of 16 m and a data rate of 20 kb/s while consuming only 1.3 mW, though an off-chip antenna, inductor, crystal, and battery are required with a claimed added cost of \$1. The design in [76] is similar to [74] in that it communicates at 45 kb/s and requires an off-chip filter, antenna, crystal and battery though with a claimed power consumption of 2.7 mW the topology is presumably compatible with on-chip power sources. The topology is said to target applications that require a 10 m communication range. The authors of [77] claim a design that communicates at 1 Mb/s using OOK modulation, while consuming 2.6 mW and requiring only a SAW filter and an antenna in terms of off-chip components. While all of the designs boast of higher data rates than the lock-and-roll receiver while being implemented in similarly inexpensive bulk CMOS, their frequency of communication prevents the use of an integrated antenna and not all of them could be powered using an on-chip ultracapacitor without substantially growing their die area.

The topologies defined in [78], [79], and [80] all communicate in (or near) the noisy 2.4 GHz band. The design in [78] uses FSK modulation and targets a 10 m communication range using a 40 kb/s data rate. The receiver consumes 7.5 mW and is implemented in a 0.18 μm CMOS process on the same die as a transmitter where both

are connected to an off-chip antenna using an integrated transmit/receive switch. The design achieves a high level of integration with reasonably low power consumption, though an off-chip antenna (and presumably nothing more) is required. The cost of the solution is presumably low. The solution outlined in [79] boasts an impressively low power consumption of only $400\ \mu\text{W}$ for the receiver, though it operates with a 10% duty cycle in order to achieve this result. The circuits communicate at 1.9 GHz and the authors claim that an average power consumption of $< 1\ \text{mW}$ is required in order to enable the use of energy-scavenging techniques for powering the topology. While implemented in a standard CMOS technology, the system depends heavily on the use of two off-chip BAW resonators and a bondwire inductor and the data rate and communication range of the transceiver are not mentioned. In [80] the designers made extreme sacrifices for the goal of achieving low power consumption. Their design is implemented in $0.13\ \mu\text{m}$ CMOS and they operate it from a mere 400 mV supply. The limited headroom they allotted themselves results in a design where no two transistors can be stacked one on top of the other and one is left to wonder how robust the circuits are to process mismatch and tolerance. Nevertheless, they claim a receiver that communicates at 400 kb/s, using binary FSK modulation, while consuming $750\ \mu\text{W}$ (a duty cycle is not mentioned). The complete transceiver is implemented on a single die including an integrated output match (though an off-chip balun might be needed) with the exception of the antenna, and the authors believe the topology can be powered by a solar cell, presumably charging an on-chip capacitor. There is no mention in [80] of the communication range of the system. These three examples, while still operating at communication frequencies that make integration of the antenna difficult, demonstrate that careful and clever design can cheat the tradeoff between operating frequency and power consumption. While not quite facilitating the potential for absolute integration, the designs in [78] and [80] are presumably cheap to manufacture.

A frequently used metric for comparison when it comes to analyzing low-power, short-range communication topologies is the energy consumed to receive (or transmit), one bit of information. While this metric clearly fails to put a value on the level of integration achieved, the physical size, or the manufacturing costs associated with the designs being compared, for completeness the author chose to borrow the comparison published in [80] and to add the lock-and-roll receiver into the analysis. Figure 7.14

shows the result, whereby many of the topologies just discussed are compared (along with [81], [82], [83], [84], [85], [86], [87], [22], and [23]) strictly in terms of their power consumption and the maximum data rate with which they are compatible.




Figure 7.14: Receiver Comparison Strictly Considering Energy/Bit

While the lock-and-roll receiver may not stand out as a clear winner in this analysis, recognize that it is operating at more than twice the communication frequency of all the results plotted, enabling the use of an on-chip antenna while still achieving a power consumption that is compatible with the output of an integrated power source.

Finally, there is at least one particular published design that deserves to be recognized as rivaling the lock-and-roll receiver in terms of the potential level of integration that was achieved. In [88], the authors outline a transceiver for use in wireless sensor networks which uses OOK modulation to communicate at 10 GHz with a data rate of 10 kb/s, and having a communication range of 30 cm. The receiver, which consists of no more than an amplifier and an energy detector, is said to consume 400 μ W. The power consumption is clearly compatible with an on-chip

power source, and an integrated antenna (though not attempted by the authors of [88]) should be feasible at 10 GHz. However, as the system can only communicate at a range of 30 cm using large patch antennas, the use of an integrated antenna (which would have much lower antenna gain) would likely reduce the communication range such that the transceiver would have few uses. Additionally, the transceiver in [88] is manufactured in a GaAs process which will almost certainly result in an overall cost that exceeds that of the lock-and-roll receiver, which by comparison, is implemented in low-cost bulk CMOS.

7.8 Lock-and-Roll Receiver Test Chip Summary

The lock-and-roll receiver test chip serves to demonstrate the unique RX topology that is proposed in Chapter 3, while putting into practice the IC design techniques that are outlined in Chapters 4 through 6 for achieving ultra-low power, highly integrated communication devices. The complete PLL, LNA with integrated input match, and all the necessary auxiliary circuits that facilitate testing of the topology are implemented on the experimental die.

The divider in the PLL that accomplishes the fixed divide-by-64 functionality required by the topology is a six-stage cascaded design, whereby each stage performs a divide-by-two operation. The first three stages of the design make use of TSPC logic in order to minimize the overall power consumption of the receiver, while switching fast enough to divide down the output frequency of the VCO under all operating conditions. Subsequent stages use simple CMOS flip-flops to divide by two, minimizing the overall power consumption yet again, once the frequency of the signal path is compatible with the CMOS logic available in the design kit.

The flexibility of the test chip, whereby the loop can be operated with the inverting loop-filter buffer either enabled or disabled, requires that a mux circuit be added between the PFD and the primary charge pump. The muxing function is easily accomplished using four CMOS transmission gates that assure the correct closed-loop polarity under all operating scenarios.

To enable testing of the lock-and-roll receiver, two different buffer circuits were designed for use on the unique test chip. The VCO output buffer circuit is a self-biased CMOS inverter with a DC blocking capacitor at the input. The circuit provides the

drive strength required at the input of the divider and the $50\ \Omega$ output buffer that drives the VCO_{Out} test pad, while adding a minimal (and symmetrical) additional load on the tank of the VCO. The $50\ \Omega$ output buffer is a three stage, common-source amplifier design which is used to drive the spectrum analyzer that monitors the VCO output during testing, and the oscilloscope used for monitoring the final demodulated output bitstream from the receiver on the $Bits_{Out}$ pad.

The die was designed to be probed exclusively, using two 8-pin RF probes and four DC probes. There are two additional rows of pads, each nine pads wide, which were uniquely designed such that the necessary open, short, load, through calibration structures, required for configuring the network analyzer to accurately characterize the input match at 5.2 GHz, are available. By rotating the die 180° midway through the calibration process, all four of the required structures can be measured and the die area consumed by the calibration structures is thereby minimized.

While probing the die proved successful for characterizing the input match, measuring the performance of the complete receiver using probes was difficult, and so the die was mounted (and bonded) to a PCB designed by the author for another project at 5.2 GHz, which offered all the required bias controls and signal access points for RX testing following minimal modifications to the PCB.

The output bitstream measured with the RX test chip demonstrates accurate demodulation of the FM input signal, though there is deterministic rail-to-rail output noise which can be attributed to an undersized integrating capacitor on the test chip. The measured power consumption of the receiver agrees well with simulation and the average power consumption that is predicted for a 20% duty cycle is validated to within 5%.

Comparing the measured performance of the lock-and-roll receiver with published alternatives is difficult given the different tradeoff choices that are made by those designing circuits for applications having few standards and regulations. There are RX topologies that consume less power than the lock-and-roll receiver, others that are capable of higher data rates, and some that are compatible with either an on-chip power source or an antenna. The author is unaware of any previously published results, however, suggesting a receiver design that is manufactured in bulk CMOS while operating at a sufficiently high frequency to be compatible with an integrated

antenna while operating from an integrated power source with a die area of $< 4 \text{ mm}^2$. The lock-and-roll receiver is a unique design which can make this claim.

