NOTE TO USERS

This reproduction is the best copy available.

UMI[®]

Low Power Sub-sampling Radio Receiver Architecture

by

Christopher Andrew DeVries

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 Engineering
Department of Electronics
Faculty of Engineering
Carleton University
Ottawa, Ontario, Canada
December 2007

Copyright © 2007 by Christopher A. DeVries

Library and
Archives Canada

Bibliothèque et
Archives Canada

Published Heritage
Branch

Direction du
Patrimoine de l'édition

395 Wellington Street
Ottawa ON K1A 0N4
Canada

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

ISBN: 978-0-494-36780-3

Our file Notre référence

ISBN: 978-0-494-36780-3

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.

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

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

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

Abstract

Radio transceivers have evolved over the last several decades from simple, discrete component based circuits into low power, highly integrated circuits. This transformation has been driven largely by the requirement for lower cost and lower power transceivers with more functionality. New applications for wireless digital communications have emerged that have created demand for these types of transceivers. In highly integrated wireless transceivers that are intended for mobile communications, the power consumption of the receiver is often the most important concern to the radio designer.

This research is focused on architectures for low power integrated radio receivers. The architecture of the sub-sampling radio receiver is identified as a leading candidate for this type of radio. The sub-sampling receiver is analyzed in detail to show design constraints and fundamental relationships specific to this architecture. A specific implementation of the sub-sampling receiver is proposed and block level requirements are derived.

The fundamental concept and requirement of bandpass filtering to enable a sub-sampling architecture to work is explained in detail. The analysis also shows that the bandpass filtering required is high Q and requires high precision tuning. A Q-enhanced filter with direct tuning is proposed as the solution to the bandpass anti-aliasing requirements of the sub-sampling receiver.

Acknowledgements

I would like to thank all those who helped me my thesis research. Specifically I would like to thank Dr. Ralph Mason, my thesis supervisor who has been instrumental in aiding my research and giving me guidance.

I would like to thank my family for their continued support and sacrifice. My wife, Amanda has always been an excellent supporter and aided me throughout my thesis research. My children Maya and Julian have provided the necessary distractions during the thesis writing process. I would also like to thank my parents and brothers who have continued to show great interest and support in what I am doing.

I would like to thank Carleton University and specifically the Electronics Department who have played a crucial role in enabling me to do my research, especially in the early stages. I would also like to thank the Canadian Micro-electronics Corporation (CMC) for providing assistance and allowing student access to process technology through their grant program.

I would also like to thank my employer, Kleer Corporation for allowing me the time away from work in order to prepare this thesis.

Finally, I would like to thank all my friends and colleagues for their support and input over the years.

Table of Contents

Abstract.....	ii
Acknowledgements	iii
Table of Contents	iv
List of Tables	vi
List of Figures.....	vi
List of Acronyms.....	ix
Chapter 1 Introduction.....	1
1.1 Thesis Motivation	2
1.2 Thesis Outline	2
1.3 Contributions.....	3
Chapter 2 Radio Technology.....	5
2.1 Introduction.....	5
2.2 Types of Wireless Systems	5
2.3 Link Budget	8
2.4 Modulation.....	10
2.4.1 Spread-Spectrum.....	12
2.4.2 Concerns for the Analog Radio Designer	13
2.5 Conclusions.....	14
Chapter 3 Receiver System Requirements	15
3.1 Introduction.....	15
3.2 Power Consumption.....	16
3.3 Sensitivity / Noise Figure.....	18
3.4 Interference Performance.....	20
3.4.1 Selectivity	20
3.4.2 Inter-modulation Distortion	21
3.4.3 Dynamic Range.....	23
3.4.4 Image Rejection	25
3.5 EVM.....	26
3.6 Phase Noise.....	30
3.6.1 Self-Interference	30
3.6.2 Reciprocal Mixing	31
3.7 Inter-Symbol Interference.....	33
3.8 Overall Performance budgets.....	33
3.9 Conclusions.....	35
Chapter 4 Receiver Architectures	36
4.1 Introduction.....	36
4.2 Integrated Receiver Architecture Descriptions.....	37
4.2.1 Generalized Receiver	37
4.2.2 Super-heterodyne	39
4.2.3 Direct Conversion	41
4.2.4 Low-IF	42
4.2.5 Oscillator Receivers	43

4.2.6	Sub-Sampling.....	45
4.3	Receiver Architecture Summary.....	50
4.4	Conclusions.....	53
Chapter 5	Sub-Sampling Receiver Architecture.....	55
5.1	Introduction.....	55
5.2	Sub-Sampling Frequency Plan.....	55
5.3	Noise and Image Folding.....	58
5.4	Basic Architecture.....	59
5.5	Selectivity & Image Rejection.....	61
5.5.1	Generalized 2 nd order Filter Requirement.....	62
5.5.2	Higher order Anti-aliasing Filtering.....	68
5.6	Gain & Noise Figure.....	70
5.7	Phase Noise.....	77
5.8	Conclusions.....	81
Chapter 6	Q-enhanced Filters.....	82
6.1	Introduction.....	82
6.2	Basic Q-enhanced Filter.....	82
6.3	Noise and Linearity.....	86
6.4	Tuning.....	88
6.4.1	Frequency Tuning.....	89
6.4.2	Q Tuning.....	91
6.5	Conclusions.....	96
Chapter 7	Sub-sampling Receiver Implementation.....	97
7.1	Introduction.....	97
7.2	Receiver Description.....	97
7.3	Block and Level Breakdown.....	99
7.3.1	Gain and Noise.....	100
7.3.2	Linearity and Dynamic Range.....	105
7.3.3	Image Rejection.....	111
7.4	Detailed Implementation.....	112
7.4.1	Programmable Bias Levels.....	112
7.4.2	Input Match and Attenuator.....	113
7.4.3	LNA / Filter 1.....	114
7.4.4	Filter 2.....	116
7.4.5	Frequency Control DAC.....	117
7.4.6	Sub-sampling Mixer.....	121
7.4.7	Programmable Gain Amplifier.....	126
7.4.8	ADC.....	129
7.5	Experimental Results.....	140
7.5.1	Test chip #1: RF front-end.....	142
7.5.2	Test-chip #2: Sub-sampling Mixer, PGA and ADC.....	150
7.5.3	Combined Receiver Results.....	161
7.6	Conclusions.....	163
Chapter 8	Conclusions and Recommendations.....	165
8.1	Possibilities for Future Work.....	168

References.....	170
------------------------	------------

List of Tables

Table 2-1: Summary of Wireless System Classifications.....	7
Table 2-2: Example Link Budget.....	9
Table 3-1: Breakdown of Current Consumption for Thermostat.....	17
Table 4-1: Receiver Architecture Summary	51
Table 7-1: Gain and Noise by Block.....	101
Table 7-2: Gain and IIP3 Breakdown by Block.....	106
Table 7-3: Inter-modulation Calculation	110
Table 7-4: RF Front-end Current Consumption Breakdown	147
Table 7-5: RF front-end Linearity at Different Frequency Offsets.....	149
Table 7-6: Summary of Receiver Key Parameters	161
Table 7-7: Comparison of Receiver Power Consumption	164

List of Figures

Figure 2-1: Examples of Signal Constellations	12
Figure 3-1: Block Diagram of Wireless Thermostat Interface	17
Figure 3-2: Inter-modulation Resulting from Non-linearities	22
Figure 3-3: Effect of Filtering on Effective IP3 at a Frequency Offset.....	23
Figure 3-4: Example of Error Vector.....	27
Figure 3-5: EVM with only Noise (above), EVM with Noise and Interference (below) .	29
Figure 3-6: Phase Noise Effect on Symbol.....	31
Figure 3-7: Example of Reciprocal Mixing in a Receiver.....	32
Figure 3-8: Combined Effect of Distortion, Noise and Interference	34
Figure 4-1: General Receiver Subsystems.....	38
Figure 4-2: General Receiver Architecture.....	38
Figure 4-3: Super-Heterodyne Receiver	40
Figure 4-4: Direct Conversion Receiver.....	41
Figure 4-5: Low IF Receiver.....	43
Figure 4-6: Super-Regenerative Receiver by Chen, et. al. [44].....	44
Figure 4-7: Injection Locked Receiver by Popplewell, et. al. [45].....	45
Figure 4-8: Sub-sampling Receiver	46
Figure 4-9: Nyquist Sampling.....	47
Figure 4-10: Sub-Sampling.....	48
Figure 5-1: Sub-sampling Frequency Conversion	56
Figure 5-2: Noise Folding and Images in Sub-sampling	59
Figure 5-3: High Level Diagram of RF Sub-sampling Receiver.....	60
Figure 5-4: Example of the Filter Q and Sampling Frequency Effects on Image	62
Figure 5-5: Image Attenuation versus Q/n ratio (x).....	65

Figure 5-6: k^{th} Image Attenuation versus Q/n ratio (x).....	67
Figure 5-7: k^{th} Image Attenuation versus Q/n ratio (x) with External Attenuation	68
Figure 5-8: Ideal Model of Sampler (Sub-sampling Mixer).....	71
Figure 5-9: Noise Figure for Sampler as a Function of Sampling Frequency	76
Figure 5-10: Comparison of Frequency Conversion Techniques	81
Figure 6-1: Simplified Q-enhanced Filter Circuit Diagram.....	84
Figure 6-2: Differential Q-enhanced Filter	84
Figure 6-3: Basic Block Diagram of Frequency Tuning	89
Figure 6-4: Basic Block Diagram of Q Tuning	93
Figure 6-5: Filter Q Tuning Example Plot.....	95
Figure 7-1: Top Level Diagram	99
Figure 7-2: Input and Interfering Signal Levels at each Block Input	106
Figure 7-3: Image Rejection of Receiver Implementation	112
Figure 7-4: Programmable Bias Used Throughout Test Chips.....	113
Figure 7-5: Circuit Diagram of Input Attenuator.....	114
Figure 7-6: LNA and Q-enhanced Filter.....	115
Figure 7-7: Frequency DAC	118
Figure 7-8: Improved Switch Configuration for Frequency DAC.....	120
Figure 7-9: Unit Capacitors for Frequency DAC	121
Figure 7-10: Circuit Diagram of Sub-Sampling Mixer.....	123
Figure 7-11: Sub-sampling Mixer with Input and Output Buffers	125
Figure 7-12: PGA Circuit Diagram.....	127
Figure 7-13: Single Bit, Discrete Time $\Sigma\Delta$ ADC Block Diagram	131
Figure 7-14: First Order Loop Filter.....	133
Figure 7-15: NTF for Different Orders of Low-pass ADCs	133
Figure 7-16: NTF for Different Orders of Band-pass ADCs.....	134
Figure 7-17: Maximum Theoretical SNR for 1-bit ADCs.....	136
Figure 7-18: Block Diagram of 4 th Order Bandpass ADC.....	136
Figure 7-19: Conceptual Block Diagram of Pseudo-Two Path ADC.....	137
Figure 7-20: ADC Integrator Circuit Diagram	138
Figure 7-21: Simulated ADC Response with a 40 MHz Sampling Clock.....	139
Figure 7-22: Microphotograph of Test Chip 1	140
Figure 7-23: Microphotograph of Test Chip 2.....	141
Figure 7-24: Receiver Blocks Included in Test Chips.....	141
Figure 7-25: Frequency DAC Step Size	144
Figure 7-26: Receiver Input Match.....	145
Figure 7-27: Receiver Noise Figure versus LNA Bias	146
Figure 7-28: Composite Transfer Function of Receiver RF Front-end	147
Figure 7-29: In-band IM3 plot	149
Figure 7-30: Test Chip #2 Simplified Test Board Schematic.....	152
Figure 7-31: Sub-sampling Mixer Isolation.....	152
Figure 7-32: Noise Spectrum with Signal at RF Input Compared to the ADC only with a 25 MHz Sampling Frequency	154
Figure 7-33: PGA Theoretical and Measured Gain	155
Figure 7-34: ADC Spectrum with 2-tone, 0dBm Input Signal to ADC.....	156

Figure 7-35: ADC Spectrum with PGA Input at High Gain.....	157
Figure 7-36: ADC SNR and SNDR.....	158
Figure 7-37: ADC Spectrum with -20 dBm (22.3 mVrms) Sub-sampling Mixer Input.	160
Figure 7-38: Example Received Symbol Constellation.....	163

List of Acronyms

Acronym	Definition
ADC	Analog to Digital Converter
AGC	Automatic Gain Control
AGN	Additive Gaussian Noise
ASK	Amplitude Shift Keying
ATTEN	Attenuator
BER	Bit Error Rate
BPSK	Binary Phase Shift Keying
CMFB	Common-Mode Feedback
CMOS	Complimentary Metal-Oxide Semiconductor
CPE	Consumer Premise Equipment
CRC	Cyclic Redundancy Check
CSR	Cambridge Silicon Radio
DAC	Digital to Analog Converter
DC	Direct Current
DCR	Direct Conversion Receiver
DSSS	Direct Sequence Spread Spectrum
EDR	Extended Data Rate
EVM	Error Vector Magnitude
FHSS	Frequency Hopping Spread Spectrum
FSK	Frequency Shift Keying
GPS	Global Positioning System
IC	Integrated Circuit
IEEE	Institute of Electrical and Electronics Engineers
IF	Intermediate Frequency
IIP3	Input (third order) Intercept Point
IM	Inter-modulation
IP3	Third order Intercept Point
IR	Image Rejection
ISI	Inter-Symbol Interference
ISM	Industrial, Scientific and Medical (unlicensed frequency band)
LAN	Local Area Network
LC	Inductor / Capacitor Tank circuit
LCD	Liquid Crystal Display
LNA	Low Noise Amplifier
LO	Local Oscillator
LPF	Low Pass Filter
MOSFET	Metal Oxide Semiconductor Field Effect Transistor
MSK	Minimum Shift Keying
NF	Noise Figure
NMOS	N-type Metal Oxide Semiconductor

Acronym	Definition
NNR	Noise to Noise Ratio
NTF	Noise Transfer Function
OOSK	On-Off Shift Keying
OQPSK	Offset Quadrature Phase Shift Keying
OSR	Over-sampling Ratio
PAN	Personal Area Network
PC	Personal Computer
PER	Packet Error Rate
PGA	Programmable Gain Amplifier
PLL	Phase Locked Loop
PMOS	P-type Metal Oxide Semiconductor
PSD	Power Spectral Density
PSK	Phase Shift Keying
QAM	Quadrature Amplitude Modulation
QPSK	Quadrature Phase Shift Keying
RC	Resistor – Capacitor
RF	Radio Frequency
RFID	Radio Frequency Identification
RX	Receive(r)
SAW	Surface Acoustic Wave
SFDR	Spurious Free Dynamic Range
SIR	Signal to Interference Ratio
SNDR	Signal to Noise + Distortion Ratio
SNR	Signal to Noise Ratio
STF	Signal Transfer Function
TV	Television
TX	Transmit(ter)
UHF	Ultra High Frequency
VCO	Voltage Controlled Oscillator
VDD	Voltage Drain Drain (supply voltage in MOS circuits)
VLSI	Very Large Scale Integration
WAN	Wide Area Network
WCDMA	Wideband Code Division Multiple Access

Chapter 1

INTRODUCTION

The history of radio communications spans over 100 years. For most of that history, the information communicated was analog in nature and the radio devices were usually large and bulky. This meant that for most applications, the mobility of the radio transceiver was limited. Radio transmission today is used for a vast array of applications from traditional analog radios to digital radios like the modern Global Positioning System (GPS).

In the last few decades, the integrated circuit has changed the face of radio communications greatly. Today, most radio transceivers are used for transmission of digital signals. Today's transceivers are highly integrated, often including digital signal processors (DSPs) in a single integrated circuit. Advanced IC technologies have enabled the radio to operate with much less power, in addition to allowing most of the radio circuitry to be fully integrated. The result of all these advances (or the cause of them through demand) is that today's integrated radio transceivers are very small and can be powered by a small battery. This has allowed the development of many new devices such as wearable personal electronics, wireless computer peripherals and wireless control and automation devices.

Modern IC technology allows many digital functions to be placed on one chip in a small area and at very low power consumption. A microprocessor and a large amount of memory can be easily integrated into a portable MP3 player or an automation controller

without causing a significant battery drain. When a radio is added to the system to allow wireless functionality, even though it may have relatively low power itself, it still often consumes a large portion of the overall battery capacity. This in turn motivates the radio designer to continue to lower the power consumption through architecture choices and circuit design.

1.1 Thesis Motivation

The motivation of this thesis is centered on receiver architectures intended for low cost and low power applications. The thesis identifies the sub-sampling receiver with tuned band-pass anti-aliasing as a leading candidate for this application. The thesis gives the reader the tools required to evaluate and design a sub-sampling receiver through the use of theoretical analysis and a specific receiver implementation.

1.2 Thesis Outline

This thesis consists of six main sections. The first four chapters, including the introduction, provide the background material related to the proposed thesis topic. Chapter 2 discusses radio technology and wireless systems and specifically how they affect the requirements for analog radio design. Chapter 3 outlines radio receiver system requirements such as noise figure and interference performance. Chapter 5 looks at different receiver architectures and makes a comparison between the most common modern integrated receiver architectures.

The last three chapters deal specifically with sub-sampling receivers. Chapter 5 concerns the general architecture and requirements of the sub-sampling receiver. This section goes through calculation details and design tradeoffs specific to this architecture. Chapter 6

covers the topic of Q-enhanced filters. The noise and linearity of the filters is derived and an algorithm is presented for tuning the filters that will meet the high Q and frequency accuracy requirements of the sub-sampling receiver. Chapter 7 is a more specific example of an RF (radio frequency) sub-sampling receiver. This section outlines the block and level breakdown of a receiver designed for a 900 MHz RF input with example specifications.

1.3 Contributions

Several important contributions are made in this thesis. They are as follows:

- **Low Power Receiver architecture:** This thesis presents an architecture suitable for a general purpose digital radio receiver that is lower power than a comparable direct conversion or Low-IF receiver. The power savings compared to these traditional integrated architectures is expected to be 40-50%. An example receiver is implemented that consumes $\frac{1}{4}$ of the power of a Low-IF (Bluetooth) receiver.
- **Anti-aliasing requirements for efficient Sub-sampling Receiver implementation:** This thesis shows the derivation for the fundamental requirement of filtering in the sub-sampling receiver. The sub-sampling ratio, anti-alias filter Q and image rejection of the sub-sampling receiver are linked in a closed-form equation.
- **Q-enhanced filter for sub-sampling receiver anti-aliasing:** The Q-enhanced filter is presented as a key to providing the anti-aliasing in a sub-sampling receiver. The high Q required for sub-sampling is achieved through direct digital tuning

and the Q-enhanced filter noise and linearity are accommodated in the receiver architectural design.

- Circuit design contributions related to a sub-sampling receiver: A novel frequency control for LC filters is developed to allow precise frequency stepping in the anti-aliasing filter. A sub-sampling mixer circuit specifically designed with high RF isolation in hold-mode and optimized for use in a sub-sampling receiver.
- Direct digital tuning methodology: The concept of direct digital tuning of a filter enabling the use of a high Q band-pass anti-aliasing filter. The method uses the programmable sample frequency of a sub-sampling receiver to measure the center frequency and Q of a band-pass anti-aliasing filter directly.

The contributions listed above are tied together. The presented sub-sampling receiver architecture is unique in the fact that it makes use of a Q-enhanced LC filter with a high Q and samples the RF signal directly. The derivation for the performance of a sub-sampling receiver given anti-alias filter Q and sub-sampling ratio proves the requirement of the particular sub-sampling architecture presented. The example receiver presented is suitable for a general purpose digital radio transmitting at low to medium data rates and consumes less power than existing reported radios.

Chapter 2

RADIO TECHNOLOGY

2.1 Introduction

Prior to the discussion of digital radio receivers, an outline of radio technology in general will be made. Radio technology refers to the use of electromagnetic radiation to transfer information from one point to another. A system of devices that are used to transmit and receive data in this manner is often called a wireless system. The first section of this chapter will briefly discuss the different ways in which radio technology is used. This is done by dividing wireless systems up into different categories. By considering the possible categories of wireless systems one will have a better picture of the requirements that are placed on the devices used in that system. The second section of this chapter deals with the fundamentals of the radio environment and the third section of this chapter is an introduction to the modulation methods used for wireless communication.

2.2 Types of Wireless Systems

The number and variety of radio technologies and applications is vast. These systems can be divided in many different functional ways. Some possible ways to categorize different wireless systems are:

- **Analog versus Digital:** This simply refers to the signals being modulated on the radio carrier. Traditional systems used mostly analog modulation, while typical modern systems employ digital modulation techniques.

- **Broadcast versus Peer to Peer:** Broadcast wireless systems include traditional analog radio and TV as well as digital /satellite radio and TV, global positioning system (GPS) and others. Peer to Peer refers to links where the signals communicated between two devices are private. This includes cellular telephone, wireless LAN, and microwave wireless internet links. Peer to Peer can be further broken down into wide, local and personal area networks (WANs, LANs and PANs). WANs connect many different clients over a large area such as wireless internet connections. LANs can be small or large networks, generally owned or restricted to one entity such as an company, institution or residence. PANs refer to very small networks used for communicating between different personal electronic devices. This could include computer peripherals, personal entertainment devices or home automation and control. PAN networks are generally special purpose.
- **Industrial / Commercial Equipment versus Consumer Premise Equipment (CPE):** This differentiation actually refers to the devices in a particular network. CPE refers to the devices used by end user or consumers, for example cell phones. The commercial equipment is often immobile and made to service many customers, such as a cell-phone base-station. The distinction between these two groups is significant to the radio designer because on the consumer side power consumption and price are often very important criteria, whereas on the industrial / commercial side the device performance is paramount.

- **Mobile versus Stationary:** Some systems are intended for use with mobile applications and hence designed to be small, for example cell phones or wireless headphones. In other situations, the device is less mobile or entirely stationary as in the case of a digital satellite receiver. The difference between the demands of these devices is that the mobile application requires a battery and small form factor, whereas the stationary application can use line power. The mobile device therefore demands much lower power consumption and generally a higher level of integration.

Table 2-1: Summary of Wireless System Classifications

Wireless System Classification	Benefits to the Radio System Designer
Analog vs Digital	Defines the modulation signal and often indicates the complexity of the overall system
WAN / LAN / PAN	Mostly useful for higher protocol layers, but gives some indication of system complexity and importance of performance and power consumption
Industrial & Commercial vs CPE	Can be a major indicator of cost as a criteria
Mobile vs Stationary	A major indicator of the importance of power consumption

All of these possible categories are useful to determine the proper radio system for a given application. It is also necessary for the radio circuit designer to consider the specific intended application as it will determine architecture choices and tradeoffs.

Applications that demand low cost, small form factor and very low power consumption are becoming more and more common. The higher level of integration in today's radio devices has caused designers to consider new radio architectures to address the demands of these new applications and make use of modern integrated circuit technologies.

Modern IC technologies allow more digital processing within the same area and power constraints than what was previously possible. Radio designers today must create ways to relax the analog design specifications through the use of new architectures that increasingly employ digital processing and digital control of the analog radio.

The classification of radio systems above leads to the definition of the physical layer of the wireless system. The next few sections will discuss the design of the entire wireless system from a high level. This high level design is done with the uses and classification of the wireless system in mind.

2.3 Link Budget

The term “Link Budget” is often given to a single communication link in a wireless system when considering the signal path between the transmitter and the receiver. The link budget is an account of all the elements that will degrade the signal and ultimately determine the maximum distance of separation possible between two devices. Table 2-2 shows a simple example of a link budget. The first section details the operating frequency as well as the power and sensitivity characteristics of the receiver and transmitter. The next section accounts for the antennas used by the devices and allows some of the link budget to account for signal loss due to multi-path fading. The remainder of the budget is then assumed to be available for propagation loss, which ultimately determines the maximum distance between devices. The following equation calculates the distance (d) between devices versus the loss (L) in the link:

$$L = 20 \log \left(\frac{4\pi d}{\lambda} \right) \quad (2.1)$$

where λ is the wavelength of the radio signal. Note that equation 2.1 is the loss in free space in dB. Most real world scenarios include other effects and more complex equations model these situations.

Table 2-2: Example Link Budget

Operating Frequency	1	GHz
Transmit Power	10	dBm
Receive Sensitivity	-90	dBm
Available Link Budget	100	dB
Transmit Antenna Gain	2	dB
Receive Antenna Gain	-3	dB
Multi-path Fading Budget	-30	dB
Remaining Link Budget	69	dB
Maximum Distance	67.3	m

The type of network or wireless system as described above usually gives us an indication of what type of environment the radio system will be working in. The environment plays a big role in calculating the link budget and therefore determining the requirements of the devices operating in a wireless system. The following list shows some of the things that affect the radio environment:

- Maximum range or distance between devices in a wireless system
- Mediums and obstructions through which the radio signals must travel
- Multiple signal pathways through reflections which can cause both constructive and destructive interference of signals (often just referred to as “Multi-path”)

- Frequency band available for radio transmission either licensed or multi-user unlicensed bands
- Number of devices that can potentially interfere by transmitting in the same frequency band at the same time

These parameters will often affect how the specifications for the radio transceiver are derived. For example, a cellular telephone system will likely have to take into account many different multi-path situations where a cell-phone user could either be in a city location where there are many multi-path signals or out in the open where only direct line of sight transmission is possible. On the other hand, a stationary system such as a wireless internet connection to the home has a stationary line of sight pathway to an immobile transmission tower. In the latter case a directional antenna can be used and the system will likely require less link budget to achieve the same distance.

2.4 Modulation

Modulation in radio transceivers can be broadly broken down between analog and digital modulation. In analog modulation, the signal being transmitted is analog in nature, as in the case of traditional television audio and video signals. Digital modulation, as the name suggests, deals with the transmission of digital signals. Many, if not most, modern communication systems are digital, even if the desired signal being transmitted is analog in nature. A good example is cell phone signals which are almost exclusively transmitted as digitized signals today. The digital system can utilize sophisticated and aggressive noise suppression and interference handling algorithms can be used to make the transmission more robust than an analog system. In a digital communication system,

signal processing can be used to verify the correct transmission of a signal and make appropriate adjustments. Although there are several different analog modulation methods, the remainder of this discussion will deal specifically with digital modulation. Note that this discussion of modulation types only introduces some of the most commonly used methods. A complete discussion of the mathematics and structure behind the modulation is outside the scope of this thesis [22].

Each modulation technique includes a set of symbols that describes it. The number of symbols used for the modulation indicates how much information is stored in each symbol. The simplest modulation forms are binary and have 2 symbols (one representing 0 and the other representing 1). Higher order modulations have 2^m different possible symbols, where m is the number of bits represented by each symbol. This is also referred to as M-ary modulation. For example, 16 QAM (Quadrature Amplitude Modulation) has 16 symbols and each symbol represents 4 bits. The symbol rate in this example would be $\frac{1}{4}$ of the bit rate.

Each modulation technique also works along a defined number of dimensions. This can be only 1 dimension (for example, the amplitude of the carrier in Amplitude Shift Keying - ASK) or it can be many dimensions as in M-ary Frequency Shift Keying (FSK). The symbols of a modulated signal can be represented in a constellation graph using the dimensions of the modulation, provided there are 3 or fewer dimensions. Figure 2-1 shows two examples of different symbol constellations. On the left is a binary modulation where each symbol represents only 1 bit and there is only 1 dimension. On the right is a more complicated constellation with 3 dimensions (each axis drawn

represents one orthogonal base for the symbols). In this case 8 symbols are used, creating an M-ary modulation with $M=8$. For this example each symbol would represent 3 bits.

Figure 2-1: Examples of Signal Constellations

The shape and nature of a signal constellation will determine what the minimum requirements for a receiver are. The distance between symbols will determine what the minimum Signal to Noise Ratio (SNR) is required to receive a signal with a pre-determined Bit Error Rate (BER). Also, the phase difference between signals will give an upper bound on how much phase noise can be introduced by a receiver before the signal is corrupted and the BER level goes up.

2.4.1 Spread-Spectrum

Spread-spectrum modulation schemes can improve the interference robustness and extend the range of a particular link without increasing the transmission power. This improvement is accomplished by increasing the amount of bandwidth used to transmit the signal without increasing the data rate of transmission. The same signal is essentially

transmitted over more bandwidth in a redundant manner. If the redundancy is done correctly, the demodulated signal will suppress any noise or interference that was received with it. The actual mechanisms for accomplishing this are beyond the scope of this thesis, but the results are important when considering and deriving system requirements for the receiver.

In direct sequence spread spectrum (DSSS), bits are translated into a pseudo-random series of symbols. Redundancy is placed in these sequences since there are many more combinations of symbol sequences than the bits that they represent. When demodulated, the signals will be multiplied by the symbol sequence again. Because the pseudo-random sequence is not correlated with the noise and interference, these signals are suppressed with respect to the desired signal which is correlated. The result is that the required SNR ($\text{SNR}_{\min}$) for the receiver is reduced and the receiver analog bandwidth is increased. The downside is that the symbol rate required to obtain the same bit rate is much higher.

Frequency hopping spread spectrum (FHSS) keeps the symbol and bit rates the same as described in the modulation sections above, but changes the frequency of the carrier in a random fashion so that if there is interference degradation at one frequency, the signal will still be received at other frequencies. As with DSSS, the overall transmitted bandwidth is greater than a non-spread signal, but the benefit is improved interference robustness.

2.4.2 Concerns for the Analog Radio Designer

The discussion of modulation and spread spectrum techniques introduced above only scratches the surface of a very large topic. The analog radio front-end designer does

however require some knowledge of the operation and tradeoffs of these techniques.

Often there is a tradeoff between the modulation and the complexity and performance of the analog radio.

The bandwidth efficiency of a digital system is often used to determine how well a band of frequencies is being used by a wireless system. The efficiency is the number of bits/s per unit of bandwidth (bits/s/Hz). There exists a fundamental limit to this efficiency, called the Shannon Limit [22]. As this limit is approached, the complexity of the system typically increases, both in the analog and digital portion of the transceiver. As a result, a compromise is often made between the bandwidth efficiency and the specifications of the radio. For example in order to relax the interference requirements of a receiver, the channel separation of a radio system can be increased. This also allows for the use of modulation techniques that have higher side-lobes. Ultimately, the specifications of a wireless system should take into account the realities of the circuit implementations required to achieve functionality.

2.5 Conclusions

This chapter has introduced some of the fundamental concepts of the wireless system.

The analog receiver designer does not often consider the high level issues introduced here, but a general understanding of the type of wireless system, link budget and modulation methodology will help in understanding the tradeoffs and derivations involved in setting the system requirements for the receiver.

Chapter 3

RECEIVER SYSTEM REQUIREMENTS

3.1 Introduction

Before looking at different receiver architectures, it makes sense to first discuss the requirements of a receiver and the metrics used to evaluate those requirements. These requirements will influence the receiver architecture choices. Each of the parameters described in this section are ultimately derived from the physical layer specifications for which receiver is being designed. For example, the specific modulation method and signal bandwidth places a requirement on the minimum SNR required to demodulate a signal. This, combined with the transmission distance, transmission environment, etc. is used to determine the required sensitivity of a receiver. It is useful for the radio designer to understand the modulation and coding scheme used for a given radio specification (described briefly in Chapter 2) and how these will affect the receiver performance.

Other constraints can also influence the specifications and sometimes conflict in setting the specifications of a receiver. The maximum allowable power consumption in a desired application often conflicts with the performance of certain circuit blocks, for example, phase noise and signal distortion. The choice of receiver architecture influences how well the constraints of the receiver are optimized.

3.2 Power Consumption

The nature of many wireless end-user products is that they are battery powered. The lifetime of the battery powered device that uses a wireless link to communicate to a network is also often limited by the power consumption of the radio. To determine the power consumption requirement, the usage profile of a product must be analyzed.

To illustrate this point consider a wireless thermostat that would typically be placed in a room of the house that was not originally wired to communication with the furnace / air conditioning system. Thus it will be battery powered and will communicate to the heating and cooling control system via a wireless link. A system level block diagram for the interface of the thermostat might look like that given in Figure 3-1.

The battery supply voltage is likely not a design variable, thus one is most interested in the current consumption, as opposed to the power consumption of the device. Table 3-1 shows the breakdown of current consumption in the thermostat interface. Even using a low power Bluetooth radio, the power consumption of the receiver is typically dominant. In fact, if the device is expected to last 3 years without battery replacement, the battery requirements become unreasonable. An “AA” size alkaline battery has roughly 2200 mA·hr of capacity with a low battery voltage of 0.9 V. To maintain a supply above 1.5 Volts and provide the necessary capacity, 6 AA size batteries would be required. Alternatively a DC-DC converter could be included and then the likely requirement would be at least 3 AA batteries. It is more reasonable that a specialty high capacity lithium or lithium-ion battery would be necessary and would be expensive and still quite large.