Chapter 8

Conclusion

Despite the current market demands for complicated wireless technologies that communicate large amounts of information at high data rates over long distances, there is also a demand for ultra-low power devices that achieve complete integration while being inexpensive to manufacture in high volume. Chapter 1 outlines the typical requirements of RFID tag devices that must be extremely cheap to manufacture for the purpose of tracking commercial goods, though the data rate, and to a lesser extent the communication range, may be reduced in an effort to minimize cost and overall size. Additionally, Chapter 1 presents the requirements of a medical radiation sensor which would also benefit from a wireless solution that is completely integrated onto a single die so as to eliminate wires which block therapeutic radiation. Though the overall cost of a medical sensor is less of a priority than in the case of an RFID tag, a wireless solution that is free of batteries (which contain heavy metals that can deflect radiation) is also important. These two application examples both highlight the need for a short-range communication system that is completely integrated onto a single die, including the antenna and the power source, where being able to manufacture the solution in an inexpensive bulk CMOS process is a clear benefit as the solution cost is minimized.

Designing CMOS circuits that function as wireless communication devices involves the consideration of numerous tradeoffs. As the previous works summarized in Chapter 2 demonstrate, historically the most successful designs have been the ones that best balanced the tradeoffs so as to optimize the circuits for their intended applications. Circuits that are to be powered by large batteries need not be overly power efficient, and if physical size is not a concern then circuits might be connected to large off-chip antennas to achieve very high communication ranges at the expense of overall

size and power consumption. The fundamental relationships between antenna gain and antenna size and between antenna gain and the frequency of the signal suggest that if one hopes to minimize the overall size of a solution by making use of a small antenna, the operating frequency should be made high. Unfortunately, the power consumption of a circuit is generally proportional to the speed at which it operates and so communication solutions that are optimized for low power typically operate at lower signal frequencies and make use of physically large antennas to compensate.

Chapter 3 proposes the novel lock-and-roll transceiver which is a unique system that has been optimized simultaneously for low power consumption and small size, facilitating complete integration in bulk CMOS. The lock-and-roll receiver occupies 1 mm^2 and consumes $285 \text{ } \mu\text{W}$ of power when duty cycled. The system consumes such little power that it could potentially be powered using an on-chip power source, even though the system communicates at 5.2 GHz using an on-chip antenna. Compared to previously published low-power transceivers, the lock-and-roll transceiver is the first known transceiver that facilitates complete integration of the circuits, the antenna, and potentially the power source onto a single bulk CMOS die. The solution is capable of a communication range of 1.75 m at a data rate of 5 kb/s when one chip with the -22 dBi on-chip antenna communicates with another chip making use of a 6.7 dBi patch antenna. This communication range can be increased without any change to the hardware configuration if one is willing to lower the data rate.

The lock-and-roll receiver is based on an integer-N PLL which can be operated in open-loop and closed-loop modes. In open-loop mode, the design must maintain a near constant loop filter voltage so as to minimize the frequency drift of the oscillator. The oscillator circuit is injection locked to the input signal which is amplified by the LNA circuit that interfaces to the integrated antenna, and the remaining loop components serve to demodulate the FM data. Implementing the sub-circuits within the lock-and-roll receiver so as to meet the requirements of the overall system involves putting into practice numerous low-power CMOS design techniques.

Chapter 4 demonstrates how the locking bandwidth of an integrated VCO circuit can be optimized for the lock-and-roll receiver application while simultaneously minimizing the circuit's power consumption.

Chapter 5 presents approaches for designing injection-locking circuits, including the LNA circuit in the lock-and-roll receiver which is conjugately matched to the

low impedance of the on-chip antenna while achieving 20 dB of gain. As the LNA output is coupled into the tank of the VCO circuit, the output impedance of the LNA is carefully designed so as not to disrupt the VCO's oscillation frequency.

Chapter 6 outlines strategies for the design of circuits that are truly unique to the lock-and-roll transceiver, allowing the PLL loops to be opened and closed while minimizing VCO drift in the open-loop mode. Examples of these circuits are the integrated loop filter which makes exclusive use of MIM capacitors and poly resistors, the loop switch with dummy cells, the unity-gain loop buffer, and the charge pump with high output impedance.

The lock-and-roll receiver test chip that is discussed in Chapter 7 puts the design strategies from Chapters 4 through 6 into practice and implements the unique receiver proposed in Chapter 3 in a bulk CMOS process. The test chip allows for the on-chip input match to be probed and characterized, where the measured impedance lends credence to the claim that the receiver is compatible with an on-chip antenna. The measured RX output demonstrates successful demodulation of BFSK data at data rates up to about 5 kb/s, with an average receiver power consumption of 285 μ W when duty cycled in a typical application – suggesting compatibility with integrated power sources occupying less than 3 mm².

8.1 Thesis Contributions

The thesis contributions towards improving the state of the RFIC design art are as follows:

1. An exploration of previous low-power, completely (or mostly) integrated RFIC topologies for short range, low-speed data reception suggesting new approaches for overcoming their limitations.
2. Demonstration of the feasibility of the new approaches proposed in 1 using a uniquely modified and completely integrated PLL (the lock-and-roll receiver), operated in the open and closed states, as an FM demodulator to a BFSK input signal having a center frequency of 5.2 GHz.

3. Demonstration of the feasibility of a system, namely the lock-and-roll receiver, making use of an on-chip antenna for the purpose of short range communications at 5.2 GHz.
4. Achieve such low power consumption from the lock-and-roll receiver design that powering the integrated circuits with ultracapacitors, which can be charged using a solar cell, is theoretically feasible, therefore proving, to a first order, that the entire receiver can be integrated onto a single chip including the antenna and power supply.

The first contribution is accomplished by Chapters 2 and 3 which review the advantages and disadvantages of previous designs, new and novel developments in the areas of antenna, crystal, and power supply miniaturization and integration, and propose the novel lock-and-roll transceiver of which the lock-and-roll receiver circuit is a critical part.

The second contribution is accomplished by Chapters 4 through 7 where the design of the lock-and-roll receiver circuit and its individual sub-blocks are presented along with measured results, attesting to the success of the unique circuits (and design strategies) and to the new system architecture outlined in Chapter 3.

The third contribution is accomplished by Chapter 5, whereby the suitability of the conjugate on-chip input match between the LNA circuit and the integrated antenna structure is implied by the input impedance that is measured through RF probing of the LNA. An on-chip antenna was not implemented on the RX test chip due to space constraints and to enable accurate probing of the LNA's matching network. Confirmation of a successful match validates the communication range calculations outlined in Chapter 3 and thereby demonstrates the feasibility of using an on-chip antenna with the lock-and-roll receiver.

The fourth and final contribution is accomplished through the power consumption measurements of the RX test chip which are outlined in Chapter 7. When compared with the averaged power consumption calculations from Chapter 3 and published A/hr ratings for integrated ultracapacitors occupying 3 mm², the results demonstrate that powering the chip by means of an integrated ultracapacitor is theoretically feasible, thereby suggesting that a completely integrated transceiver is achievable.

8.2 Publications and Major Recognition/Awards Resulting from this Work

The author and his colleagues have published one IEEE journal paper and seven IEEE conference papers of which the author was the primary author of four, presenting simulated and measured results of the lock-and-roll transceiver topology. These publications, listed most recent to the oldest, are as follows:

1. A. Shamim, V. Karam, **P. Popplewell**, L. Roy, J. Rogers, and C. Plett, “A CMOS Active Antenna/Inductor for System-on-a-Chip (SoC) Applications,” Proceedings of the IEEE Antennas and Propagation Society International Symposium, July, 2008, pp. 1-4. [89]
2. **P. Popplewell**, V. Karam, A. Shamim, J. Rogers, L. Roy, and C. Plett, “A 5.2 GHz BFSK Transceiver Using Injection-Locking and an On-Chip Antenna,” IEEE Journal of Solid-State Circuits, April, 2008, pp. 981-990. [90]
3. **P. Popplewell**, V. Karam, A. Shamim, J. Rogers, and C. Plett, “An Injection-Locked 5.2 GHz SoC Transceiver with On-Chip Antenna for Self-Powered RFID and Medical Sensor Applications,” Symposium on VLSI Circuits Digest of Technical Papers, June, 2007, pp. 669-672. [91]
4. **P. Popplewell**, V. Karam, A. Shamim, J. Rogers, and C. Plett, “A 5.2 GHz BFSK Receiver with On-Chip Antenna for Self-Powered RFID and Medical Sensors,” Proceedings of the IEEE RFIC Symposium, June, 2007, pp. 88-89. [92]
5. V. Karam, **P. Popplewell**, A. Shamim, J. Rogers, and C. Plett, “A 6.3 GHz BFSK Transmitter with On-Chip Antenna for Self-Powered Medical Sensor Applications,” Proceedings of the IEEE RFIC Conference, June, 2007, pp. 101-104. [93]
6. **P. Popplewell**, V. Karam, A. Shamim, J. Rogers, M. Cloutier, and C. Plett, “5.2 GHz Self-Powered Lock-and-Roll Radio using VCO Injection-Locking and On-Chip Antennas,” Proceedings of the IEEE ISCAS Conference, May, 2006, pp. 5203-5206. [94]

7. A. Shamim, **P. Popplewell**, V. Karam, L. Roy, J. Rogers, and C. Plett, "*Silicon Differential Antenna/Inductor for Short Range Wireless Communication Applications*," Proceedings of the IEEE Canadian Conference on Electrical and Computer Engineering, May, 2006, pp. 94-97. [48]
8. A. Shamim, **P. Popplewell**, V. Karam, L. Roy, J. Rogers, and C. Plett, "*5.2 GHz On-Chip Antenna/Inductor for Short Range Wireless Communication Applications*," Proceedings of the IEEE IWAT Workshop, March, 2006, pp. 213-216. [95]

The author and his colleague Victor Karam have filed a U.S. Patent, a Patent Cooperation Treaty (PCT) application and a Canadian Patent on the lock-and-roll transceiver topology. The author and Victor Karam secured funding support from Carleton University's Foundry Program for \$25,000 to cover patenting expenses.

As a result of his involvement in the research, the author has received the following recognitions:

1. Awarded one of only two (worldwide) IEEE Solid-State Circuits Society Pre-doctoral Fellowships for 2006-2007
2. Awarded the Ottawa Center of Research and Innovation's (OCRI) Student Researcher of the Year Award for 2007
3. Invited along with his co-authors, based on his presentation of paper 3 listed above at VLSI 2007, to submit a paper on the lock and roll transceiver topology to the April, 2008, special issue of the Journal of Solid-State Circuits (JSSC)

Additionally, the author and his colleague Victor Karam have competed in two prestigious local business case competitions, pitching the commercialization potential of the lock-and-roll transceiver with favourable results:

1. Won the 2007 Technology Venture Challenge (TVC)
2. Finished second place in Carleton University's 2007 Wesley Nicol Business Case Competition

8.3 Future Work

While the lock-and-roll receiver test chip proved valuable for demonstrating the potential of the lock-and-roll receiver topology, there is room for improvement and opportunities for further research to build on the success of the project towards achieving the goal of a completely integrated, self-powered transceiver. Some of the logical next steps are as follows:

1. Fabricate an updated version of the lock-and-roll receiver whereby the input match is fine tuned (given the measured results) and physically connected to an instance of the single turn loop antenna on a CMOS die such that a true communication range test can be performed using the lock-and-roll RX and TX communicating together (note that to do so an updated version of the lock-and-roll TX is also required given that the tunable range of the TX as measured [52] was unable to achieve 5.2 GHz). The on-chip integrating capacitor at the receiver's output should also be increased to clean up the noise on the output bitstream as suggested in section 7.5.1.
2. Fabricate an integrated ultracapacitor power cell for verification of the power sourcing capabilities suggested in literature, perhaps on its own for initial testing, and then on top of the lock-and-roll receiver along with an integrated solar cell to facilitate testing of the receiver's performance when operating from an integrated power source.
3. Replace the ideal reference for the lock-and-roll receiver PLL with a completely integrated one, and also with a quartz crystal reference circuit for comparison, evaluating the inevitable performance vs. integration tradeoff that is likely to result from using a completely integrated or semi-integrated reference circuit.
4. Combine the lock-and-roll RX and TX circuits on a single die with an input matching circuit that allows both transmit and receive paths to share the same integrated antenna. If the LNA circuit in the lock-and-roll receiver were modified slightly such that the gate terminals of the gain devices were shorted together when the LNA was disabled using a physically large MOSFET device with an R_{on} of about $2\ \Omega$, the very same matching network that was designed

to achieve a conjugate match when the LNA is enabled would translate the impedance seen at the input of the match (i.e. from the antenna) to look like a modestly high- Q inductor. In fact this makes sense intuitively as if the right side (LNA side) of the input match is essentially shorted out when disabled, the shunt capacitor is also shorted out and all that is seen by the antenna is the series combination of the two matching inductors. The Q of the inductors used in the match simulated to be about 14 at 5.2 GHz, which is roughly 50% higher than the estimated Q of the inductor/antenna (see section 3.2). Thus, if the LNA circuit was changed in this manner and the antenna was connected directly to both the core of the VCO in the lock-and-roll transmitter and the input match of the lock-and-roll receiver, the transmitter's VCO would essentially see only an extra parallel inductance in the tank that could potentially be budgeted for and included in the tank design – possibly maintaining the same center frequency and transmit power as the original topology. The tuning range may be reduced as the ratio of inductance to capacitance will clearly have been altered, but a re-balancing of the ratio of fixed to tunable capacitance might overcome this effect. Naturally the impedance presented by the transmitter to the LNA's input match, when the transceiver is in receive mode, must also be considered, but perhaps a similar impedance translation game can be played.

Appendix A

Oscillator Design Fundamentals – Supplementary Information

A.1 Barkhausen Criteria Derivation using the Linear Model

The Barkhausen criteria is best explained with the use of the simple model for a feedback system. Consider the transfer function of the classical feedback system shown in Figure A.1.




Figure A.1: Simple Model of a Feedback System

The transfer function can be derived as follows:

$$Out(s) = A(s)[In(s) + B(s)Out(s)] \quad (A.1)$$

$$Out(s)[1 - A(s)B(s)] = A(s)In(s) \quad (A.2)$$

$$\frac{Out(s)}{In(s)} = \frac{A(s)}{1 - A(s)B(s)} \quad (A.3)$$

From (A.3) one can conclude that if $A(s)B(s) = 1$ then the denominator of the gain expression goes to zero and we have infinite output signal regardless of the input signal. In fact, this result defines the condition for oscillation.

More formally, however, the frequencies at which this condition occurs are referred to as the poles of the system. In solving for the poles of the system we simply determine the requirements on $A(s)$ and $B(s)$ such that the denominator of (A.3) is zero. By noting that s represents $j\omega$ and that $A(s)$ and $B(s)$ are complex expressions, we can conclude that for

$$1 - A(s)B(s) = 0 \quad (\text{A.4})$$

we require

$$|A(j\omega)||B(j\omega)| = 1 \quad (\text{A.5})$$

and




$$\angle[A(j\omega)B(j\omega)] = 2n\pi \quad (\text{A.6})$$

where n is a positive integer number.

Equations (A.5) and (A.6) formally define the Barkhausen criteria for oscillation, which essentially concludes that for sustained oscillations to exist at any particular frequency, the gain around the loop must be unity and the phase must be a positive integer multiple of 2π .