Figure 3-1: Block Diagram of Wireless Thermostat Interface

Table 3-1: Breakdown of Current Consumption for Thermostat

Block		Current Consumption	Duty Cycle	Average Current Consumption
Radio	RX Mode	30 mA ¹	0.5%	150 μ A
	TX Mode	35 mA ¹	0.01%	3.5 μ A
Micro-processor	Active Mode	30 μ A ²	1%	0.3 μ A
	Watch Mode	1.2 μ A ²	99%	1.2 μ A
ADC		1.5 mA ²	0.05%	1 μ A
LCD		5 μ A ²	100%	5 μ A
Buttons / Thermostat [†]		0	N/A	0
Total				161 μ A
Required battery size for 3 year operation				4126 mA·hr

[†] Estimated from CSR BlueCore4 Bluetooth Datasheet

² Taken from Hitachi H8/38004 Microcomputer Data sheet

[†] The thermostat draws current from the ADC

It is obvious from the above analysis that the reduction of the RX mode radio power consumption is a priority for these types of applications as its power consumption accounts for 93% of the system total. A large battery is required to meet the specifications of this system.

For form factor and cost reasons, smaller batteries would make a product like this much more useful. Small lithium batteries would provide a low-cost and long shelf-life (> 5 years) and are therefore a desirable target. The common CR-2032 is a coin cell used commonly for PC clock and memory backup with 220 mA·hr of capacity. If the system could utilize two of these batteries, the cost and form factor would make it much more favourable. This would require an RX current of 10 mA or less to achieve the 2 year operating target.

The same analysis could be done for many different mobile systems with different scale factors involved, but the common end conclusion is that the radio power consumption is very significant and the receiver power consumption is typically one of the most important contributors to this power consumption.

3.3 Sensitivity / Noise Figure

Sensitivity is often the first specification considered when analyzing receiver performance, even though it is not usually the most important. It simply refers to the minimum signal level that is detectable by the receiver. It gives no indication of how

large or close interference can be or how much non-ideality in the transmitted signal can be handled by the receiver.

Noise Figure is often used in analog RF receivers to separate the performance of the analog portion of the receiver from the chosen modulation scheme and digital baseband implementation. The sensitivity (P_{sens}) can be derived from the noise figure as follows [9]:

$$P_{sens} = P_{RS} + NF + SNR_{min} + 10\log(B) \quad (3.1)$$

Where all quantities above are in dB. P_{RS} refers to the noise in the source and assuming the receiver is properly matched to the antenna, this value is only dependent on the temperature ($P_{RS}=kT$ or -174 dBm/Hz at room temperature). The sensitivity equation becomes:

$$P_{sens} = -174 + NF + SNR_{min} + 10\log(B) \quad (3.2)$$

B is the bandwidth of the signal in the system, or more specifically, the noise bandwidth of the channel. SNR_{min} refers to the minimum Signal to Noise ratio (SNR) required before a signal can be received. This value is very specific to the radio system in question. It requires knowledge of the modulation type (as described in 2.3), the acceptable bit error rate, and any impairments in the digital baseband demodulator. As an example, coherent QPSK requires a minimum SNR of 8.3 dB to achieve a bit error rate (BER) of less than 0.01% [22]. This assumes a perfect demodulator. In practice, margin will need to be built in to allow for non-idealities in the demodulation such as carrier offset, symbol recovery and phase noise. A more realistic number would be 11 dB.

Spread-spectrum systems sometimes complicate this relationship by introducing another level of complexity. The physical bandwidth of the analog portion of the receiver is generally higher than the actual data bandwidth but a lower signal to noise ratio is required. This can be accounted for in the sensitivity equation by either reducing the $\text{SNR}_{\min}$ value and leaving the bandwidth (B) as the true analog bandwidth used or using the same $\text{SNR}_{\min}$ as a non-spread system, but adjusting the bandwidth to reflect the true bandwidth of the demodulated signal. The first option is usually better for analyzing the analog front-end of the receiver since the use of the true analog bandwidth is required for most of the analog signal analysis.

3.4 Interference Performance

There are many criteria used to describe the interference performance of a receiver in one way or another. In each case an attempt is being made to put a specification in place that will guarantee the functionality of the receiver given certain specific interferences.

3.4.1 Selectivity

Selectivity refers to the receiver's ability to filter out unwanted signals. It is particularly important when connecting one portion of a receiver to another (for example, an analog front-end to a digital demodulator). When the signal is passed from one block to another, a minimum attenuation at certain offsets is expected. Selectivity is often reported for the analog front-end of a receiver since it is a specification required by the analog to digital converter (ADC) and the digital baseband circuitry. When the entire receiver is considered, including the analog front-end, ADC, digital filtering and demodulation circuits, the selectivity is no longer important. More important is the idea of

“interference selectivity”, which is simply a measure of how large an interfering signal would have to be in order to disrupt reception of a desired signal, and is generally described as a ratio of the interfering signal to the desired signal. For example, if an interferer at 20 MHz offset with a power of -20 dBm causes reception of a desired signal at -65 dBm to just fail, then the Interference selectivity is said to be 45 dB.

3.4.2 Inter-modulation Distortion

Inter-modulation distortion refers to the effect of non-linearities in the receiver causing harmonics of interference signals to mix into the channel of interest. Specifically, this happens with odd-order harmonics where two or more frequencies mix together and land on top of the desired channel. To illustrate this consider a transfer function that, although ideally is completely linear, in reality has non-linear components:

$$v_{out}(t) = a_1 v_{in}(t) + a_2 v_{in}^2(t) + a_3 v_{in}^3(t) + \dots \quad (3.3)$$

The term a_1 is the linear gain of the block and the terms a_2 , a_3 , etc. are the non-linear coefficients that can cause inter-modulation. The even terms like a_2 do not produce tones in band and if differential circuits are used are generally much smaller. The odd terms, especially the 3rd order terms are what will be of concern because some of the products will fall in-band and can cause interference. To see how this works, consider two distinct tones entering a receiver block, defined as:

$$v_{in}(t) = V_1 \cos(2\pi f_1 t) + V_2 \cos(2\pi f_2 t) \quad (3.4)$$

When this input signal is cubed, the resulting output will have frequencies of $3f_1$, $2f_1 - f_2$, $2f_2 - f_1$ and $3f_2$. The frequencies of $3f_1$ and $3f_2$ will be very high and not of concern, but

the other terms will fall near f_1 and f_2 . In fact, they fall at the same distance away on either side of the two tones as the two tones are separated, as shown in Figure 3-2. This calculation is expanded upon further in section 7.3.

Figure 3-2: Inter-modulation Resulting from Non-linearities

The 3rd order intercept point is often used to evaluate the non-linear performance of a block. It is defined as the point where the inter-modulation products ($2f_1-f_2$ and $2f_2-f_1$) are equal to the fundamental output at f_1 and f_2 when two equal amplitude tones are applied to a block. This point is a purely theoretical idea since the gain compresses long before this point is actually reached, but it does give a ratio that can be used to find the overall inter-modulation products. Also, an overall IP3 can be calculated between cascaded blocks (see Chapter 7). The higher the IP3, the lower the inter-modulation distortion will be for a given block or group of blocks. When designing the receiver architecture, the IP3 can be used as a specification for each sub-block in the receiver.

When considering the IP3 of several blocks in succession, it is also important to consider the filtering effect of each block. The interfering signal causing the inter-modulation distortion is usually offset from the desired signal and depending upon the receiver architecture, could be filtered to some degree by blocks in the receiver. Figure 3-3 shows the resulting IP3 in two different situations. In the first case (a), two blocks are in series

followed by a filter that has an attenuation of 20 dB at a 20 MHz offset. The resulting IP3 is -30 dBm. In the second case (b), the filter is moved in between the first and second blocks. Although the in-channel IP3 will be the same as in case (a), if one is really concerned about interfering signals at offsets of 20 MHz and more, the filter improves the IP3 substantially. In case (a) the inter-modulation was almost entirely caused in Block 2 both because its IP3 is lower and because of the gain of Block 1. In case (b) the interference at 20 MHz offset is attenuated by 20 dB before it sees the higher non-linearity of Block 2. In this case Blocks 1 and 2 are equal contributors to the overall IP3, yielding a total IP3 of -13 dBm.

Figure 3-3: Effect of Filtering on Effective IP3 at a Frequency Offset

3.4.3 Dynamic Range

“Dynamic Range” as a term is not specific enough since it can be defined in many different ways. In general, it refers to the range of signal levels that the receiver accepts in its functional limits. The functional limits themselves are defined by the modulation type and system in use. Ultimately, a certain minimum SNDR (Signal to Noise plus Distortion Ratio) is required and that requirement is derived from the digital modulation

scheme (see section 2.3). Some examples of how dynamic range is defined are as follows:

- Maximum detectable signal versus the minimum detectable signal, at any time. This dynamic range refers exclusively to the signals desired by the receiver. This is an indication of the breadth of the automatic gain control (AGC) range. It can also be affected by the modulation used as some types of modulation can handle large amounts of compression, thereby allowing a much higher maximum signal, while others require high linearity.
- Dynamic range in the presence of interference. This is the same as “Interference Selectivity” described above. The dynamic range is defined as the ratio of the maximum interfering signal to the desired signal, where the receiver still functions. This is sometimes referred to as instantaneous dynamic range.
- Spurious Free Dynamic Range (SFDR) refers to the dynamic range in the presence of one or more interferences. In some circuits, the interference frequencies are unrelated to the input frequencies, as is sometimes the case with data converters. For radio frequency designs, the SFDR is commonly used to describe the effects of IM3 distortion on the signal. It differs from the interference rejection in that two interfering signals are required at equally spaced offsets from the desired signal to identify the worst case situation. As with IP3, it can be defined as in-band, ignoring filtering effects or it can be defined as dependent on frequency offset if the receiver has significant in-band filtering.

The SFDR is most often used for the purpose of evaluating or comparing the analog portion of receivers, even though the other definitions can be of equal importance to the fully integrated receiver. The SFDR is often considered a worst-case instantaneous dynamic range and so although it can be pessimistic, it is used to evaluate analog receiver blocks or the entire receiver at fixed gain settings. The AGC dynamic range is of more importance to the entire receiver including the digital baseband portion, as the digital portion and method of choice for demodulation has an effect on this range.

3.4.4 Image Rejection

Image rejection is a special case of interference caused by the mixing or frequency conversion of the receiver. The action of mixing to convert one frequency to another causes more frequencies than the desired frequency to be down-converted. This can be seen in the mathematical expression for the mixing function:

$$\cos(\omega_1 t) \cos(\omega_2 t) = \frac{1}{2} \cos(\omega_1 t - \omega_2 t) + \frac{1}{2} \cos(\omega_1 t + \omega_2 t) \quad (3.5)$$

The result from this equation gives two frequency components but after Low Pass Filtering (LPF) only the first term, the difference between the two frequencies, remains. Because the cosine is a real valued signal, it has both negative and positive frequency components (and the sign of the frequency is not important). Hence:

$$LPF(\cos(\omega_1 t) \cos(\omega_2 t)) = \frac{1}{2} \cos(|\omega_1 t - \omega_2 t|) \quad (3.6)$$

If one of the frequencies is fixed as the mixing frequency (called the LO or local oscillator), there are two possible frequencies that give the same IF (Intermediate

Frequency). If a signal exists at the undesired frequency, it will interfere with the desired signal, possibly corrupting it.

Image rejection refers to the amount that the image frequency is attenuated before it reaches the mixing stage that would cause this interference. It is slightly different than the interference selectivity quantity in that the minimum SNR required has not been taken into account. The interference selectivity at the image frequency can be defined as:

$$\text{Interference Selectivity @ Image Frequency} = IR - SNR_{\min} \quad (3.7)$$

where IR is the Image Rejection in dB and $SNR_{\min}$ is the minimum SNR required to receive a signal. As with any other Interference Selectivity, it is generally defined at levels well above the noise floor (P_{sens}) of the receiver, otherwise the noise will add to the interference and the result will be worse.

3.5 EVM

Signal distortion can take on a number of different forms. In a radio system that employs phase modulation, a common method for quantifying this distortion is the Error Vector Magnitude or EVM [22]. The EVM is a measure of the deviation of a digitally modulated signal from its ideal. The error vector itself is defined as the difference between the actual received and ideal symbol and is shown in Figure 3-4.

Figure 3-4: Example of Error Vector

Each symbol received will generally have a different associated error vector. The EVM is the rms value of the magnitude of these vectors and can be expressed as:

$$EVM = \frac{\sqrt{\frac{1}{N} \sum_{k=1}^N [(I_{r,k} - I_v)^2 + (Q_{r,k} - Q_v)^2]}}{|\bar{v}|} \quad (3.8)$$

Where the ideal symbol for each received symbol is defined as:

$$\bar{v} = I_v + jQ_v \quad (3.9)$$

The evm is normally expressed as a percentage.

EVM can be caused by many different receiver non-idealities, including the addition of noise. It is, in effect, a measure of the overall quality of a signal once processed by the receiver. Even thermal noise added in the receive channel will increase EVM. In the case of noise exclusively causing the EVM, the EVM is essentially the ratio of the noise to the signal. To that extent, the EVM is the inverse of the SNR. Normally, EVM is measured for the purpose of interferences other than thermal noise. The top graph in Figure 3-5 shows how the SNR increases as the EVM decreases when the input power to

a receiver is increasing. In this example, the thermal noise floor is -90 dBm (and can be seen as the point where the SNR is close to 0 dB) and no other interference sources are considered. In the lower graph of Figure 3-5, the EVM is plotted versus Signal to Interference Ratio (SIR). The SIR is a measure of all the interference sources of a signal, including noise. In this example the EVM at very high input powers is a constant 8%. This can be caused by some signal-independent effect such as phase noise. At low input powers, the thermal noise dominates as an interference source and the EVM is almost entirely defined by the thermal noise. At higher input powers, the SNR stops increasing because it is limited by the minimum EVM of the system. As can be seen from this relationship, the EVM is most useful in describing distortion in the receiver that is proportional to the signal, and it is normally quoted as a specification that ignores thermal noise or assumes a sufficiently large input signal.

Figure 3-5: EVM with only Noise (above), EVM with Noise and Interference (below)

EVM is caused by any change in the symbol pattern that cannot be adjusted for in the receiver, because typically these changes are random in nature.

3.6 Phase Noise

Phase noise is always an important design criterion for a radio receiver. It refers to the noise that is introduced in the mixing and / or sampling action of the receiver due to the random changes in the phase or period of the reference. It is considered separately from other AGN (Additive Gaussian Noise) sources because the noise is essentially modulating the incoming signal. Phase noise has two main effects in a receiver that will be discussed below.

3.6.1 Self-Interference

Phase Noise causes self-interference whereby the incoming signal is modulated by the noise in the reference signal and is seen as interference. This effect is an important contributor to the EVM of the receiver. Figure 3-6 shows the difference between the effect of AGN noise added to the signal (on the left) to that of phase noise and additive noise. The phase noise only affects the phase of the signal, but reduces the margin of the signal overall. An important difference between the effect of phase noise and the effect of additive noise is that the phase noise amplitude increases as the signal level increases (in other words, the ratio of signal to phase noise is constant). For this reason it is important to have a large enough margin with phase noise because it contributes to the EVM at all signal levels. An increase in phase noise past this margin will cause the receiver to stop functioning. An increase in additive noise on the other hand would only have the effect of reducing the receiver's sensitivity.

Figure 3-6: Phase Noise Effect on Symbol

3.6.2 Reciprocal Mixing

The second effect of phase noise is in the presence of a large interfering signal. The interfering signal is mixed with the reference signal in the receiver and the side “skirts” created cause an interference with the desired signal. This is sometimes referred to as Reciprocal Mixing. The reciprocal mixing effect may or may not matter, depending on whether the interference at that level has limited the interfering signal level through inter-modulation distortion or gain compression, as discussed in more detail in section 3.4.

Phase noise in radio circuits is usually expressed in the frequency domain. In general, the power spectral density (PSD) of the phase noise is inversely proportional to the square of the offset in frequency and the phase noise of circuit blocks or systems are specified at certain offsets from the carrier. From this information, one can usually predict the effect of reciprocal mixing. Figure 3-7 shows how the phase noise of the reference used for mixing “spreads” the energy of the interfering signal such that some of it appears on top of the desired signal (shaded).

Figure 3-7: Example of Reciprocal Mixing in a Receiver

The required interference performance can often drive the specifications for phase noise. In practice, self-interference and reciprocal mixing are both used to derive phase noise requirements because they refer to phase noise at very different offsets from the carrier. In the case of self-interference, the frequency offset is within a channel and the specifications usually have the most effect on the specs for the PLL used in the receiver to lock the oscillator used for down conversion. Reciprocal Mixing on the other hand

deals with frequency offsets that are often outside of the bandwidth of the PLL and thus drive the specifications of the local oscillator itself.

3.7 Inter-Symbol Interference

Inter-Symbol Interference or ISI results from filtering in the receiver. In an ideal system, a brick-wall filter would be used in the receiver to filter the channel of interest, minimizing the required bandwidth. This type of filter has a significant impulse response on a received symbol that affects the time domain response far outside of the symbol's designated timing window. If the timing of the receiver were perfect, the zero crossings of all interfering symbols would be zero right at the same moment the receiver needs to sample the current symbol. In practice the timing is far from exact and a brick-wall filter would cause excessive interference (ISI) from one symbol to the next [21].

The implications for filter design are usually only considered in the digital section of a modern receiver where the bandwidths are tight enough to cause problems with ISI. However, in some cases where low bandwidth analog filtering is implemented to reduce the required dynamic range on the analog portion of the receiver, both the analog and digital filtering must be considered in evaluating the ISI.

3.8 Overall Performance budgets

At any one time the receiver will suffer from several of the performance degradations described above. It is not possible to combine and test a particular performance criteria for all types of interference at one time, but it is important to consider the effects of all of the degradations that will take place in a given situation. For example, when interference is present at the receiver's image, it is not sufficient to make sure that the SNDR is

greater than $\text{SNR}_{\min}$ while considering only the interfering signal level as the distortion. In fact, the budget allowed for phase noise, ISI and white noise needs to be considered too. In the case of $P_{\min}$, the minimum receiver input power is calculated with no interference present at all. In this situation, it is not possible to have any interference at all or the bit error rate (BER) will go up. Figure 3-8 shows an example where noise, distortion and interference have been plotted together to give the total range of possible received symbols. It can be expected that symbols from real receivers, if plotted on a complex graph, would have a scatter plot similar to the bounds shown here. Note that in this example, there is very little margin left before some symbols cross the decision threshold. Past this point they would be detected incorrectly and would create bit or symbol errors.

Figure 3-8: Combined Effect of Distortion, Noise and Interference

3.9 Conclusions

In this chapter the primary metrics used to evaluate the performance of a receiver have been defined. The different effects are inter-related and very often lead to design tradeoffs in the receiver. The circuit architecture chosen for a receiver plays a very important role in determining how these requirements can be met and where the tradeoffs are made. Inevitably, a particular architecture will provide superior performance on some of these metrics, while performing poorly in others. The goal of the receiver designer in this case is to match the requirements with the most suitable architecture.

In mobile and battery powered wireless systems, power consumption is of paramount concern. Although a certain minimum level of performance is necessary in other metrics, a low enough power consumption often determines whether a certain application is even viable.

Chapter 4

RECEIVER ARCHITECTURES

4.1 Introduction

In this chapter different receiver architectures are explored and compared in terms of their potential for low power consumption. The method for frequency conversion and filtering in the receiver is what differentiates one architecture from another. This chapter will first describe the architectures most often used in integrated receivers. Next an overview of recent example receivers will be compared to see how power consumption compares as well as other relevant performance criteria. The difficulty in comparing one receiver implementation to another is that they are rarely designed with the same performance targets in mind so that differences in power consumption can be difficult to attribute solely to an architectural choice. Also, it is often the case that one architecture is suited to a certain set of bandwidth (data rate) and performance criteria, yet a different architecture would be chosen if the criteria were changed. The sub-sampling receiver presented in this thesis fits within a particular range of these criteria. In the final portion of this chapter the low power potential of the sub-sampling receiver will be shown by comparing it with published receivers that also fit in this range.

4.2 Integrated Receiver Architecture Descriptions

4.2.1 Generalized Receiver

The primary function of a receiver is to convert radio frequency (RF) signals to their original format (i.e. demodulation). In the context used here, the receiver is limited to the analog portion of a digital receiver. As such, the modern analog portion of a digital radio receiver rarely includes the actual demodulation. For the analog circuit, the most important function of the receiver is to convert the frequency of the signal from the incoming RF frequency to a baseband signal. The baseband signal is still modulated, but is centered on DC (i.e. has no carrier frequency). Secondary functions of the analog receiver include signal amplification, filtering unwanted signals, and analog to digital conversion. It is useful to first discuss the general building blocks and subsystems of which most receivers are comprised.

Figure 4-1 shows two main subsystems for a receiver architecture. The first subsystem is the frequency conversion or “heterodyne” stage. The word heterodyne itself means “two different frequencies that are combined to produce two new frequencies, the sum and difference of the original frequencies”. This is commonly referred to as “mixing” the two signals. This stage serves to convert the signal frequency from a high frequency to a lower frequency using a local oscillator (4) and a mixer (3). The incoming frequency is called the radio frequency or RF and the output of the heterodyne stage is called the intermediate frequency or IF. A gain stage (1) is commonly included to improve the noise performance of the system and is often called the Low Noise Amplifier (LNA). A filter (1) is sometimes included to remove interference from the input signal and to

possibly remove mixing products if the sub-system follows a previous frequency conversion. The second sub-system is the analog baseband subsystem and is almost always present in a digital radio receiver. The main block of this system is the Analog to Digital Conversion (ADC) (7). A gain block is included (6) that is often designed as a programmable gain. This serves to control the signal level at the input of the ADC, thereby reducing the required dynamic range of the ADC. The baseband filtering (5) is used to removing mixing products from the previous frequency conversion and also to do some or all of the channel selection. This reduces the required dynamic range of the ADC. These subsystems are combined to create a receiver in various combinations. In general terms, a receiver is comprised of one or more Frequency Conversion Subsystem and one Analog Baseband Subsystem, as shown in Figure 4-2.

Figure 4-1: General Receiver Subsystems

Figure 4-2: General Receiver Architecture

The choice of frequencies for the conversion in the heterodyne stage or frequency conversion subsystem has a number of important tradeoffs [9]. The heterodyne stage

introduces an image frequency as described in section 3.4.4. In order to achieve sufficient image rejection, the signal must be filtered before the frequency conversion. This filter is sometimes called the image reject filter. Given the high frequency of the RF signal, the bandwidth of the image reject filter will generally be large relative to the bandwidth of the signal (filter Quality Factor or Q is limited to small values for practical reasons and that limits the realizable bandwidth of the image rejection filter). The result is that in order to get a high image rejection, it would be preferable to have the LO frequency far from the RF frequency.

On the other hand, the intermediate frequency (IF) is determined by the difference of the LO and RF frequency. At the IF, the channel selection filter needs to filter out only the channel of interest. If the IF frequency is high, the required Q of this filter is therefore high. As a result, the IF frequency should to be as low as possible in order to make this filter design more reasonable. This conflicts with the desire to keep the image frequency as far from the channel frequency as possible.

4.2.2 Super-heterodyne

The super-heterodyne receiver alleviates the problem described above by incorporating two separate frequency conversion stages. In the super-heterodyne architecture the first IF is still at a relatively high frequency, making the image rejection of the first mixing stage easy to obtain. The band-pass filter of the IF stage does not need to have complete channel selection so its Q requirements are relaxed. The second mixing stage converts the signal to the second IF or baseband where the channel selection is done. Note that there is now a second image created by this second mixing stage so the filtering in the

first image will need to provide sufficient image rejection. The overall architecture is shown in Figure 4-3 [11].

Figure 4-3: Super-Heterodyne Receiver

The super-heterodyne receiver has been in wide usage since its invention by Armstrong during World War I [36]. The frequency plan of the traditional super-heterodyne receiver normally involves a fixed first IF frequency and fixed second LO. Traditionally, to implement sufficient filtering at the first IF a separate ceramic or SAW (surface acoustic wave) [12] filter was used, since integrated technology did not allow integration of a filter. Because the filter was not tunable, each channel which was received needed to be moved to the same IF frequency. This dictated that two separate synthesizers were required to generate LO₁ and LO₂. In modern variations of the super-heterodyne receiver, the first filter can be mostly integrated and tunable, allowing the first IF frequency to follow the RF channels. This enables a single synthesizer to generate both LO frequencies if they are chosen correctly. This configuration is sometimes referred to as a “walking IF” architecture [24].

4.2.3 Direct Conversion

The Direct Conversion Receiver (DCR) omits the first IF stage of the super-heterodyne receiver above and converts the signal directly to DC, meaning there is no image frequency in this frequency conversion stage. The frequency conversion must be done in a complex manner to differentiate between the I and Q phases of the incoming signal. The channel selection is done using low-pass filters in the baseband. The general DCR is shown in Figure 4-4.

Figure 4-4: Direct Conversion Receiver

There are some obvious advantages to the direct conversion receiver. Having no image frequency and a single mixing stage allows this receiver to have fewer RF components. The architecture also accommodates the integration of the low frequency channel filtering so fewer off-chip components are required.

The difficulty in this architecture arises from the channel filtering design. Normally there is little or no other filtering in front of the channel filters. There is also not much gain at RF, otherwise the IP3 / linearity would suffer greatly. Thus, the noise in the filters as well as the linearity of the first stages of filtering are very important design criteria. It is

very difficult to get low noise performance at DC, particularly working with MOSFET transistors due to the $1/f$ noise [25]. Also, the DC offset of the filter must be very low otherwise it will saturate the gain required by the baseband. This typically requires sophisticated trimming and cancellation schemes.

4.2.4 Low-IF

The low-IF receiver can be seen as a slight modification to the direct conversion receiver in that the signals is converted to a frequency just above DC to alleviate some of the problems described above [26]. This is also sometimes referred to as an image reject receiver architecture [9] and is shown in Figure 4-5. An image frequency is created by the conversion to the low-IF, but the architecture maintains an I and Q path so that the image frequency can be cancelled using the correct phase shifting and signal addition as shown. In this operation, the desired signal frequency is adding constructively while the image frequency is subtracted out. This architecture can have the filters for channel selection implemented on-chip as with the DCR and reduces the effects of $1/f$ noise and DC offset.

Figure 4-5: Low IF Receiver

The problem with the Low-IF receiver is that in order for the image rejection to work, the I and Q paths have to match very well in amplitude and phase shift. This is typically difficult to accomplish and image rejections of more than 20 dB is hard to achieve. The image rejection could be further improved by adding a notch pre-filter before the mixing stage, but the filtering of the image is often limited because it is still at RF frequencies [27].

4.2.5 Oscillator Receivers

The architectures described thus far have used oscillators as a source of mixing signals (local oscillators or LOs). “Oscillator receivers” refers to receivers that use an oscillating circuit in the signal path as a source of signal gain and possibly phase locking. This group includes super-regenerative receivers, injection locking and PLL (phase locked loop) based receivers.

This class of receivers can often achieve very low power consumption because the oscillator in the signal path provides a high gain to the received signal with a relatively low power draw. Also, an RF synthesizer is not usually required or is composed of the signal path oscillator itself. On the other hand, the modulation scheme must be constant amplitude such as OOK or FSK as the input signal amplitude is clipped or hard-limited by the oscillation circuit. When considering process and temperature variations, there are practical issues of implementing an oscillator that has a tight control of the injection locking bandwidth during reception that is both sufficiently wide to allow the signal to be amplified, yet narrow enough to reject adjacent channels.

In the case of super-regenerative receivers, the oscillator is switched from a filtering mode to oscillating mode using a “quench” signal [43]. The control and shaping of this signal is critical and can lead to design complexities. The adjacent channel rejection of the super-regenerative receiver is poor because the oscillator can be easily locked by a close-by interferer that is larger than the desired signal. The receiver presented in [44] (shown below in Figure 4-6) has very low power consumption of 3.6 mW with a maximum data rate of 500 kBits / s, but it is hard to determine the adjacent channel rejection or dynamic range potential of this receiver. The adjacent channel is reported to be 10 MHz away, but one would expect that interference near the adjacent channel or closer would have to be very low or the receiver will not operate properly. There are likely applications whereby these issues would not be limiting, but for most PANs using the unlicensed ISM frequency bands, interference will limit the performance of the super-regenerative receiver and make it an unsuitable choice.

Figure 4-6: Super-Regenerative Receiver by Chen, et. al. [44]

A second example of an “oscillating receiver” is shown below in Figure 4-7. In this example, the oscillator is injection locked by the incoming RF signal, as in the super-

regenerative receiver. To ensure that the oscillator operates at the correct frequency, a phase locked loop (PLL) is used periodically to lock the oscillator. FSK demodulation comes as a by-product of the phase and frequency detection of the PLL. The receiver's data rate is limited by the locking bandwidth of the VCO and in this case is reported as 5 kbits / s with an overall power consumption of 5 mW. The power consumption is low, but the data rate is also very low. A higher power receiver could have a raw data rate as much as 500 times the data rate of this receiver and through duty cycling would ultimately be much lower power. Also, the modulation method is not flexible with this architecture and thus this receiver cannot be used for many standard-based radios.

Figure 4-7: Injection Locked Receiver by Popplewell, et. al. [45]

4.2.6 Sub-Sampling

A sub-sampling receiver is one in which the incoming RF signal is down-converted to baseband using a sample and hold circuit (sometimes referred to as a sub-sampling mixer or simply sampler). The sampler is clocked at a lower frequency than incoming RF

signal, but higher than the bandwidth of the signal. This type of a receiver is similar to the super-heterodyne receiver, however the first mixing stage is replaced with the sub-sampling stage and the second mixing is done directly to baseband, but in the digital domain. The IF signal is a sample-domain signal at exactly $\frac{1}{4}$ of the sampling frequency as described in Chapter 5.

Figure 4-8: Sub-sampling Receiver

4.2.6.1 Sub-sampling Mechanism

In a regular sampling scheme (sometimes called Nyquist Sampling), the sampling frequency is at least 2 times higher than the signal being sampled as dictated by the Nyquist Theorem [13]. Figure 4-9 shows an example of a signal sampled in this manner. The input signal is shown as a solid line and the resulting discrete-time sampled signal is shown with markers.

Figure 4-9: Nyquist Sampling

In a sub-sampling system, the sampling frequency is lower than the desired signal frequency. At first this may seem to disobey the Nyquist Criteria, but in fact all that is required is that the sampling frequency is greater than twice the signal bandwidth and not necessarily greater than the signal center frequency. Once the signal is sub-sampled, it is indistinguishable from a signal that was within the Nyquist Frequency. Figure 4-10 shows an example of sub-sampling where the input signal is the solid line and the sampled signal is shown by markers. The offset between the input frequency and an integer multiple of the sampling frequency in Figure 4-10 is equal to the input frequency of Figure 4-9. In other words, the act of sampling has converted the high frequency signal in Figure 4-10 to its sample-domain baseband equivalent, which is indistinguishable from the sample-domain equivalent of the signal in Figure 4-9. The sampled signals from these two examples are exactly the same and this effect is also referred to as aliasing. While the idea of sub-sampling seems to go against the Nyquist

Theorem, the important point is that the bandwidth of the desired signal must still be less than $\frac{1}{2}$ of the sampling frequency.

Figure 4-10: Sub-Sampling

The sub-sampling receiver allows the use of a much lower frequency synthesizer. Other architectures generally require an RF PLL to convert the incoming signal to an IF or directly to baseband. With a sub-sampling receiver, the reference frequency used to convert the signal is much lower. This ultimately results in much lower power consumption.

The analog baseband portion of the receiver can be simplified as well. Provided that a solution is found for an anti-alias filter before the sub-sampling mixer, the analog baseband dynamic range requirements are low (Note: this is covered extensively in Chapter 5 and Chapter 6). Also, the discrete-time signal resulting from the sampling lends itself well for integration with a discrete-time sigma-delta ADC.

The downside of the sub-sampling receiver is the required bandpass filtering before sub-sampling. This filter will need to be tunable for each channel and at least a moderate

value of Q is required to overcome the high noise figure of the sampler and improve the image rejection of the sub-sampling stage. Also, the sampling frequency must be high enough such that these image frequencies are sufficiently far away to meet interference requirements. A high sampling frequency will make the design of the sample-domain baseband difficult. These concerns are addressed further in Chapter 5.

An example of the key circuit blocks of a sub-sampling receiver is shown by Pekau and Haslett [34]. The circuits described here are essentially the first 3 blocks of the basic sub-sampling receiver diagram shown in Figure 4-8. This sub-sampling mixer was designed to sample a 2.4 GHz RF signal with a 100 MHz sampling clock. The reported power consumption is over 75 mW, but most of the power is going to output buffers that would not be required in a fully integrated receiver. The authors included bandpass filtering with the intention of making this a receiver front-end, but the Q of the LNA load is only 4.7. This means that the first image of the receiver (at a 50 MHz offset) will see almost no attenuation. In fact, the 3 dB bandwidth of the filter includes close to 10 images. Also, the effect of noise folding is contributing significantly to the reported noise figure of 22 dB. In order to make this front-end suitable for the receiver of Figure 4-8, an external band-filter would be required to further reduce the images and noise out of band. Alternatively, a band-pass filter Q much higher than 4.7 would improve the receiver performance without the external filtering. The proper choice and relationship of band-pass filter Q , image rejection and noise figure is given in detail in Chapter 5.