A.2 The $-G_m$ Oscillator and Gain Margin

Figure A.2 shows what is likely the most common LC oscillator design used in modern integrated circuit design, referred to as the “ $-G_m$ ” oscillator. Like the less popular Hartly or Colpitts oscillators [46] shown in Figure A.3 and Figure A.4, the $-G_m$ oscillator is made up of an LC resonator tank and a gain block which provides feedback. Where the Hartly and the Colpitts oscillators provide feedback to an intermediate node in the LC resonator, having tapped inductors and capacitors respectively, the $-G_m$ oscillator is typically a better balanced topology, sampling the voltage across the whole resonator and injecting current into the same node(s) being sampled to sustain oscillations.

Figure A.2: $-G_m$ LC OscillatorFigure A.3: Hartley LC OscillatorFigure A.4: Colpitts LC Oscillator

Considering the $-G_m$ oscillator model, all the losses in the circuit can be lumped into a single parallel resistance represented in the RLC tank as seen in Figure A.2, labelled as R_L . Similarly, all the sources of gain (often made up of two or four transistors in a CMOS design) can be lumped together and represented by the term G_m in the model. At resonance, the impedance of the resonant tank is simply $R_L \Omega$. The equivalent admittance is $1/R_L$ S. The admittance of the gain block, however, which actually sources current for an applied voltage because of the way it is connected (and is therefore behaving like a negative resistance), is $-G_m$ S. As the gain block is in parallel with the RLC resonator from the point of view of output voltage taken across the resonator, the equivalent net admittance is $1/R_L - G_m$, or a net impedance of $1/(1 - R_L G_m)$. What this simple analysis tells us is that if the combined transconductance of the transistors that make up the gain blocks within the $-G_m$ oscillator is larger than $1/R_L$, then the net impedance of the circuit is negative and oscillations will grow. Similarly, if the net transconductance is less than $1/R_L$, the oscillations will decay or cease to develop at startup and the circuit is a lousy oscillator.

Another way of analyzing the circuit is to break the feedback loop as shown in Figure A.5 and to calculate the open-loop gain.




Figure A.5: $-G_m$ LC Oscillator, Open-Loop Analysis

Referring to Figure A.5, the open-loop gain at resonance can be calculated as

$$V_{out} = I_{tank}R_L \quad (\text{A.7})$$

$$= V_{in}G_mR_L \quad (\text{A.8})$$

$$\frac{V_{out}}{V_{in}} = G_mR_L \quad (\text{A.9})$$

Again, the result is the same. If $G_m \gg 1/R_L$ then the loop gain is greater than unity and the condition promotes oscillations, if not then the loop gain is less than unity and oscillations will decay and eventually die out.

As the only initial input to an oscillator at startup is thermal noise, the loop gain must be greater than unity at some frequency in order for oscillations to develop, and the amount by which the gain exceeds unity (or 0 dB) is called the gain margin of the design.

For example, a typical cross coupled NMOS $-G_m$ integrated oscillator might have a gain block made up of two transistors with an effective transconductance of 0.02 A/V. The circuit might make use of a 3.4 nH on-chip inductor with a Q of 7, and a total tank capacitance of about 1.3 pF to yield a resonant frequency of 2.4 GHz. Note that the concept of tank Q and resonance is discussed in more detail in section A.4. The resulting R_L for the circuit will likely be about 350 Ω in this case. A reasonable estimate of the open-loop gain is therefore $\frac{V_{out}}{V_{in}} = 0.02 * 350 = 7$ V/V, or 17 dB, thus the gain margin is 17 dB. Typically, good VCO designers will optimize their circuits for a gain margin of 10 to 20 dB such that even over process corners and temperature fluctuations the VCO is guaranteed to startup. If one were to require a higher gain margin from this design example, the bias current through the transistors or their physical size could be increased in order to increase G_m . Similarly, increasing the inductor's Q , if possible, would also improve the gain margin. If the gain margin is thought to be excessively high because simulations over corners and temperature don't show G_m or R_L to drop to problematic levels, then bias current can likely be reduced in favour of better power consumption, or the physical size of the transistors could be reduced to conserve die area. Clearly the tradeoffs outlined in Figure 2.1 apply to LC oscillator design.

A.2.1 Transistor Operating Point and the Effect on Output Impedance

Consider the standard enhancement NMOS device curves [57] shown in Figure A.6.




Figure A.6: NMOS Modes of Operation

Recognizing that the small signal output impedance ($r_{out} \equiv r_{ds}$) of an NMOS device is inversely proportional to the slope of the curves in Figure A.6, i.e. $r_{ds} = \partial V_{DS} / \partial I_{DS}$ one can conclude that the output impedance is much higher when the device is operating in the saturation region – a well known property of the NMOS device. This result can be confirmed visually with Figure A.6 by realizing that when the output impedance is high, the output current will vary little with a moderate change in output voltage (as $\partial I_{DS} = \partial V_{DS} / r_{ds}$) if r_{ds} is large, thus the flattest part of the I_{DS} vs. V_{DS} curve represents the region of highest output impedance for a given V_{GS} . The finite output impedance in the saturation region (beyond pinch-off) results from the channel length modulation effect which is well documented in many texts that discuss CMOS device theory, [57], [96]. In general, the accepted empirical models [67] for the NMOS transistor's output impedance are

$$r_{ds} = \frac{\partial V_{DS}}{\partial I_{DS}} \quad (\text{A.10})$$

$$= (\mu_n C_{ox} \left(\frac{W}{L}\right) (V_{GS} - V_T - V_{DS}))^{-1} \quad (\text{A.11})$$

$$\approx \frac{1}{g_m} \big|_{V_{DS}=0} \quad (\text{A.12})$$

in the triode region, where W and L are the transistor width and length respectively, V_T is the device's threshold voltage, μ_n is the electron mobility in silicon, C_{ox} is the oxide capacitance per unit gate area, and the last result clearly only applies for the common case of $V_{DS} = 0$.

In the saturation region, at a V_{DS} beyond where channel pinch-off occurs ($V_{DS} \geq V_{GS} - V_T$), the channel length modulation factor (λ) must be taken into account and the output impedance is generally calculated as

$$r_{ds} = \frac{\partial V_{DS}}{\partial I_{DS}} \quad (\text{A.13})$$

$$= (\lambda \left(\frac{\mu_n C_{ox}}{2}\right) \left(\frac{W}{L}\right) (V_{GS} - V_T)^2)^{-1} \quad (\text{A.14})$$

$$\approx \frac{1}{\lambda I_{DS}} \quad (\text{A.15})$$

A.3 MOSFET Noise Theory

There are three sources of noise inherent of CMOS transistors: thermal noise, shot noise, and flicker noise.

A.3.1 CMOS Thermal Noise

CMOS transistors suffer from thermal noise which is typically broken down, for the purpose of modeling, into drain current noise and gate noise. The drain current noise arises from the fact that a MOSFET device is essentially a voltage controlled resistor and resistors are well known to contribute thermal noise (spectral density) of $4kTR$ V^2/Hz where k is Boltzmann's constant (1.38×10^{-23} J/K) and T is the temperature

in units of degrees Kelvin. As such, drain current noise in a MOSFET can be shown [97] to be roughly modeled by a noise current source connected between drain and source with spectral density of

$$\overline{I_{nd}^2} = 4kT\gamma g_{ds0}\Delta f \quad (\text{A.16})$$

where coefficient γ is 1 at $V_{DS} = 0$ and approaches 2/3 at saturation, g_{ds0} is the drain-source conductance with $V_{DS} = 0$, and where (A.16) simplifies to

$$\overline{I_{nd}^2} = 4kT\gamma g_m\Delta f \quad (\text{A.17})$$

for long channel devices.