4.3 Receiver Architecture Summary

The receivers described above are summarized in Table 4-1. The benefits of the sub-sampling receiver create the possibility of building a receiver with less power consumption than the Direct Conversion or Low IF architectures. This is largely due to the fact that the sub-sampling receiver does not require an RF PLL or quadrature (I & Q) baseband paths. This comparison can be made more clearly by contrasting some specific examples and estimating the power savings converting from Direct Conversion or Low-IF to a sub-sampling architecture.

Table 4-1: Receiver Architecture Summary

Receiver Architecture	Benefits	Drawbacks
Super-heterodyne	- IF frequency is high enough to avoid image problems	- Often requires off-chip IF filter - High IF makes filter design difficult - Full integration not easily achieved
Direct Conversion	- No image problem - Can be feasibly fully integrated	- DC offset and 1/f very difficult problems - Baseband filter requires very high dynamic range at the input
Low-IF	- DC offset and 1/f problems greatly reduced	- Image frequency is close, requiring image reject receiver
Oscillating Receivers	- Low power consumption	- Constant envelope modulation required - Data rate is limited, leading to higher comparable power consumption - Interference performance is poor
Sub-sampling	- Low power consumption - Data rate not limited	- Band-pass anti-aliasing required; high Q anti-aliasing requires precise tuning

An example of a Direct Conversion Receiver is shown in [46] where a multi-standard cellular transceiver was designed. The receiver power consumption is reported as 91 mW with almost 50% of that coming from the fractional-N PLL Synthesizer required to generate the signals to mix with the incoming RF signal. In this receiver, the synthesizer maximum output is greater than 4 GHz to accommodate the multi-standard operations in frequency bands ranging from 800 MHz to 2200 MHz. A sub-sampling architecture can cover all these bands with a lower frequency synthesizer because of the nature of the sub-sampling frequency conversion. One can conservatively expect a reduction in the synthesizer current of 2/3 or 30 mW. To back up this assertion, consider the PLL in [15] that consumes 16 mW to operate at only 900 MHz. This is already a 2/3rd reduction compared to the power consumption of approximately 45 mW above. In a sub-sampling

implementation, the synthesizer frequency would be chosen to be even lower (likely approximately 200 MHz) so the assertion of saving $2/3^{\text{rd}}$ of the PLL power is a good one and in fact, is conservative. It is difficult to make direct comparison between two synthesizers designed for different applications, but it is easy to find many examples of lower frequency PLLs that consume much less power than this estimate assumes [16][17]. Further power savings would also come from the mixers which are reported to consume a total of 6 mW, and from the baseband circuitry. As discussed above, the direct conversion baseband must be designed to deal with $1/f$ noise and DC offset problems, thus a compatible baseband for a sub-sampling receiver would consume less power. It would not be unrealistic to expect a comparable sub-sampling receiver to be 40% of the power of this example or less.

The receiver reported in [48] is a Low-IF architecture for Bluetooth with data rates in the range of 1-3 Mbits/s. The reported power consumption of the receiver is 39 mW. The paper does not give a breakdown by block of the receiver so it is difficult to make assumptions on the potential savings going to a sub-sampling architecture. The RF synthesizer would be a fairly large portion of the receiver budget, as with the direct conversion example above and a lower frequency synthesizer would again be much lower power. Some savings can be expected from the quadrature LO generation and the mixers in the comparable sub-sampling receiver, but maybe not as much from the baseband circuitry. On the other hand, the Low-IF receiver would require very good I/Q matching in the baseband and would still encounter a poorer interference rejection at the image frequency which wouldn't exist in the analogous sub-sampling receiver. In this case a

50% reduction in power consumption can be expected if this receiver used a sub-sampling architecture using similar arguments as with the direct conversion comparison above.

The technology used to implement the integrated circuit radio sometimes has an effect on the power consumption of different architectures. In some cases, one architecture will benefit more or less than another from more advanced (i.e. smaller geometry) silicon technologies. This is most often related to how much of the radio is implemented in digital circuitry versus analog. Digital circuits in general benefit in terms of speed and power more from technology scaling than analog circuits. Thus, the sub-sampling receiver having less power consumption in analog circuit blocks would therefore benefit at least as much or more from technology scaling than a similar Direct Conversion or Low-IF receiver.

4.4 Conclusions

This chapter presented the major receiver architectures used today to create integrated receivers. It is difficult to say that one receiver architecture is superior to all others for all applications, but it can be concluded that the sub-sampling receiver, in terms of power consumption, can be superior to Direct Conversion and Low-IF architectures for a range of applications. Also, compared to some specialized “oscillator based” receivers, the sub-sampling receiver is more flexible and when the limited data rate of the oscillator based receivers is considered, the sub-sampling receiver can be superior in terms of power consumption as well. The main hurdle to making the sub-sampling receiver a reality is implementing effective anti-aliasing filtering at RF frequencies. In the next

chapters the requirements for that filter and what filter implementations and innovations are required will be discussed.

Chapter 5

SUB-SAMPLING RECEIVER ARCHITECTURE

5.1 Introduction

In this chapter the concept of a sub-sampling receiver is explored in further detail. As described in section 4.2.5, the sub-sampling receiver has particular advantages for a low power receiver. One can now analyze the basic sub-sampling architecture in terms of the receiver system requirements outlined in Chapter 3. The sub-sampling receiver has unique relationships in frequency plan and performance such that the standard receiver analysis techniques do not give the system designer a complete picture. In this chapter these relationships are derived and presented in a format to make the system design of the sub-sampling receiver possible.

5.2 Sub-Sampling Frequency Plan

According to the discussion in chapter 4, the only requirement for the sampling signal is that it is greater than 2 times the bandwidth of the signal. To make an efficient receiver however, there are several things to consider when sub-sampling. The actual frequency of the resulting sampled-signal will be determined by the offset of the incoming RF frequency and the closest harmonic of the sampling frequency. The resulting “sampled-IF” signal is shown in Figure 5-1.

Figure 5-1: Sub-sampling Frequency Conversion

After sub-sampling, further conversion of the signal will be required to obtain the baseband I and Q components. This is done by employing a complex mixer and can be described using the following equations:

$$x_I = x_{IF} \cdot \cos(2\pi \cdot f_{IF}t) \text{ and } x_Q = x_{IF} \cdot \cos(2\pi \cdot f_{IF}t - \pi/2) \quad (5.1)$$

where x_{IF} is the output of the sub-sampling mixer and x_I and x_Q are the baseband I and Q signals. Since x_{IF} is a discrete-time signal, equation 5.1 should be rewritten:

$$x_{I,k} = x_{IF,k} \cdot \cos\left(2\pi \cdot \frac{f_{IF}}{f_s} \cdot k\right) \text{ and } x_{Q,k} = x_{IF,k} \cdot \cos\left(2\pi \cdot \frac{f_{IF}}{f_s} \cdot k - \frac{\pi}{2}\right) \quad (5.2)$$

So for a given ratio of f_{IF}/f_s , the signal can be converted to I and Q baseband by multiplying by two different sequences. This conversion could be done in the digital domain, allowing the design of only one ADC following the sub-sampling mixer.

The disadvantage in using an arbitrary ratio of f_{IF}/f_s is that the resulting number sequences will not be simple fractional numbers. Also, if the sampling frequency is fixed for all channels, the ratio will change. This will result in a larger and more power hungry

digital conversion circuit. A better solution would be to choose a fixed ratio for f_{IF}/f_s that results in a simple baseband conversion. By choosing a ratio of $1/4$, the baseband conversion is made very simple. This ratio represents the exact mid-point of the Nyquist band and a carrier at this frequency can be represented by only 3 values (-1, 0 and 1). If a substitute of $1/4$ for is made for f_{IF}/f_s into equation 5.2, then:

$$x_{I,k} = x_{IF,k} \cdot (1,0,-1,0\dots) \text{ and } x_{Q,k} = x_{IF,k} \cdot (0,1,0,-1\dots) \quad (5.3)$$

Restricting the frequency plan to this ratio simplifies the circuit implementation immensely. Equation 5.3 indicates that conversion to baseband I and Q signals is accomplished by first sorting every other signal into two different signals. Second, each resulting signal (now with an effective sampling frequency of $f_s/2$) is multiplied by a sequence of (1,-1,1,-1...) to obtain I and Q. One is left with the choice of implementing this conversion digitally, where the circuitry is now trivial, or one can still separate the signals in analog before the ADC. This is accomplished by using two sample and hold circuits that run at $1/2 f_s$ from each other and are out of phase by one sample. The final step of multiplying by (1,-1,1,-1...) can be accomplished either digitally, or by swapping the positive and negative terminals (in a fully differential circuit) on every other sample.

Restricting the choice of sampling frequency creates the requirement that the sampling frequency changes for each channel. The resulting relationship between channel frequency (f_{RF}) and the sampling frequency (f_s) is:

$$f_s = \frac{f_{RF}}{n \pm \frac{1}{4}} \quad (5.4)$$

Where n is the sub-sampling ratio. Note that the range of sampling frequencies required is not necessarily that wide. The different sampling frequencies required between different channels are determined roughly by the channel spacing divided by n .

5.3 Noise and Image Folding

The sub-sampling mixer is essentially just a sample and hold circuit. The act of sampling has the effect of “folding” many images down to the signal band. As described in the previous section, this well known effect is also called aliasing [15]. The aliasing in the sub-sampling mixer has two negative effects. Firstly, all of the noise at the input of the filter is folded down and appears in the Nyquist band. This has the effect of increasing the noise density in the output signal by a factor of at least $2n$, where n is the sub-sampling ratio. In baseband sampling where the sampling frequency is at least twice the maximum signal frequency, a low-pass anti-alias filter is used to limit the amount of noise folded by the sample and hold. In the case of sub-sampling, a band-pass filter is required to perform the same operation.

The second effect of the folding is to create images that can interfere with the desired signal. Figure 5-2 shows the image frequencies resulting from sub-sampling. Following the guidelines for the choice of f_s from section 5.2, the images will be offset from the signal frequency by $f_s/2$. To reduce this effect, some form of band-pass filtering is necessary in combination with a minimum requirement on the sampling frequency, f_s .

Figure 5-2: Noise Folding and Images in Sub-sampling

5.4 Basic Architecture

One is now prepared to look at the basic architecture of the sub-sampling receiver. The sub-sampling receiver architecture enables the use of very low power frequency synthesis and analog to digital conversion. A bandpass filter with a suitably high Q will be required to make use of this architecture. Furthermore, pre-tuning the center frequency of the filter for a specific channel is also necessary so that the noise and image performance can be maintained over a large band of channels. The exact filter topology used will depend somewhat on the specifications of the overall system, but the requirement for tunability and the relatively high Q requirements mean that the filter will be at least partially integrated and will have an active component. The active component can be made without inductors by using a G_m -C filter topology [31], but it is often more efficient to include an inductor in an RF filter. These filters are often referred to as Q -enhanced filters [33][34]. The actual tuning of the center frequency and possibly the Q of

the filter can present a challenge but there are several methods for dealing with this problem [32][29]. The structure and tuning of Q-enhanced filters is covered in Chapter 6.

The basic structure of the receiver is shown in Figure 5-3. The input match (1) is used to interface the integrated receiver to the antenna. This typically means transforming a high impedance LNA input to a 50 Ohm antenna. The LNA (2) provides gain to allow the overall system to achieve a low noise figure. An anti-alias band-pass filter (3) is necessary to alleviate the effects of the sub-sampling as was discussed in section 4.2.5. The sub-sampling mixer (4) samples the incoming RF signal at the sampling rate f_s . Finally, an ADC (5) is required to convert the signal to digital. For simplicity, some blocks are missing from this diagram including the programmable gain required in front of the ADC for gain control and the digital baseband processor that completes the receiver. In many cases the final filtering and demodulation is processed digitally. Also missing is the source for the sampling signal, f_s . Depending on the overall system this would be some form of PLL derived from the system crystal oscillator. This block diagram is sufficient for a general analysis of the performance of a sub-sampling receiver.

Figure 5-3: High Level Diagram of RF Sub-sampling Receiver

5.5 Selectivity & Image Rejection

The selectivity in-band will be determined largely by the linearity of the LNA, the band-pass filter and Q of the filter. It is essentially a similar situation to that shown in Figure 3-3 where the filtering is placed before the sub-sampling mixer and the baseband circuitry, thereby reducing the effect of the linearity of these later blocks.

The choice of sampling frequency will be tightly coupled to the choice of Q and the required image rejection at an offset of half the sampling frequency. Figure 5-4 shows an example of a 2nd order bandpass filter with a Q of 200, a sampling frequency of 45 MHz and a center frequency of 920 MHz. In this case the image rejection is approximately 20 dB in the first image, seen at 22.5 MHz above and below the channel frequency. In order to improve this image rejection, either the sampling frequency will need to be increased, or the Q of the bandpass filter will need to be larger.

Sometimes, rather than focus only on the image rejection, it might be more important that the images are a certain minimum distance away so that the image rejection specification can be relaxed. For example, if there is significant filtering external to the chip (e.g. an external band-select filter) or if the nature of the environment is that interference is not expected outside a certain band, it would be highly desirable to have the images at least that far away. Images further away than the first image will see a higher rejection and possibly much higher if there is significant off-chip filtering.

Figure 5-4: Example of the Filter Q and Sampling Frequency Effects on Image

5.5.1 Generalized 2nd order Filter Requirement

A more general statement on the filtering required for a sub-sampling receiver can be made. First assume that the filter is a simple 2nd order bandpass filter (e.g. single tank circuit LC filter) with the following transfer function:

$$H(s) \propto \frac{\omega_0 s}{s^2 + \frac{\omega_0}{Q}s + \omega_0^2} \quad (5.5)$$

Equation 5.5 can be used to show the attenuation of the filter at a given frequency offset by dividing the amplitude at the center frequency by the amplitude at the desired offset:

$$Attenuation = \frac{H(j\omega_0)}{H(j\omega)} = \frac{f_0 f}{\sqrt{Q^2(f_0^2 - f^2)^2 + (f_0 f)^2}} \quad (5.6)$$

Where f_0 is the center frequency and the channel frequency and f is the frequency at which the attenuation is calculated. Using equation 5.4 the center frequency is expressed in terms of the sampling frequency and the sub-sampling ratio, n :

$$f_0 = f_s \left(n + \frac{1}{4} \right) \quad (5.7)$$

The filter attenuation at the first image frequency is of primary concern. The first image frequency will appear at an offset of $f_s/2$ from f_0 , so choosing the first lower image:

$$f = f_s \left(n - \frac{1}{4} \right) \quad (5.8)$$

It is also useful to express the Q of the filter as a function of the sub-sampling ratio, n as well. Defining x as follows:

$$x = \frac{Q}{n} \quad (5.9)$$

Combining equations 5.7, 5.8 and 5.9 into equation 5.6:

$$Attenuation = \frac{\left(n + \frac{1}{4} \right) \cdot \left(n - \frac{1}{4} \right)}{\sqrt{(xn)^2 \left(\left(n + \frac{1}{4} \right)^2 - \left(n - \frac{1}{4} \right)^2 \right)^2 + \left(\left(n + \frac{1}{4} \right) \cdot \left(n - \frac{1}{4} \right) \right)^2}} \quad (5.10)$$

Expanding further:

$$Attenuation = \frac{\left(n^2 - \frac{1}{16} \right)}{\sqrt{x^2 n^4 + \left(n^2 - \frac{1}{16} \right)^2}} \quad (5.11)$$

Since n is the sub-sampling ratio and is expected to be an integer of at least 3 or greater, $(n^2-1/16)$ can be approximated as n^2 . The term n drops right out of equation 5.11 and thus:

$$Attenuation = \frac{1}{\sqrt{x^2 + 1}} \quad (5.12)$$

Equation 5.12 looks very similar to a low-pass filter function and the attenuation does happen to drop off by 20 dB per decade with respect to x . Figure 5-5 shows a plot of the image attenuation versus x . Note that the first image below the signal was used for this calculation. The image attenuation at the image higher than the signal will be slightly less than equation 5.12 predicts, but if the sub-sampling ratio is reasonably high, the difference will be very small.

As one would expect, the Q of the anti-alias filter must be set in conjunction with the choice of sub-sampling ratio. Equation 5.12 can be used to determine the minimum Q required given a desired image attenuation and sub-sampling ratio. For example, in order to get an image rejection of -20 dB, the Q of the anti-aliasing filter must be at least 10 times the sub-sampling ratio. The sub-sampling receiver designer has only one option to relax the Q requirements of the anti-aliasing filter and that is to decrease the sub-sampling ratio (i.e. increase the sampling frequency, f_s). If the images are placed out of the signal band, the image attenuation specification can often be relaxed, thereby allowing the ratio x to be lower. Even so, to design a sub-sampling receiver with a sub-sampling ratio of 10 or more, a high Q band-pass filter will be required.

Figure 5-5: Image Attenuation versus Q/n ratio (x)

The derivation above can be extended to include all the images generated when sub-sampling. These images will be attenuated to a greater extent than the first image, but in some systems if they land in a certain interference band they could be of greater importance. To do this, one can redefine the image frequency from equation 5.8 as:

$$f = f_s \left(n + \frac{1}{4} - \frac{k}{2} \right) \quad (5.13)$$

Where k is an integer and f is now equal to the frequency of the k^{th} image away from f_0 .

Substituting this into equation 5.6 and doing some simplification:

Figure 5-6: k^{th} Image Attenuation versus Q/n ratio (x)

Although the higher order images are attenuated to a greater degree than the first image, the attenuation does not drop off as quickly as some applications might require. With k values of 3 or greater, the image frequency is far from the signal band, assuming the first image is designed to be outside of the signal band. Significant interference is possible at these offsets, but it is normally realistic to expect further attenuation of these images by the input match of the receiver, the antenna and possibly an external filter used for band selection. A combination of these sources of filtering are accounted for in the receiver design by placing a specification on the external filtering required at given frequency offset. An example of this is shown in Figure 5-7 where the external filtering has a

specified attenuation of 10 dB at the second image, 20 dB at the third image and 30 dB for images further out. The image attenuation for these images further away is now much greater and would allow the receiver to function in the presence of very large signals operating at these frequencies.

Figure 5-7: k^{th} Image Attenuation versus Q/n ratio (x) with External Attenuation

5.5.2 Higher order Anti-aliasing Filtering

Naturally, the anti-aliasing filtering will be better if the filter is higher than 2nd order band-pass. This is often possible for lower frequency filters, where on-chip component matching and tuning can make the poles and zeros of a filter very precise. At RF frequencies it is much more challenging to construct higher order filters. There are

examples of 4th order Q-enhanced filters [32] using a transformer and two equivalent tank circuits rather than a simple L-C. The addition of more tank circuits requires more Q-enhancement transconductors to boost the Q. On the other hand, increasing the Q of a Q-enhanced filter comes at a very small power consumption penalty. The only drawback is a reduction in the dynamic range of the Q-enhanced filter.

Higher order RF filters would also require more circuit components, most importantly inductors. If these are off-chip components, their cost is of concern and on-chip the area required to accommodate the higher order filter can become prohibitive. The final consideration for higher order filters is the tuning procedure required to make up for the process variation. The low bandwidth required by the sub-sampling receiver requires that the filter is tuned precisely, necessitating a direct tuning method [5]. The tuning method described in Chapter 5 would be much more complicated for even a 4th order filter. Interaction between the two tank circuit sections would be difficult to predict and the overall Q and center of the filter would be hard to tune independently unless a complex algorithm was employed.

Despite the above concerns, it is still possible to employ a higher order filter as the anti-alias filter of a sub-sampling receiver. The analysis above could be repeated for a 4th order LC filter with all poles at the same frequency, or a more complicated formula could be derived for an arbitrary filter. It is much more likely that architecture designs will favour the simpler 2nd order structure and if higher image rejection is required, the operating Q of the filter will be increased.

5.6 Gain & Noise Figure

Initially one can assume that the filter (item 3 in Figure 5-3) is ideal, meaning it adds no noise and has unity gain. For simplicity, the sampler gain can be assumed to be 1 also. In practice the receiver's gain is distributed through several blocks (i.e. not just in the LNA) and there are generally losses associated with passive sub-sampling mixers, but these losses can be minimal (on the order of 1 dB) and will not affect the overall analysis greatly.

The overall noise figure of the simplified receiver in Figure 5-3 is given by the Friis equation as [9]:

$$NF = NF_{Frontend} + \frac{(NF_{smp} - 1) + (NF_{ADC} - 1)}{G_{Frontend}} \quad (5.16)$$

Ideally one will have to match the gain of the receiver and the noise of the ADC such that the noise figure of the ADC is not significant. This is generally true of receivers operating at their maximum gain setting. Some kind of gain control is usually used to place the incoming signal sufficiently above the noise floor of the ADC. Rewriting the overall noise figure gives:

$$NF = NF_{frontend} + \frac{NF_{smp} - 1}{G_{frontend}} \quad (5.17)$$

5.6.1.1 Sampler Noise Figure

The sampler noise figure can be broken down into two parts. First, there is the noise that is added to the signal by the internal circuitry of the sampler itself, also referred to as the intrinsic noise. Secondly, the noise from previous stages of the receiver, at the input of

the sampler is folded down by the sub-sampling action and is therefore accounted for by the noise figure of the sampler itself. This is called the extrinsic noise. The sampler can be modeled as an ideal switch and capacitor with a real resistance, as shown in Figure 5-8.

Figure 5-8: Ideal Model of Sampler (Sub-sampling Mixer)

The resistance, R_{sw} represents the finite resistance of the switch when it is on. It is ultimately the source of the intrinsic noise of the sampler, sometimes referred to as kT/C noise [50]. This result can be found by first using the noise density of the resistor:

$$V_n^2 = 4kTR_{sw} \quad (5.18)$$

Since the sampler's RC forms a low pass filter, the actual input referred noise is filtered by this RC. To get a total noise this is integrated over all frequencies [18]:

$$V_{no} = \sqrt{4kTR_{sw} \int_0^{\infty} \frac{d\omega}{1 + (\omega R_{sw} C_s)^2}} = \sqrt{4kTR_{sw} \frac{1}{4R_{sw} C_s}} = \sqrt{\frac{kT}{C_s}} \quad (5.19)$$

As a result, the resistor is not part of this equation. So hence the noise figure requirement of the sampler will only specify a minimum size of capacitance independent of switch design. Note that so far it is assumed that all of the noise at the output of the sampler is included in the signal to noise ratio. In practice, there is further filtering of the signal

following the sampler which can be analog or digital, sometimes referred to as the baseband filtering. This filtering is quite significant because it effectively reduces the amount of noise that is of concern at the output of the sampler. Similar to an over-sampled ADC, this is referred to as over-sampling. The reason for this term is that the sampler's frequency is higher than it needs to be given the Nyquist Criterion. To account for this effect, the sampler's noise output can be modified by taking into account the over-sampling ratio (OSR) [15]:

$$\overline{V_{noi}^2} = \frac{kT}{C_s} \cdot \frac{f_{BW}}{f_s/2} = \frac{kT}{C_s} \cdot \frac{1}{OSR} \quad (5.20)$$

Where f_{bw} is the effective noise bandwidth of the baseband filtering and f_s is the sampling frequency. This can also be expressed as a noise spectral density as:

$$\overline{e_{no}^2} = \frac{kT}{C_s} \cdot \frac{2}{f_s} \quad (5.21)$$

On top of the sampler's intrinsic noise, one has the effect of folding noise from the preceding stages. When sampling at the Nyquist rate, it can be assumed that the incoming signal is band-limited to the Nyquist frequency and noise folding would be ignored. But in the case of sub-sampling, the incoming signal has a much higher frequency and it is the same action of folding that is being used to down-convert our RF frequency.

Because the noise calculations for preceding stages would have taken into account the noise bandwidth of the baseband filtering (f_{bw}), noise outside of this bandwidth would normally be ignored. For example, if an LNA has a quoted noise figure, it assumes that

the incoming signal is considered to be over a fixed bandwidth and the noise in that bandwidth is compared to the incoming signal. In other words, only the noise density of the signal of interest is typically considered. When the signal is sub-sampled, noise at intervals of $f_s/2$ is folded down, so in calculating this effect an assumption about the incoming noise at these frequencies must be made. For simplicity, start by assuming that the incoming noise spectral density is constant, and the bandwidth at the input of the sampler is higher than that of the sampler (note that these assumptions would be neither typical nor practical, but adjustments will be made later). In this case, the incoming noise is increased as follows:

$$\overline{V_{no}^2} = \overline{V_{ni}^2} \cdot \frac{2 \cdot f_{3dB}}{f_s} \quad (5.22)$$

Where $f_s/2$ is the Nyquist frequency and f_{3db} is the corner frequency set by R_{sw} and C_s . Assume that in designing the sampler, f_{3db} would have to be sufficiently above our RF frequency to allow for process variation. Since a band-pass anti-aliasing filter is employed to filter out unwanted interference at the input of the sampler, the noise spectral density will be shaped by this filter. The previous section outlined how the attenuation of the k^{th} image is related to the Q of the band-pass filter and the sub-sampling ratio, n . This approximation can be used to find out the true extent of the noise folding effect. Using the result in equation 5.15, and summing the power of all the folded noise terms gives:

$$\overline{V_{no2}^2} \approx \overline{V_{ni}^2} \cdot \sum_{k=-\infty}^{\infty} \frac{1}{x^2 k^2 + 1} \quad (5.23)$$

Where the summation for all k represents the power of all images after filtering. The center of the channel band is found at $k=0$ and as expected the gain of the filter is 1. Note that the finite bandwidth of the sampler has been ignored, but it would further reduce higher image frequencies. The sum can be approximated as the following:

$$\overline{V_{no2}^2} \approx \overline{V_{ni}^2} \cdot \left(1 + 2 \cdot \sum_{k=1}^{2n} \frac{1}{x^2 k^2 + 1} \right) \quad (5.24)$$

Where it is assumed that images above $2n$ are very small and would not contribute significantly to the overall noise summation. The second term is now a summation that captures the power of all image frequencies while the term 1 represents the in-band noise power. The second term of equation 5.23 can be further approximated using the solution to Basel's problem [42]:

$$\sum_{k=1}^{2n} \frac{2}{x^2 k^2 + 1} < \sum_{k=1}^{\infty} \frac{2}{x^2 k^2} = \frac{2}{x^2} \cdot \frac{\pi^2}{6} \approx \frac{3.3}{x^2} \quad (5.25)$$

Equation 5.25 can now be used to simplify equation 5.24 as:

$$\overline{V_{no2}^2} \approx \overline{V_{ni}^2} \cdot \left(1 + \frac{3.3}{x^2} \right) \quad (5.26)$$

Now that the noise sources of the sampler have been characterized, one can proceed to derive a noise figure for it. Using the simplifying assumption that the sampler gain is 1:

$$NF = \frac{SNR_{in}}{SNR_{out}} = \frac{\overline{V_i^2} / \overline{V_{ni}^2}}{\overline{V_o^2} / \overline{V_{no}^2}} = \frac{\overline{V_{no}^2}}{\overline{V_{ni}^2}} \quad (5.27)$$

Adding the two sources of noise together from equations 5.20 and 5.23, and using the approximation from equation 5.25 to derive the noise figure of the sub-sampling mixer gives:

$$NF = \frac{\overline{V_{no}^2}}{V_{ni}^2} = 1 + \frac{3.3}{x^2} + \frac{\frac{kT}{C} \frac{1}{OSR}}{V_{ni}^2} \quad (5.28)$$

Using the noise spectral density for the output noise (equation 5.21), add the input noise as that caused by the source impedance and the noise figure becomes:

$$NF = \left(1 + \frac{3.3}{x^2}\right) + \frac{1}{2R_{source}C_s f_s} \quad (5.29)$$

For relatively small values of x , the first term in equation 5.29 is small. For example, when $x=4$ the noise figure due to the first term only is 0.8 dB. In the previous section it was determined that for reasonable image rejection performance, the Q/n ratio (x) will likely have to be at least 10. If this is the case, the first term can be ignored and the noise figure of the sampler then reduces to:

$$NF = 1 + \frac{1}{2R_{source}C_s f_s} \quad (5.30)$$

For some reasonable values, this still works out to a very large number. Figure 5-9 shows that even with a sampling capacitor of 0.8 pF, a sampling frequency of 50 MHz yields a noise figure of almost 25 dB. This will require a very high gain from the LNA in order to achieve an overall noise figure that is acceptable. For example, a typical Bluetooth receiver [47] has an overall noise figure of approximately 6 dB to achieve a sensitivity of

-88 dBm. If a sub-sampling architecture was chosen where the sub-sampling mixer noise figure was 25 dB and the front-end noise figure was 5 dB, then approximately 26 dB of gain would be required by the front-end to meet the sensitivity requirements. This can be generalized by re-arranging equation 5.17 as:

$$G_{frontend} = \frac{NF_{samp} - 1}{NF_{tot} - NF_{frontend}} \quad (5.31)$$

Figure 5-9: Noise Figure for Sampler as a Function of Sampling Frequency

5.6.1.2 LNA gain and Noise Figure

The high noise figure of the sampler will cause the required gain of the front-end to increase. To obtain a low noise figure, the gain of the front-end will have to be equal to or greater than the noise figure of the sampler. This means a gain of upwards of 30 dB. The combination of an LNA and Q-enhanced filter is ideally suited to meet this requirement the design and implementation of which is covered in Chapter 6. The main challenge is to achieve the gain and noise figure required while maintaining a reasonable linearity.

5.7 Phase Noise

A sub-sampling architecture has no significant effect on the receiver phase noise. The common assumption made when analyzing the jitter or phase noise in sub-sampling receivers is that the sub-sampling mixer is treated like a traditional ADC sample and hold [4]. The jitter power is considered over the entire Nyquist Sampling band and a requirement is placed on the sampling clock such that this jitter must be less than the required noise floor of the ADC. This method for considering jitter is justified when there is little or no over-sampling of the signal band, but as discussed earlier, the sub-sampling receiver by its nature requires the signal to be over-sampled to give a reasonable image rejection. As a result, for the sub-sampling receiver it is only necessary to consider the jitter that will fall in the signal band.

The advantage of the sub-sampling architecture is that it does not require an RF synthesizer, but the phase noise specifications on the sampling clock are not different than what would be required of the RF synthesizer. This is best illustrated by comparing the two specifications.

First, start by defining the phase noise of a sampling (or mixing) signal as $\varphi(t)$. The signal used to derive the sampling clock is itself a sinusoidal signal:

$$s(t) = A \sin(\omega_s t + \varphi(t)) \quad (5.32)$$

Where ω_s is the sampling frequency ($2\pi f_s$). If the phase noise is replaced with a definition of jitter ($\tau(t)$) using the relationship:

$$\tau(t) = \frac{\varphi(t)}{\omega_s} \quad (5.33)$$

then:

$$s(t) = A \sin(\omega_s(t + \tau(t))) \quad (5.34)$$

Equation 5.34 is a general definition of a mixing signal with an arbitrary jitter. In a sampling system, the zero crossing of the mixing signal is the only part of interest. A sampling signal defined by the zero crossings of 5.34 can be defined as:

$$x_s(t) = \sum_{n=-\infty}^{\infty} \delta(t - nT_s + \tau(t)) \quad (5.35)$$

Where T_s is the period of the sampling clock. If an incoming signal, v_{in} , is sampled then the sampled version of v_{in} , defined here as v_s ,

$$v_s(t) = v_{in}(t) \sum_{n=-\infty}^{\infty} \delta(t - nT_s + \tau(t)) \quad (5.36)$$

In other words, each sample of v_s can be described as:

$$v_{s,n} = v_{in}(nT_s - \tau(t)) \quad (5.37)$$

A slight approximation can be made by simplifying equation 5.37 as:

$$v_{s,n} = v_{in}(nT_s - \tau(nT_s)) \quad (5.38)$$

Where it is assumed that the jitter can be sampled at the precise sampling instant, nT_s , rather than the jittery sampling instant. One now essentially have a discrete series of random jitter samples, defined by $\tau(nT_s)$.