Gate noise also arises from the thermal agitation of channel charge which translates to noise on the gate current due to the gate to channel capacitance – clearly less of a concern at lower frequencies where the capacitive coupling effect is reduced. Similar to drain noise, the gate noise in a MOSFET can be shown [98] to be roughly modeled by a noise current source connected between gate and source (in parallel with conductance g_g) having spectral density of

$$\overline{I_{ng}^2} = 4kT\delta g_g\Delta f \quad (\text{A.18})$$

where coefficient δ is twice as large as γ (i.e., $\delta=4/3$ in saturation), and where g_g can be calculated as

$$g_g = \frac{\omega^2 C_{gs}^2}{5g_{d0}} \quad (\text{A.19})$$

Thomas Lee[99] points out that measurements of short-channel devices suggest that equations (A.16) and (A.18) are relatively optimistic in their estimate of drain and gate noise respectively (i.e., they suggest lower noise than is the case in reality), unless the parameter δ is assumed to be 2, 3 or even larger. Lee also points out that as the drain and gate noise models share a common origin, as both stem from the thermal excitation of channel charge, they are likely well correlated and therefore δ and γ should maintain their 2:1 relationship regardless of the channel length.

Whether the absolute values of parameters δ and γ in equations (A.16) and (A.18) should be fine tuned in each specific case is somewhat beyond the realm of

concern for the typical VCO designer. What is important to note from this discussion and from equations (A.16) and (A.18), however, is that as the MOSFET devices are made longer, the values of δ and γ reduce and so does the drain and gate noise proportionally (for constant g_m). Designers looking to minimize the drain and gate noise contributions of the CMOS devices themselves towards the overall phase noise response of their VCO designs must keep this result in mind when choosing the optimal transistor device size. Of course, increasing the length of a transistor in favour of having less noise also means that the width must be increased as well if the gain margin is to remain constant (as transistor gain is greatly dependent on the width to length ratio) and increasing both width and length increases the gate area (and thus the gate capacitance) of the device which will affect the center frequency and or tunability of the VCO (see section A.4), and so once again the designer is forced to consider the tradeoffs of Figure 2.1.

A.3.2 CMOS Shot Noise

Shot noise, first described in 1918 by the German physicist Walter Schottky [100], arises from the somewhat sporadic behaviour of electrons as they pass over a potential barrier such as when electrons pass from the emitter to the base region in an NPN bipolar transistor. The term “shot noise” arises not from the name of the physicist that documented it, but from the fact that shot noise on an audio source is said to resemble the sound of buckshot falling on a hard surface. Two conditions are necessary for shot noise to develop, first there must be current flow, and second there must be a potential barrier for the electrons to cross, and it’s the somewhat random timing with which the carriers leap across that barrier that leads to the white noise profile that is characteristic of shot noise. In CMOS transistors, however, only the gate leakage current is a source of shot noise and as gate leakage current is a very small quantity in modern CMOS devices, shot noise is rarely a worry within the context of an overall design, and is therefore often overlooked by even the most cautious CMOS oscillator designer.

A.3.3 CMOS Flicker Noise

Flicker noise, also known as 1/f noise, is a rather mysterious type of noise because no universal mechanism for its creation has been identified [99]. As the 1/f name suggests, flicker noise is characterized by a spectral density that decreases with increasing frequency. Some academics argue that the spectral density also increases without bound as frequency decreases, as the simple 1/f relation would suggest, but our inability to measure noise at infinitesimally small frequency offsets makes absolute verification of this theory all but impossible. Thus, the 1/f relation is assumed at all frequencies leading to the general formula of

$$N^2 = \frac{K}{f^n} \Delta f \quad (\text{A.20})$$

where N^2 is the rms noise in voltage or current, K is a device specific empirical parameter (generally dependent on bias), and the exponent n is usually close to unity. In transistors, 1/f noise is greater in devices where current travels along a horizontal surface (like the horizontal channel under the gate of a MOSFET) than vertically or otherwise. Electron trapping due to defects and impurities at the surface of the channel is typically thought to be the main source of 1/f noise in MOSFETs. The corner frequency is a term used to describe the intersection between the 1/f noise profile of a device and its thermal noise floor. Thus all else being equal, devices with a higher corner frequency have higher overall noise (assuming a constant thermal noise floor). In bipolar devices the corner frequency is typically as good as a few tens or hundreds of herz (or better), while in MOSFET devices the corner frequency is usually much worse, often lying in the range of a few tens of kilohertz or even as bad as a few megahertz.

In MOSFET devices, the 1/f drain noise current is given by

$$i_n^2 = \frac{K}{f} \frac{g_m^2}{WLC_{ox}^2} \Delta f \quad (\text{A.21})$$

where K is typically $10^{-28} \text{ C}^2/\text{m}^2$ for PMOS devices (when buried channels are used) and can be as much as 50 times higher than that for NMOS devices. Equation (A.21) shows that larger MOSFET devices exhibit lower flicker noise and this is due to their larger gate capacitance which smoothes out their channel charge profile.

A.4 LC Resonance, Unloaded and Loaded Tank Q

In LC oscillator designs, it is the LC resonant circuit itself, or the tank circuit, that controls the frequency of oscillation and to a large extent the amplitude of the oscillator output depending on the Q , or quality factor, associated with the tank.

A.4.1 Parallel LC Resonance

Figure A.7 shows a typical parallel RLC tank circuit, with input current I_{in} and output voltage V_{out} measured across the tank. The admittance of the tank, I_{in}/V_{out}




Figure A.7: Parallel RLC Resonant Tank

can be calculated as

$$I_{in}/V_{out} = \frac{1}{R_p} + j\omega C + \frac{1}{j\omega L} \quad (\text{A.22})$$

$$= \frac{1}{R_p} + j\left(\omega C - \frac{1}{\omega L}\right) \quad (\text{A.23})$$

Studying (A.23) one can conclude that the admittance of the tank approaches infinity (equivalent to the impedance of the tank approaching zero) when ω is either 0 or infinite. In other words, the capacitor is dominant and shorts the impedance to zero at high frequencies, while the inductor is dominant and shorts the impedance to zero at low frequencies. The frequency at which the inductive and capacitive admittances

cancel occurs at

$$\omega_0 C - \frac{1}{\omega_0 L} = 0 \quad (\text{A.24})$$

$$\omega_0 = \frac{1}{\sqrt{LC}} \quad (\text{A.25})$$

where ω_0 is known as the resonance frequency of the tank. Note that the mathematical analysis of series *RLC* circuits can easily be carried out in much the same way, yet as the topic is addressed in numerous texts already [99], [46] and is not directly applicable to the oscillator circuits discussed in this thesis, the exercise is left to the reader. At parallel resonance, the capacitor and inductor would appear to cancel each other from an impedance point of view (while in fact the AC currents through either one of those passive components is often quite large when considered independently of the other) leaving only parallel resistance R_p visible to the outside world. Rarely is R_p a physical resistor that is added in parallel with the *LC* tank (as doing so reduces the *Q*), but typically it is a lumped element representation of the imperfect inductor and capacitor, essentially representative of their losses. This realization leads us to the discussion of *Q* factor.

A.4.2 Unloaded *Q* Factor

The classical definition of the term quality factor states that the *Q* of any network is equal to the ratio of stored energy relative to energy lost by the network [99]. For the parallel *RLC* circuit in Figure A.7, the stored energy sloshes back and forth at resonance between the inductor and the capacitor at the resonant frequency, ω_0 . When the tank voltage is at its peak, all the energy stored in the tank is momentarily transferred to the capacitor and using the simple calculation for energy stored in a capacitor one can calculate the total energy being stored (and sloshed) in the tank as

$$E_{\text{stored}} = \frac{1}{2} C V_{pk}^2 = \frac{1}{2} C (I_{pk} R_p)^2 \quad (\text{A.26})$$

The energy that is dissipated by the resistor can be calculated as

$$E_{\text{loss}} = P_{\text{avg}} / \omega_0 = \frac{I_{\text{rms}}^2 R_p}{\omega_0} = \frac{1}{2} I_{pk}^2 R_p \sqrt{LC} \quad (\text{A.27})$$

which results in the Q factor expression

$$Q = \frac{E_{stored}}{E_{loss}} = \frac{1}{\sqrt{LC}} \frac{\frac{1}{2}C(I_{pk}R_p)^2}{\frac{1}{2}I_{pk}^2R_p} = \frac{R_p}{\sqrt{L/C}} \quad (\text{A.28})$$

This Q factor is often referred to as the unloaded Q (sometimes represented as Q_U) within the context of LC VCO design as this is the Q of the tank, alone, unloaded by the transconductance of the VCO circuit. When one considers the transfer function of the unloaded RLC tank circuit, the two sided (or full) 3-dB bandwidth (BW_{3dB}) can be shown [99] to be calculated (given a few reasonable assumptions and simplifications) as $\frac{1}{R_P C}$, much like that of a simple RC circuit but where one considers the frequency relative to ω_0 rather than to DC. As such, normalizing the 3-dB bandwidth with respect to ω_0 leads to the realization that

$$\frac{BW_{3dB}}{\omega_0} = \frac{\sqrt{LC}}{RC} = \frac{\sqrt{L/C}}{R} = 1/Q \quad (\text{A.29})$$

and thus one can simply estimate the Q of a parallel RLC circuit upon reflection of the transfer function or gain response of the network with no prior knowledge of the resistor, inductor, or capacitor values themselves or their losses.

Until this point, the concept of Q factor has been reviewed as it applies to a parallel RLC circuit, but in fact the Q of any network or individual passive device can be calculated, and the Q of a combined LC circuit is equivalent to the Q of the inductor in parallel with that of the capacitor. Recall that R_P is rarely a physical resistor placed in parallel with the resonant LC tank but is more typically just a representation of the equivalent combined parallel losses of the inductor and the capacitor themselves. In the case of modern IC design, the losses associated with capacitors are typically much less than those of inductors when both are manufactured on chip and so the Q of the inductor is almost always lower and will therefore dictate the Q of the overall resonant circuit. Note that the Q of an inductor can be calculated [46] as

$$R_P = Q_L \omega L \quad (\text{A.30})$$

where R_p is the parallel resistance representative of the inductor's inherent losses.

A.4.3 Loaded Q Factor

Where the unloaded quality factor, Q_U , is related to the bandwidth of the unloaded resonant tank according to (A.29), the term loaded Q , or Q_L , is often used to describe the relationship between the bandwidth of the overall VCO (when the tank is loaded by the transconductor such as the case shown in Figure A.2) and the resonant frequency ω_0 . As before, the relationship follows

$$Q_L = \frac{\omega_0}{BW_{3dB,L}} \quad (\text{A.31})$$

The 3-dB bandwidth of a loaded oscillator is in fact very narrow, resulting in typical Q_L values of 80 dB or higher. Recalling the discussion on oscillator gain and frequency response from section 2.3.1, Figure 2.4 illustrates the relationship between Q_L and Q_U which can be summarized [9] as

$$Q_L = \frac{Q_U P_{out}}{F k T B_U} \frac{2}{\pi} \quad (\text{A.32})$$

Where the unloaded Q_U is indicative of the bandwidth of the unloaded LC resonant tank by itself relative to the resonant frequency, the loaded Q_L is representative of the bandwidth of the complete oscillator circuit, or in other words, it dictates the phase noise of the circuit.

References

- [1] “Thomson Nielsen MOSFET Dosimetry,” See <http://www.thomson-elec.com/mosfetdosimetry.htm>.
- [2] B. Razavi, *RF Microelectronics*, Prentice Hall, Inc., Upper Saddle River, NJ, USA, 1998.
- [3] N.C. Chu, K. F. Tsang and P. C. L. Yip, “A low power consumption receiver at LF using frequency shift keying technique,” in *Proceedings of the IEEE Radio Receivers and Associated Systems Conference*, 1995, pp. 67–70.
- [4] B. Fong, A. C. M. Fong and G. Y. Hong, “A low power receiver architecture for mobile biomedicine systems,” in *Proceedings of the IEEE Conference on Electron Devices and Solid-State Circuits*, 2005, pp. 541–544.
- [5] R. B. Waterhouse, S. D. Targonski and D. M. Kokotoff, “Design and performance of small printed antennas,” *IEEE Transactions on Antennas and Propagation*, vol. 46, no. 11, pp. 1629–1633, 1998.
- [6] J. G. Proakis, *Digital Communications*, McGraw-Hill, New York, USA, 2001.
- [7] H. J. Bergveld, K. M. M. van Kaam *et al*, “A low-power highly-digitized receiver for 2.4-GHz-band GFSK applications,” in *Proceedings of the IEEE Radio Frequency Integrated Circuits Symposium*, 2004, pp. 347–350.
- [8] P. H. R. Popplewell, R. E. Amaya, M. Cloutier and C. Plett, “Calibration-free on-chip inductor coupling experiment with injection-lockable VCOs,” in *Proceedings of the IEEE Bipolar/BiCMOS Circuits and Technology Meeting*, September, 2004, pp. 261–264.
- [9] P. H. R. Popplewell, “Using oscillator gain and injection-locking to measure on-chip inductor coupling,” M.S. thesis, Carleton University, 2004.

- [10] W. P. Robins, *Phase Noise in Signal Sources*, Peter Peregrinus Ltd., London, UK, 1984.
- [11] M. Cloutier, "Oscillators," unpublished work discussed in private communications, 2003.
- [12] D. B. Leeson, "Simple model of a feedback oscillator noise spectrum," *Proceedings of the IEEE*, vol. 54, pp. 329–330, February, 1966.
- [13] R. E. Amaya, P. H. R. Popplewell, M. Cloutier and C. Plett, "EM and substrate coupling in silicon RFICs," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 9, pp. 1968–1971, September, 2005.
- [14] R. E. Amaya, P. H. R. Popplewell, M. Cloutier and C. Plett, "Analysis and measurements of EM and substrate coupling effects in common RF integrated circuits," in *Proceedings of the IEEE Custom Integrated Circuits Conference*, October 2004, pp. 363–366.
- [15] R. Adler, "A study of locking phenomena in oscillators," *Proceedings of the IRE*, vol. 34, pp. 351–358, June, 1946.
- [16] B. Razavi, "A study of injection locking in oscillators," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 9, pp. 1415–1424, September, 2004.
- [17] C. DeVries and R. Mason, "A 0.18 μ m CMOS, high Q-enhanced bandpass filter with direct digital tuning," in *Proceedings of the IEEE Custom Integrated Circuits Conference*, 2002, pp. 279–282.
- [18] B. Pham and A. Dinh, "Dual-mode tunable Q-enhanced filter for narrowband and UWB systems," in *Proceedings of the Canadian Conference on Electrical and Computer Engineering*, 2006, pp. 2318–2321.
- [19] G. Jiandong and A. Dinh, "A 2.6 GHz tunable CMOS bandpass filter using Q-enhanced circuit," in *Proceedings of the Pacific Rim Conference on Communications, Computers and Signal Processing*, 2003, pp. 57–61.
- [20] H. Ahmed, C. DeVries and R. Mason, "RF, Q-enhanced bandpass filters in standard 0.18 μ m CMOS with direct digital tuning," in *Proceedings of the IEEE International Symposium on Circuits and Systems*, May, 2003.