If v_{in} is a sinusoidal signal, defined as:

$$v_{in}(t) = A \cdot \sin(\omega_{RF}t) \quad (5.39)$$

Then the sampled output, v_s would be:

$$v_{s,n} = A \cdot \sin(\omega_{RF}(nT_s - \tau(nT_s))) = A \cdot \sin(\omega_{RF}nT_s - \omega_{RF}\tau(nT_s)) \quad (5.40)$$

Going back to the original definition of phase noise from equation 5.33, and substituting it back into equation 5.40 yields:

$$v_{s,n} = A \cdot \sin\left(\omega_{RF}nT_s - \frac{\omega_{RF}}{\omega_s}\varphi(nT_s)\right) \quad (5.41)$$

So, the relationship between the phase noise in the sampling clock to the resulting phase noise in the sampled signal is:

$$\varphi_{out,n} = \frac{f_{RF}}{f_s}\varphi_s(nT_s) = N \cdot \varphi_s(nT_s) \quad (5.42)$$

Where $\varphi_{out,n}$ is the output phase noise, φ_s is the phase noise of the sampling signal, and N is the sub-sampling ratio (f_{RF}/f_s). So the phase noise amplitude simply increases linearly with the sub-sampling ratio, or alternatively the phase noise power goes up by a factor of N^2 when sub-sampling. This is the same result that was obtained in [1]. In comparison, a system where the signal conversion is accomplished by using an RF synthesizer, the phase noise will be essentially multiplied up by the same factor. In-band phase noise amplitude in a synthesizer is proportional to the divide ratio of that synthesizer [60]. Since the synthesizer in a sub-sampling receiver is lower in frequency compared to a direct conversion or Low-IF receiver, by the sub-sampling ratio n , the phase noise in the

synthesizer would be expected to be lower by a factor of n^2 . Therefore the overall in-band phase noise in the end is the same between the sub-sampling and traditional receivers. Out of band phase noise can be a problem with reciprocal mixing where a large interfering signal can be mixed with the synthesizer phase noise to interfere with the desired signal. Again with a sub-sampling receiver, the phase noise power will be multiplied by n^2 . Since the interferer is likely outside of the synthesizer bandwidth, it is the VCO in the synthesizer loop that is likely the source of this phase noise. Leeson's formula shows that the phase noise amplitude of a VCO is proportional to its absolute center frequency [9]. The VCO in the synthesizer for a sub-sampling receiver will therefore have n^2 less noise than the equivalently designed RF VCO of a traditional receiver. The total effect of phase noise will again be the same once the signal is sub-sampled. Figure 5-10 illustrates on the left a traditional RF synthesizer used for frequency conversion and on the right a sub-sampling mixer. With the RF synthesizer, the jitter added to the signal will be proportional to f_{LO} [4]. With the sub-sampling mixer, the jitter is proportional to f_{RF} , as shown above. The LO frequency may be lower than the RF for one conversion stage, but in the end the RF signal will be converted to baseband, so an equivalent amount of jitter will have been added between the two systems.

Figure 5-10: Comparison of Frequency Conversion Techniques

5.8 Conclusions

This chapter has covered the fundamental design relationships in a sub-sampling receiver; namely the required Q of the band-pass anti-aliasing filter given the sub-sampling ratio and the resulting noise figure of the sub-sampling mixer. Additionally, the phase noise resulting for the use of sub-sampling was analyzed and it was determined that there is no phase noise penalty for the sub-sampling receiver. The relationships shown in this chapter are important tools for designing a sub-sampling receiver and will be used in the example sub-sampling receiver implementation in Chapter 7. To implement a high Q anti-alias band-pass filter, the Q -enhanced filter was suggested as a solution. Chapter 6 covers the analysis, design and tuning of this type of filter in more detail.

Chapter 6

Q-ENHANCED FILTERS

6.1 Introduction

To enable the use of a sub-sampling receiver, a bandpass filter is required for anti-aliasing. This will both reduce the noise figure of the sub-sampling mixer and work to relax the minimum required sampling frequency as shown in Chapter 5. This chapter will first give the background on the basic structure of the Q-enhanced filter. The noise and linearity will be analyzed as it is needed in the system design of the sub-sampling receiver. Finally, the methodology for directly tuning the Q-enhanced filter is covered in detail. It is important that the filter is tuned directly in order to meet the high Q required to make the sub-sampling receiver possible. As seen in Chapter 5, in order to have a reasonably high sub-sampling ratio, the Q of the anti-alias filter must also be high.

6.2 Basic Q-enhanced Filter

LC filters are a desirable circuit to use in a RF receiver front end, but they suffer from very low quality factor (Q). An integrated inductor normally has a Q of about 5, while off-chip inductors can range from 10 to 30 or slightly more for higher quality wire-wound inductors. When operating at RF frequencies, Qs of this level do not significantly filter in-band interference. If a higher Q can be used, the effect is a relaxing of the specifications for linearity for subsequent stages and reducing the sampling frequency required to obtain a certain image rejection.

Active circuits can be used to compensate for the losses in the inductor. These filters are commonly referred to as a Q-enhanced LC filter. The circuit for the Q-enhanced filter is relatively simple, employing only two transconductors. One transconductor is used for the input and the other to make up the inductor compensation or negative resistance. A major advantage of Q-enhanced filters is that it is possible to design a Q-enhanced filter for higher frequency operation and lower power than the equivalent G_m -C filter [2].

A basic Q-enhanced filter design is illustrated in Figure 6-1. This structure is based on a simplified second order topology using parallel mode realization where g_0 represents the approximate equivalent parallel loss resistance of the inductor and accompanying circuitry. Transconductance g_m indicates the negative resistance designed for offsetting the circuit losses. Transconductance g_i converts the input signal to a current and together with the Q-enhancement transconductor complete the second order filter structure.

The enhanced quality factor of the filter can be shown to be [5]:

$$Q = \frac{\omega_0 C}{g_0 - g_m} = \frac{1}{1 - \frac{g_m}{g_0}} \cdot Q_o \quad (6.1)$$

Where Q_o is the initial Q of the resonant circuit (without enhancement and typically limited by the inductor Q) and ω_o is the filter center frequency given by,

$$\omega_o = \frac{1}{\sqrt{LC}} \quad (6.2)$$

Through suitable setting of g_m , the filter enhanced Q can be theoretically very large.

Furthermore, setting g_m larger than g_0 can convert the filter into an oscillator.

Figure 6-1: Simplified Q-enhanced Filter Circuit Diagram

Figure 6-2: Differential Q-enhanced Filter

Figure 6-2 shows a more accurate differential circuit implementation of a Q-enhanced filter. The input transconductance is comprised of transistors $M_{1a,b}$ and $M_{2a,b}$. This

transconductance sets the gain of the filter as the current is fed into the tank circuit as shown by the filter's transfer function:

$$H(s) = \frac{\omega_0 s}{s^2 + \frac{\omega_0}{Q}s + \omega_0^2} \cdot \left(g_i \sqrt{\frac{L}{C}} \right) \quad (6.3)$$

Where g_i is the input transconductance. The noise figure of the filter will be largely dependent on this input stage. The input stage can be designed as the LNA of the receiver if the filter is used as the first active stage. The enhancement transconductor is composed of transistors $M_{3a,b}$ and M_4 and resistor R_1 . The transconductor is a basic differential pair with resistive degeneration. By connecting the transconductor in a positive feedback mode, it becomes a negative conductance (in Figure 6-1 this was referred to as “ $-g_m$ ”). The voltage at the current source to the enhancement transconductor, V_Q , controls the Q of the filter. Degeneration resistor R_1 and transistor M_4 are included in the transconductor to improve the linear range of the filter. M_4 operates in the triode region as a resistor and is included to allow the amount of degeneration to be varied to find an optimum balance between dynamic range and power consumption. The resistor R_1 in combination with transistor M_4 forms the degeneration. The high Q operation depends on a difference of $g_0 - g_m$ and thus the gain is very sensitive to changes in the g_m of this transconductor. The use of degeneration decreases the sensitivity of the overall transconductance to the input / output signal level. It can be shown that the overall transconductance, G_m with degeneration is [9]:

$$G_m = \frac{g_m}{1 + g_m R} \quad (6.4)$$

Where g_m is the transconductance of the transistors M_{3a} and M_{3b} , and R is the total degeneration resistance. Here it is seen that the transconductance has been reduced by the feedback term $(1+g_m R)$. Also, the maximum G_m is limited to $1/R$. This means that the degeneration resistance cannot be set higher than the equivalent losses in the tank circuit $(1/g_0)$ or high Q operation will not be possible. For some applications, it may be desirable to change the linearity of the transconductor for different operating conditions. A lower resistance would mean less power required to reach a given Q . Having M_4 as a variable resistor allows this flexibility in operation. The Q of the filter is controlled by the voltage at the current source to the enhancement transconductor, V_Q . This current source is split into two parts, with the degeneration between the two sections so that a DC voltage drop is not seen across R_1 . The tank circuit of the filter is comprised of a real inductor and a tunable bank of capacitors. The inductor can be on or off-chip, but since the Q_0 of the inductor is likely the dominant source of losses in the tank circuit, it will determine the required G_m (from equation 6.4) for the enhancement transconductor.

6.3 Noise and Linearity

The noise contribution from a Q -enhanced filter will largely depend on how the input transconductance is designed. In the case of an RF filter, the noise figure is extremely important, especially if the filter is the first component the input signal experiences. In this case, the input transconductance is the LNA. A detailed analysis of the design of an LNA would be required to optimize the noise figure of the overall filter in this case [10]. The LNA noise can be modeled as a noise current injected into the tank circuit. In fact, the enhancement transconductor will simply add noise to the output signal by injecting its

own noise current into the tank circuit. The noise in a general transconductor can be approximated with [31]:

$$i_n^2 = 4kTFg_m \quad (6.5)$$

Where F is a factor accounting for the specific implementation of the transconductor and is usually between 1 and 2. In the Q-enhanced filter there are two transconductors that add noise to the signal: the input transconductance g_i and the enhancement transconductor, g_m . As well, thermal noise from the resistive losses, R in the tank circuit will be added. The total noise current entering the tank circuit is then:

$$i_{nt}^2 = 4kT(F_i g_i + F g_m + R) \quad (6.6)$$

For high Q operation the value of g_m is set almost equal to the losses in the tank (g_0) and these losses are approximately $1/R$. This means equation 3.7 can be approximated with:

$$i_{nt}^2 = 4kT(F_i g_i + (F + 1)g_m) \quad (6.7)$$

So in the case of an LNA, g_i will be much larger than g_m and the noise will be dominated by the input transconductance. In this situation, having a lower Q_0 and therefore higher g_m will not have a detrimental effect on noise (although will require higher power consumption).

The Q-enhanced filter suffers from a low linearity because of its open-loop nature [3]. Also, the current in the enhancement transconductor is determined almost entirely by the losses in the tank circuit and as such, any large signals in the filter will cause non-linearities that cannot be compensated for by raising the bias current. This is mitigated

somewhat by the degeneration in the transconductor, giving the filter some feedback compensation.

In the case where the noise is dominated by the input transconductance, the dynamic range of the filter can be shown to be inversely proportional to the square of the operating Q [5]:

$$DR \propto \frac{1}{Q_0 Q^2} \quad (6.8)$$

In this situation it makes sense to aim for higher losses in the tank circuit to increase the required current in the enhancement transconductor, thereby improving the linearity. In the situation where the Q-enhanced filter is not used with an LNA, noise of the input transconductor may not be dominant. In this situation, the dynamic range will not be a function of the tank circuit losses (Q_0 drops out of equation 6.8). For this case, the linearity requirements of the filter will need to be balanced with any effects on the noise floor. In practice it is not possible to generalize a filter circuit design for dynamic range without knowledge of the overall block and level breakdowns in noise and linearity. It is the entire receiver dynamic range that is of greater concern.

6.4 Tuning

Using a high Q presents a particular problem in tuning the Q-enhanced filter. Traditional Master-Slave techniques [58] are not feasible due to the matching required between two designs, and since the master design will have a huge area or component penalty by requiring a 2nd tank circuit. Realistically the Q-enhanced filter will need to be tuned directly [5].

6.4.1 Frequency Tuning

By adjusting the Q of the filter very high, the filter is turned into an oscillator. The oscillation frequency is very close to the center frequency of the filter when not in oscillation meaning that coarse frequency tuning of the filter can be done using the receiver's baseband circuitry. A simplified diagram of the important blocks is shown in Figure 6-3. The preferred frequency plan for a sub-sampling receiver, described in Chapter 5 involved converting the received signal to a digital IF centered at $f_s/4$. The digital baseband circuitry in Figure 6-3 shows this digital IF being modulated down to DC. The output of this block is complex I and Q, but is shown in the diagram as a single signal. Channel select filtering is performed on the complex baseband signal in the digital domain. The same filter used for the regular receiver should be sufficient for frequency tuning. An energy detect block is required that gives a scalar value of the signal level (absolute value of I and Q) at the output of the filter.

Figure 6-3: Basic Block Diagram of Frequency Tuning

To coarse tune the frequency of the filter, the filter is placed in oscillation, then the center frequency is varied while the energy detect level at the baseband is measured. When the

center frequency of the filter matches that of the baseband, the energy detect value will be very high. Thus, when the energy detect level passes a pre-determined threshold, the center frequency is considered tuned coarsely.

Due to aliasing, it is possible that the center frequency found using this method will still be incorrect. If the tuning range and variation of the filter's tank circuit is larger than $f_s/2$ (the distance between the channel frequency and the next alias), then it is possible for the first iteration of the tuning method to find the wrong alias. In the next step, it is therefore necessary to change the sampling frequency, f_s such that if the filter were at the wrong alias, the energy detect would no longer be high. Since the frequency plan demands that the sampling frequency be changed over a range of frequencies (to accommodate different RF channels), this should not change the requirements of the synthesizer that drives f_s . This requirement can be put into an equation as follows: Recall the sampling frequency is a function of the RF frequency as:

$$f_s = \frac{f_{RF}}{n \pm \frac{1}{4}} \quad (6.9)$$

As an example, two different sampling frequencies would both convert the same RF frequency to baseband:

$$f_{s1} = \frac{f_{RF}}{20.25}, \text{ and } f_{s2} = \frac{f_{RF}}{19.25} \quad (6.10)$$

If f_{RF} is 915 MHz, this then works out to sampling frequencies of 44.44 MHz and 47.53 MHz. The alias frequencies below and above the desired frequency of 915 MHz (when using f_{s1}) is given by:

$$f_{alias} = f_{RF} \pm \frac{f_{s1}}{2} = f_{RF} \left(1 \pm \frac{1}{20.25} \right) \quad (6.11)$$

Which works out to 960.2 MHz and 869.8 MHz. When using equation 6.11 with f_{s2} instead alias frequencies of 962.5 MHz and 867.4 MHz are found. The alias frequencies of the two sampling frequencies are more than 2 MHz apart. If in the first step of frequency tuning the filter landed in the alias frequency of f_{s1} , when switching to f_{s2} the oscillating filter would no longer cause the energy detect to surpass the threshold. The algorithm would then continue to change the center frequency of the filter and use f_{s1} until the oscillation is seen again by the energy detect. The coarse frequency tuning is complete when the oscillation is seen at both sampling frequencies.

Note that it would not be necessary to perform this coarse frequency tuning sequence more often than the first time or for more than one channel. The tuning would determine a fixed frequency setting for the filter to begin Q tuning and fine frequency tuning. The actual center frequency of the filter will be different than that found by this coarse tuning because in oscillation, the parasitic capacitance of the enhancement transistor is different than when the filter is operating on small signals. This difference leads to a slight error in center frequency that must be corrected for with the filter out of oscillation.

6.4.2 Q Tuning

The Q of the filter should again be tuned directly because of the high accuracy requirements of the sub-sampling architecture. If a filter bandwidth only slightly larger than the channel bandwidth itself is required, a slight change in Q will have a detrimental effect on the signal. Also, the gain and linearity of the receiver is dependent on the

operating Q being a certain value and as was shown by equation 6.8, the linearity in particular has a strong dependence on the Q .

Tuning of a Q -enhanced filter may be done using direct digital tuning that does not require an input reference [5]. This process involves using thermal (white) noise shaped by the filter to determine the filter's Q . Because the filter operates at a much higher Q than any preceding circuits, the shape of the noise at the filter output is a good representation of the Q . The bias control for the enhancement transconductor is then controlled using a digital to analog converter (DAC) built specifically for that purpose. The DAC adjusts the current bias to the transconductor, thereby varying the g_m and therefore Q , according to equation 6.1.

The issue with this approach is that noise at the output of the filter is not necessarily the dominant noise source when the total receiver noise is viewed at the input of the ADC. Even if it is, there is not necessarily sufficient difference between the filter output noise and the noise from the rest of the receiver. This ratio can be referred to as the Noise to Noise Ratio (NNR); the noise of interest being equivalent to the signal normally considered in a Signal to Noise Ratio (SNR). If the receiver in Figure 7-1 is designed so that the NNR from the filter output to the ADC is sufficient for Q tuning, this would be over-designing the noise and gain performance of the sub-sampling mixer, PGA and ADC. A better approach is to devise a special mode meant only for filter tuning that allows the NNR to be temporarily increased while tuning the Q of the filter.