- [21] E. Armstrong, "Some recent developments of regenerative circuits," in *Proceedings of the IRE*, August, 1922, vol. 10, pp. 244–260.
- [22] A. Vouilloz, M. Declercq and C. Dehollain, "A low-power CMOS super-regenerative receiver at 1 GHz," *IEEE Journal of Solid-State Circuits*, vol. 36, no. 3, pp. 440–451, March, 2001.
- [23] P. Favre, N. Joehl, A. Vouilloz, P. Deval, C. Dehollain and M. Declercq, "A 2 V 600 μ A 1 GHz BiCMOS super-regenerative receiver for ISM applications," *IEEE Journal of Solid-State Circuits*, vol. 33, pp. 2186–2196, December, 1998.
- [24] "FM demodulator including injection locked oscillator/divider," US Patent # 4,953,010.
- [25] "Phase-locked loop FM demodulator," US Patent # 5,140,278.
- [26] "Low energy consumption RF telemetry control for an implantable medical device," US Patent # 6,456,887.
- [27] R. N. Simons, D. G. Hall and F. A. Miranda, "RF telemetry system for an implantable bio-MEMS sensor," in *IEEE MTT-S International Microwave Symposium*, June, 2004, pp. 1433–1436.
- [28] R. N. Simons, D. G. Hall and F. A. Miranda, "Spiral chip implantable radiator and printed loop external receptor for RF telemetry in bio-sensor systems," in *Proceedings of the IEEE Radio and Wireless Conference*, September, 2004, pp. 203–206.
- [29] A. Burke, "Ultracapacitors: Why, how and where is the technology," *Journal of Power Sources*, vol. 91, pp. 37–50, 2000.
- [30] R. Kotz and M. Carlen, "Principles and applications of electrochemical capacitors," *Electrochimica*, vol. 45, pp. 2483–2498, 2000.
- [31] K. Yamamoto, A. Nakajima, M. Yoshimi, T. Sawada, S. Fukuda, K. Hayashi and T. Suezaki, "High efficiency thin film silicon solar cell and module," in *Record of the Twenty-Ninth IEEE Photovoltaic Specialists Conference*, May 19–24 2002, pp. 1110–1113.

- [32] K. Yamamoto, M. Yoshimi, T. Suzuki, T. Nakata, T. Sawada, A. Nakajima and K. Hayashi, "Large-area and high efficiency a-Si/poly-Si stacked solar cell submodule," in *Record of the Twenty-Eighth IEEE Photovoltaic Specialists Conference*, Sept. 15-22 2000, pp. 1428–1432.
- [33] R. Glidden *et al*, "Design of ultra-low-cost UHF RFID tags for supply chain applications," *IEEE Communications Magazine*, vol. 42, pp. 140–151, August, 2004.
- [34] R. Barnett, G. Balachandran, S. Lazar, B. Kramer, G. Konnail, S. Rajasekhar and V. Drobny, "A passive UHF RFID transponder for EPC Gen 2 with -14 dBm sensitivity in 0.13 μm CMOS," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, February 2007, pp. 582–583.
- [35] A. Safarian, A. Shameli, A. Rofougaran, M. Rofougaran and F. De Flaviis, "An integrated RFID reader," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, February 2007, pp. 218–219.
- [36] I. Kipnis, S. Chiu, M. Loyer, J. Carrigan, J. Rapp, P. Johansson, D. Westberg and J. Johansson, "A 900 MHz UHF RFID reader transceiver IC," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, February 2007, pp. 214–215.
- [37] C. Mikeka and H. Arai, "Development of a batteryless sensor transmitter," in *Proceedings of the IEEE Radio and Wireless Symposium*, 2010, pp. 68–71.
- [38] H. Lhermet, C. Condemine, M Plissonnier, R. Salot, P. Audebert and M. Rosset, "Efficient power management circuit: Thermal energy harvesting to above-IC microbattery energy storage," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, February 2007, pp. 62–63.
- [39] Tele Quarz GmbH & Co., "Quartz Crystal Technical Overview and Product Catalogue," 1981.
- [40] Epson Toyocom Co., "The Crystal Master - Crystal Devices Catalogue 2006.10," 2006.

- [41] M. S. McCorquodale, J. D. O'Day, S. M. Pernia, G. A. Carichner, S. Kubba and R. B. Brown, "A monolithic and self-referenced RF LC clock generator compliant with USB 2.0," *IEEE Journal of Solid-State Circuits*, vol. 42, pp. 385 – 399, February 2007.
- [42] S. Cho and A. P. Chandrakasan, "A 6.5 GHz energy-efficient BFSK modulator for wireless sensor applications," *IEEE Journal of Solid-State Circuits*, pp. 731–739, 2004.
- [43] S. Willingham, M. Perrott, B. Setterberg, A. Grzegorek and B. McFarland, "An integrated 2.4 GHz frequency synthesizer with 5 μ s settling and 2 Mbps closed loop modulation," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, 2000, pp. 200–201.
- [44] A. Yamagishi, M. Ugajin and T. Tsukahara, "A 1 V 2.4 GHz PLL synthesizer with a fully differential prescaler and a low-off-leakage charge pump," in *IEEE MTT-S International Microwave Symposium*, June, 2003, pp. 733–736.
- [45] "Ansoft - HFSS," See <http://www.ansoft.com/products/hf/hfss/>.
- [46] J. Rogers and C. Plett, *Radio Frequency Integrated Circuit Design*, Artech House, Boston, USA, 2003.
- [47] "Agilent Technologies - Advanced Design System," See <http://eesof.tm.agilent.com/products/adsoview.html>.
- [48] A. Shamim, P. Popplewell, V. Karam, L. Roy, J. Rogers and C. Plett, "Silicon differential antenna/inductor for short range wireless communication applications," in *Proceedings of the IEEE Canadian Conference on Electrical and Computer Engineering*, May, 2006, pp. 94–97.
- [49] C. Balanis, *Antenna Theory: Analysis and Design, 3rd Edition*, John Wiley & Sons, Inc., Toronto, Canada, 2005.
- [50] H. T. Friis, "A note on a simple transmission formula," in *Proceedings of the IRE*, May, 1946, pp. 254–256.

- [51] J. Rogers, C. Plett and F. Dai, *Integrated Circuit Design for High-Speed Frequency Synthesis*, Artech House, Boston, USA, 2006.
- [52] V. F. Karam, *Techniques for Low-Power CMOS Transmitter System Integration for Short-Range Radio-Frequency Communication*, Ph.D. thesis, Carleton University, 2008.
- [53] P. H. R. Popplewell, R. E. Amaya, M. Cloutier and C. Plett, "Analysis and measurements of injection-lockable oscillators for high-gain, high-Q applications," in *Proceedings of the IEEE International Analog VLSI Workshop*, October, 2004, pp. 35–40.
- [54] K. Kurokawa, "Injection locking of microwave solid-state oscillators," *Proceedings of the IEEE*, vol. 61, pp. 1386–1410, October, 1973.
- [55] X. Wang and N. J. Gomez, "Locking bandwidth equations for electrically and optically injection-locked oscillators," in *Proceedings of IEEE Optoelectronics*, Dec. 29 2004, pp. 476–481.
- [56] D. Ham and A. Hajimiri,, "Concepts and methods in optimization of integrated LC VCOs," *IEEE Journal of Solid-State Circuits*, pp. 896–909, June, 2001.
- [57] A. S. Sedra and K. C. Smith, *Microelectronic Circuits*, Holt, Rinehart and Winston, Toronto, Canada, 1987.
- [58] H. Ahmed, C. DeVries and R. Mason, "A digitally tuned 1.1 GHz subharmonic injection-locked VCO in 0.18 μ m CMOS," in *Proceedings of the IEEE European Solid-State Circuits Conference*, 2003, pp. 81–84.
- [59] C. DeVries and R. Mason, "A 0.18 μ m CMOS 900 MHz receiver front-end using RF Q-enhanced filters," in *Proceedings of the IEEE International Symposium on Circuits and Systems*, May, 2004, vol. 4, pp. 279–282.
- [60] C. DeVries and R. Mason, "Subsampling architecture for low power receivers," in *IEEE Transactions on Circuits and Systems II*, April, 2008, vol. 55, pp. 304–308.

- [61] H. T. Friis, "Noise figures in radio receivers," in *Proceedings of the IRE*, July, 1944, pp. 419–422.
- [62] S. Voinigescu, M. Maliepaard, J. Showell *et al*, "A scalable high-frequency noise model for bipolar transistors with application to optimal transistor sizing for low-noise amplifier design," *IEEE Journal of Solid-State Circuits*, vol. 32, no. 9, pp. 1430–1439, September, 1997.
- [63] C. Bowick, *RF Circuit Design*, H. W. Sams, Indianapolis, USA, 1982.
- [64] A. Hastings, *The Art of Analog Layout*, Prentice-Hall Inc., New Jersey, USA, 2001.
- [65] "Agilent Technologies - AppCAD," See <http://www.hp.woodshot.com/>.
- [66] "Eagleware Corporation [now Agilent Technologies] - PLL," See <http://www.eagleware.com>.
- [67] D. A. Johns and K. Martin, *Analog Integrated Circuit Design*, John Wiley & Sons, Inc., Toronto, Canada, 1997.
- [68] R. Schaumann, M. Ghausi and K. Laker, *Design of Analog Filters - Passive, Active RC and Switched Capacitor*, Prentice Hall, Inc., Englewood Cliffs, NJ, USA, 1990.
- [69] R. Gregorian and G. Temes, *Analog MOS Integrated Circuits for Signal Processing*, John Wiley & Sons, Inc., New York, NY, USA, 1986.
- [70] T. Soorapanth, C. Yue *et al*, "Analysis and optimization of accumulation-mode varactor for RF ICs," in *IEEE Symposium on Very Large Scale Integrated Circuits Digest of Technical Papers*, 1998, pp. 32–33.
- [71] J. Yuan and C. Svensson, "High-speed CMOS circuit technique," *IEEE Journal of Solid-State Circuits*, vol. 24, no. 1, pp. 62–70, February, 1989.
- [72] A. Porret, T. Melly, C. Enz and E. Vittoz, "A low-power low-voltage transceiver architecture suitable for wireless distributed sensors network," in *Proceedings of the IEEE International Symposium on Circuits and Systems*, May, 2000, vol. 1, pp. 56–59.

- [73] C. Enz, N. Scolari and U. Yodprasit, "Ultra low-power radio design for wireless sensor networks," in *Proceedings of the IEEE International Workshop on Radio-Frequency Integration Technology*, 2005, pp. 1–17.
- [74] H. Seo, Y. Moon *et al*, "A low power fully CMOS integrated RF transceiver IC for wireless sensor networks," *IEEE Transactions on Very Large Scale Integration Systems*, vol. 15, no. 2, pp. 227–231, February, 2007.
- [75] A. Molnar, B. Lu, S. Lanzisera, B. Cook and K. Pister, "An ultra-low power 900 MHz RF transceiver for wireless sensor networks," in *Proceedings of the IEEE 2004 Custom Integrated Circuits Conference*, October, 2004, pp. 401–404.
- [76] R. van Langevelde, M. van Elzakker *et al*, "An ultra-low-power 868/915 MHz RF transceiver for wireless sensor network applications," in *Proceedings of the IEEE Radio Frequency Integrated Circuits Symposium*, 2009, pp. 113–116.
- [77] D. Daly and A. Chandrakasan, "An energy-efficient OOK transceiver for wireless sensor networks," May, 2007, vol. 42, pp. 1003 – 1011.
- [78] C. Chen, H. Lee *et al*, "Low-power 2.4-GHz transceiver in wireless sensor network for bio-medical applications," in *Proceedings of the IEEE Biomedical Circuits and Systems Conference*, 2007, pp. 239–242.
- [79] B. Otis, Y. Chee and J. Rabaey, "A 400 μ W-RX 1.6 mW-TX super-regenerative transceiver for wireless sensor networks," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, February, 2005, pp. 396–397.
- [80] B. Cook, A. Berny, A. Molnar, S. Lanzisera and K. Pister, "An ultra-low power 2.4 GHz RF transceiver for wireless sensor networks in 0.13 μ m CMOS with 400 mV supply and an integrated passive RX front-end," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, February, 2006, pp. 370–371.
- [81] W. Kluge, F. Poegel *et al*, "A fully integrated 2.4 GHz IEEE 802.15.4 compliant transceiver for zigbee applications," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, February, 2006, pp. 372–373.

- [82] J. Chen, M. Flynn and J. Hayes, "A fully integrated autocalibrated super-regenerative receiver," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, February, 2006, pp. 376–377.
- [83] S. Verma, J. Xu, M. Hamada and T. H. Lee, "A 17 mw 0.66 mm direct-conversion receiver for 1 Mb/s cable replacement," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, February, 2005, pp. 548–549.
- [84] V. Peiris, C. Arm *et al*, "A 1 V 433/868 MHz 25 kb/s-FSK 2 kb/s-OOK RF transceiver SoC in standard digital 0.18 μm CMOS," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, February, 2005, pp. 258–259.
- [85] K. Muhammad, D. Leipold *et al*, "A discrete-time bluetooth receiver in a 0.13 μm digital CMOS process," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, February, 2004, pp. 268–269.
- [86] C. Cojocaru, T. Pamir *et al*, "A 43 mW bluetooth transceiver with -91 dBm sensitivity," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, February, 2003, pp. 90–91.
- [87] A. Porret, T. Melly, D. Python, C. Enz and E. Vittoz, "An ultralow-power UHF transceiver integrated in a standard digital CMOS process: Architecture and receiver," *IEEE Journal of Solid-State Circuits*, vol. 36, no. 3, pp. 452 – 466, March, 2001.
- [88] C. Hwang, I. McGregor *et al*, "An ultra-low power OOK RF transceiver for wireless sensor networks," in *Proceedings of the IEEE European Microwave Conference*, 2009, pp. 1323 – 1326.
- [89] A. Shamim, P. Popplewell, V. Karam, L. Roy, J. Rogers and C. Plett, "A CMOS active antenna/inductor for system on a chip (SoC) applications," in *Proceedings of the IEEE Antennas and Propagation Society International Symposium*, July, 2008, pp. 1–4.

- [90] P. Popplewell, V. Karam, A. Shamim, J. Rogers, L. Roy and C. Plett, "A 5.2-GHz BFSK transceiver using injection-locking and an on-chip antenna," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 4, pp. 981–990, April, 2008.
- [91] P. Popplewell, V. Karam, A. Shamim, J. Rogers and C. Plett, "An injection-locked 5.2 GHz SoC transceiver with on-chip antenna for self-powered RFID and medical sensor applications," in *IEEE Symposium on Very Large Scale Integrated Circuits Digest of Technical Papers*, June, 2007, pp. 88–89.
- [92] P. Popplewell, V. Karam, A. Shamim, J. Rogers and C. Plett, "A 5.2 GHz BFSK receiver with on-chip antenna for self-powered RFID tags and medical sensors," in *Proceedings of the IEEE Radio Frequency Integrated Circuits Symposium*, June, 2007, pp. 669–672.
- [93] V. Karam, P. Popplewell, A. Shamim, J. Rogers and C. Plett, "A 6.3 GHz BFSK transmitter with on-chip antenna for self-powered medical sensor applications," in *Proceedings of the IEEE Radio Frequency Integrated Circuits Symposium*, June, 2007, pp. 101–104.
- [94] P. Popplewell, V. Karam, A. Shamim, J. Rogers, M. Cloutier and C. Plett, "5.2 GHz self-powered lock and roll radio using VCO injection-locking and on-chip antennas," in *Proceedings of the IEEE International Symposium on Circuits and Systems*, 2006.
- [95] A. Shamim, P. Popplewell, V. Karam, L. Roy, J. Rogers and C. Plett, "5.2 GHz on-chip antenna/ inductor for short range wireless communication applications," in *Proceedings of the IEEE International Workshop on Antenna Technology, Small Antennas and Novel Metamaterials*, March, 2006, pp. 213–216.
- [96] P. R. Gray and R. G. Meyer, *Analysis and Design of Integrated Circuits*, John Wiley & Sons, Inc., New York, USA, 1993.
- [97] A. van der Ziel, "Thermal noise in field effect transistors," in *Proceedings of the IEEE*, August, 1962, pp. 1801–1812.

- [98] A. van der Ziel, *Noise in Solid State Devices and Circuits*, John Wiley & Sons, Inc., New York, USA, 1986.
- [99] T. H. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*, Cambridge University Press, Cambridge, UK, 1998.
- [100] W. Schottky, “Über spontane stromschwankungen in verschiedenen elektrizitätsleitern,” in *Annalen der Physik*, 1918, vol. 57, pp. 541–567.