Figure 6-4 shows 2 possible additions to allow for this. The first is an RF amplifier inserted in between the filter to be tuned and the sub-sampling mixer. This amplifier will

increase the noise level before the sub-sampling mixer in order to improve the noise figure of the rest of the receiver. This amplifier will work out as long as the noise figure of the amplifier is low enough, considering that its noise will be folded by the sub-sampling mixer, as described in section 5.3. This will typically require a large, high powered amplifier to meet these objectives. Also, it is important that the amplifier is disabled or placed in a lower gain mode while the receiver is in normal operation. A second option that may function better would be to increase the noise power at the input of the filter. This could be accomplished by switching in a “noise generator”. The only requirement is that the noise be substantially flat in spectrum, at least compared to the bandwidth of the filter. This circuit can be enabled during Q tuning only and could increase the NNR to the point where no further modifications would be necessary in the rest of the receiver.

Figure 6-4: Basic Block Diagram of Q Tuning

The actual tuning of the filter Q and fine tuning of the frequency is accomplished by measuring the energy detected in the receiver baseband at three frequency points. Note that the energy detect circuit gives a measure of the amplitude of the noise seen by the

receiver integrated over the receiver bandwidth. One measurement is taken at the desired filter center and one measurement at equal positive and negative offsets from the desired center. The filter frequency will be centered when the two offset measurements are equal in value. The actual offset used is not important, however the resolution of the energy detect and the NNR must be considered when choosing the frequency offset. Once the frequency is centered, the Q can be measured by comparing the energy detect output at the center of the filter and at one of the frequency offsets. The ratio of these two measurements is compared to a target value given a target Q setting. The required ratio can be determined using equation 6.3 at two different frequencies:

$$\left| \frac{H(j\omega_0)}{H(j(\omega_0 + \Delta\omega))} \right| = Q \cdot \frac{(\omega_0 + \Delta\omega)}{\sqrt{\left(\frac{\omega_0 + \Delta\omega}{Q}\right)^2 + (\Delta\omega + \Delta\omega^2)^2}} \quad (6.12)$$

For example, if a Q of 500 at 915 MHz is desired, the noise amplitude at the center frequency and at a 1.75 MHz offset can be measured and a 6 dB drop in noise amplitude will be seen if the Q is properly tuned. If the amplitude drop is greater, the Q is too high, if it is less, the Q is too low. A slight error has been introduced here by ignoring the fact that the receiver baseband filtering (shown in Figure 6-3) has a finite bandwidth. As long as this bandwidth is much less than the measurement offset chosen, this error will be small. The error can be accounted for by calculating the expected amplitude when integrating the receiver bandwidth, at an offset with equation 6.12.

The Q-tuning step must also include a fine tuning of the frequency in order to ensure the Q tuning is accurate. Also, the center frequency is a mild function of the operating Q of

the filter since the parasitic capacitance of the enhancement transistor will change with bias condition. For this reason fine frequency tuning and Q tuning will be done in each iteration of the tuning algorithm. Frequency adjustments can be made based on the difference between amplitudes of the noise at positive and negative offsets. If the positive offset has a higher power, the frequency must be decreased, etc.

An example of what the output spectrum would look like in a spectrum analyzer is shown in Figure 6-5. A random number generator was used to simulate the noise in the filter and two different plots are included, one with the filter tuned precisely and a second with the Q set high and the frequency low. The test frequency bands are shown as dotted lines and represent where the noise amplitude measurements are taken.

Figure 6-5: Filter Q Tuning Example Plot

6.5 Conclusions

This chapter outlined the basic structure of the Q-enhanced filter and analyzed the noise and linearity of the filter. When operating at a relatively high Q, the filter will likely be the limiting factor for linearity in the receiver so the noise and linearity analysis provided will be important when completing the system design of a sub-sampling receiver. This chapter also provided an algorithm for tuning the filter directly in both frequency and Q, which is a mandatory step if the filter is to be operated at a high Q.

Chapter 7

SUB-SAMPLING RECEIVER IMPLEMENTATION

7.1 Introduction

This chapter will detail the design of an entire analog receiver using the sub-sampling architecture described in previous chapters. First, the receiver is described on a high level, including a description of the targeted radio standard and application. A block and level breakdown is made where the gain, noise and linearity budget for the receiver is itemized such that approximate specifications can be derived for individual circuit blocks. A detailed description is given of the design of each block in the receiver. Finally, detailed experimental results for the test chips are shown and the overall receiver performance is summarized.

7.2 Receiver Description

To demonstrate the feasibility and benefits of a sub-sampling methodology for a low power wireless receiver, the analog portion of a receiver was implemented in 0.18 μ m CMOS technology. The receiver is intended for use in the 900 MHz ISM band as part of a general purpose, medium to low data rate wireless link.

A high level circuit diagram is shown in Figure 7-1 which illustrates the major functional circuit blocks. To complete the receiver, a digital baseband and a microprocessor is required as shown in the diagram. Input attenuators (1) and (4) are used for gain control. The first input attenuator is required for very large input signals as well as during the

filter tuning process, where the input signal could interfere with the filter tuning operation. The second attenuator is placed in front of the second filter (6) which operates at a high Q and therefore limits the receiver linearity.

The LNA (2) is the input transconductance for the first filter (3). The first filter is set at a nominal Q of 60-100 with minimal Q-enhancement. This provides some band selection for the receiver without causing a large degradation in linearity. Both filters are frequency tuned and Filter 2 is Q tuned using direct digital tuning [29]. The second transconductor (5) feeds current into the tank circuit of the second filter. The Q of the second filter is tuned to 500 to provide partial channel selection and corresponds to a 3dB bandwidth of 1.8 MHz. The buffer (7) provides further power gain by converting the second filter's output impedance to that of the sub-sampling mixer (8). In addition, during the Q-tuning of Filter 2, it can be switched to higher gain to allow the noise at the output of Filter 2 to be effectively sampled.

The sub-sampling mixer (3) converts the RF signal to a discrete-time IF at $\frac{1}{4}$ of the sampling frequency of 45 MHz. This IF signal is then further amplified by a discrete-time Programmable Gain Amplifier (PGA) (4) which completes the gain control of the receiver by allowing fine-tuning of the signal level at the input of the ADC. The ADC is a discrete-time Sigma-Delta ADC (5). The ADC operates on both I and Q channels in an interleaved fashion and is essentially a bandpass ADC.

The digital section of the system, not part of this test circuitry, includes the baseband modem (6) that completes the channel select filtering while simultaneously acting as the decimation filter for the Sigma-Delta ADC. This is also where the protocol specific

acquisition, gain control and demodulation would be implemented. A microprocessor (7) is included in the top level diagram as it is an important part of the tuning and control of the receiver. Low resolution DACs (9) control gain and frequency settings for all blocks in the receiver.

Figure 7-1: Top Level Diagram

7.3 Block and Level Breakdown

In order to design an example receiver, a set of wireless system requirements must be considered. The requirements below are based loosely on the IEEE 802.15.4 (Zigbee) standard [59]. They are as follows:

- Signal bandwidth of 1.2 MHz
- BPSK modulation
- Input sensitivity of -105 dBm.

- Interference at 5 MHz offset of 30 dBc and 10 MHz offset of 40 dBc

The required SNR is low because in this example the receiver uses direct sequence spread spectrum (DSSS) signaling. An ideal receiver requires 9.7 dB of SNR to decode BPSK with a bit error rate (BER) of 1×10^{-6} [22]. Some margin of interference / EVM needs to be set aside for the digital demodulator, so allowing for an SNR of 12 dB is reasonable. The actual bit rate of the demodulated signal will be 40 kbits / s after it is de-spread, while the symbol rate of the modulated signal is 600 kchips / s. This translates to almost 12 dB improvement in the SNR. Thus the resulting required SNR is 0 dB, considering the 1.2 MHz bandwidth. Using equation 3.2, a required noise figure of 9 dB is found.

Although the interference specifications are not difficult, the last requirement of the receiver is that it has low power consumption. It is intended for use in a remote, battery powered sensing application where battery life is critical. The power analysis in Chapter 3 showed a receiver with a current draw of 30 mA, yielding a required battery of over 4100 mA·hr capacity. The target for this receiver should be at most $\frac{1}{4}$ of this example, requiring a battery of approximately 1000 mA·hr for 3 years of life.

7.3.1 Gain and Noise

First, the gain and noise breakdown of the system is analyzed. Table 7-1 shows a rough breakdown of the receiver by block. G_v is the voltage gain, while G_p is the power gain of each block. When designing blocks on-chip, it is often more convenient to use the voltage gain and input referred voltage noise (V_{ni}) of the block, rather than the noise figure and power gain. To convert between these two quantities the impedance at the

interface between the blocks (Z_{in}) must be known. Below is a brief description of how each gain and noise estimate is achieved:

Attenuator 1: The input attenuator is included here because of its effect on the gain of the front-end, and therefore the overall noise figure. The matching of the receiver to the antenna is done with the attenuator and LNA as a whole. In reality, the design of the LNA and input match will be done with the attenuator in its maximum gain mode. The attenuator noise figure is shown as zero, but is in fact included with the LNA.

Table 7-1: Gain and Noise by Block

	Z_{in} (Ohms)	G_v (dB)	G_p (dB)	NF (dB)	V_{ni} (μV_{rms})
Atten 1	50	-1.5	-1.5	0	0
LNA / Filter 1	50	26.5	11.1	4.5	0.3
Filter 2	1725	31	21.8	11.9	5
Buffer	14373	-1	11.8	4.4	5
Sampler	750	-1.5	0.3	27.3	20
PGA	500	23	23.0	31	25
ADC	500				250
Totals (input referred):		73	63	6.3	

LNA / Filter 1: The LNA along with the attenuator is matched to a 50 Ohm external reference. The output impedance of the LNA is dictated by the Q of the tuned load, which is in turn set by Filter 1. Given a Q of 60, the equivalent narrow band output impedance for the inductor, L at resonance is found using [40]:

$$R_{out} = R_p = \omega L \cdot Q \quad (7.1)$$

For a center frequency of 915 MHz, an inductor value of 5 nH can be assumed, leaving a capacitance of 6 pF. The capacitance will be tuned to get the proper center frequency so the overall capacitance value must be sufficiently larger than the total parasitic capacitance in the tank circuit in order to allow tuning. The output impedance can then be calculated as 1.7k Ohm. The voltage and power gain of the LNA / Filter 1 combination respectively is given by:

$$G_v = g_m R_p \text{ and } G_p = g_m^2 R_p \cdot R_i \quad (7.2)$$

So the only remaining variable to be determined is the input transconductance of the LNA (g_m). The transconductance of a differential pair can be estimated using the equation:

$$g_m = \sqrt{2\mu_{0n}C_{ox} \frac{W}{L} I_D} \quad (7.3)$$

The LNA will be fully differential and consumes a large percentage of the receiver current budget. If a total current of 1mA is allocated, each side will have a drain current (I_D) of 0.5 mA. Giving a rough sizing of $W/L = 75 \text{ um} / 0.18 \text{ um}$ and the rough guide of 3.6×10^{-4} for $\mu_{0n}C_{ox}$, a g_m value of 12.2 mS is obtained. This allows calculation of the voltage and power gain for the LNA / Filter 1 combination. The noise figure of 4.5 dB was chosen as a goal, given the overall requirement of -105 dBm sensitivity which requires a receiver noise figure of less than 9 dB. Note that any losses in the attenuator and input match will be directly added to the overall noise figure.

Filter 2: The Q of filter 2 is almost 10 times higher than filter 1. The output impedance, calculated in the same manner, is over 14 kOhms. The result is that the input transconductance does not have to be as high in order to achieve a high voltage gain. Noise is not much of a concern because there has been some power gain in the first stage. For this reason, less current is budgeted for this filter. Using a current of 100 uA per side in the filter 2 input stage, and a transistor sizing of 20 um / 0.18 um, a g_m of 2.8 mS is obtained and gives a voltage gain of over 30 dB. There is a tradeoff between the voltage gain of Filter 2 and the linearity of the receiver overall. The high Q and impedance of Filter 2 will require the use of a precise tuning algorithm and methodology as described in Chapter 6. Note that the effects of the second attenuator are included in the design of Filter 2.

Buffer: The buffer provides further power gain by providing an interface between the high impedance of Filter 2 and the relatively low impedance of the sampler. The voltage amplitude in the buffer is shown as a slight drop because the buffer is implemented as a source follower. The power gain is then given by:

$$G_p = G_v^2 \frac{R_o}{R_i} \quad (7.4)$$

Where the input impedance (R_i) is set by the tank circuit of Filter 2 and the output impedance (R_o) is set by the design of the sub-sampling mixer or sampler. The output impedance of the buffer is combined with the series resistance of the sampler in order to determine the bandwidth of the sampler. This means that the derivations of noise figure for the sampler in section 5.2 would have to include the output impedance of the buffer.

Sampler: As stated, the sampler bandwidth should be approximately 1.3 times the RF frequency, in this case 915 MHz. A second constraint is the desired voltage noise equivalent from the sampler (or noise figure), which will be determined entirely by the size of the sampling capacitor through kT/C noise. Since a larger capacitor will require a lower output and switch impedance from the buffer (and therefore higher power consumption), the sampling capacitance should ideally be just large enough. The basic relationship that should be considered was discussed in Chapter 5, and is repeated here:

$$\overline{V_n^2} = \frac{kT}{C_s} \cdot \frac{f_{BW}}{f_s/2} \quad (7.5)$$

For this receiver a sampling frequency, f_s of around 50 MHz is chosen so that the PGA and ADC can be discrete-time and of reasonable power consumption. The signal bandwidth, f_{BW} is 1.2 MHz as stated above. For this analysis a capacitance of 200 fF gives a noise level of 15 μV_{rms} (where Table 7-1 shows 20 μV_{rms} , leaving some margin for capacitance variation). To obtain the required bandwidth, the total output impedance of the buffer plus the switch resistance is 660 Ohms. If 150 Ohms is allocated for the switch and 500 Ohms for the buffer, the g_m of the buffer works out to 0.002 Siemens. This can be obtained with a reasonably sized buffer and 100 μA of current.

PGA: The programmable gain amplifier (PGA) can be implemented as a sampled domain (switched-capacitor) gain stage since the output of the sampler is a discrete time signal. An estimated impedance is given in Table 7-1. The sampler can act as the input sampling switch for the PGA stage. The noise budget for the PGA is then allocated for the input equivalent noise of the opamp, output switches and capacitor.

ADC: The input noise of the ADC can then be permitted to be higher than the PGA because of the gain of the PGA. The ADC will contribute kT/C noise and noise from the opamps, just like the PGA, but will also contribute quantization noise due to the finite error in converting the analog signal to the digital domain. This total noise could be much higher than the PGA at the maximum gain setting, but because the quantization noise level will be fixed, it can be designed such that at the minimum AGC gain setting the noise of the ADC will dominate the receiver. For that reason, the noise specification for the ADC is the same as that of the PGA.

7.3.2 Linearity and Dynamic Range

To assess the contribution of each block in terms of its linearity, it is important to include the effect of filtering. If the model calculated thus far is taken, and using the gain and Q settings of filter 1 and 2, a picture of the signal and interference levels at each block can be created. An input signal applied near the receiver sensitivity and an interferer applied just below the receiver's interference blocking specification can be propagated through the gain and filter specifications to obtain the signal levels at each block input. Figure 7-2 shows what an incoming signal at -105 dBm and an incoming interference signal at a 10 MHz offset and -60 dBm would be at each block input. Note that P_{in} refers to the input power, while P_{int} refers to the power of the interfering signals. P_{int} and P_{in} are shown as voltages at the input of each receiver block. The third line plotted (yellow) shows the ratio of interfering power to input power. At the antenna the ratio is 45 dB, but because of the filtering of filter 1 and particularly filter 2, the ratio at the sampler through

ADC is only 20 dB. This permits much relaxed specifications for the dynamic range of these blocks.

Figure 7-2: Input and Interfering Signal Levels at each Block Input

Table 7-2: Gain and IIP3 Breakdown by Block

	Z _{in}	G _v	IIP3	IIP3
	(Ohms)	(dB)	(V)	(dBm)
Atten 1	50	-1.5	2	-41.0
LNA / Filter 1	50	26.5	0.3	-57.4
Filter 2	1725	31	0.04	-90.3
Buffer	14373	-1	0.3	-82.0
Sampler	750	-1.5	0.2	-72.7
PGA	500	23	0.5	-63.0
ADC	500	0	5	-43.0
Totals		76.5	0.0051	-58.8

The linearity of each receiver block can be analyzed, as was done with the gain and noise figure above. Table 7-2 shows the voltage gain and input impedance again along with the input referred IP3 of each block. The input IP3 of each block is an estimate as to what is possible with the planned system and chosen circuit architecture. The overall linearity (Input IP3 or IIP3) can then be calculated at the receiver's input. The value of -62.9 dBm is a result of the following formula [9]:

$$\frac{1}{V_{IP3,tot}^2} = \frac{1}{V_{IP3,1}^2} + \frac{G_{v1}^2}{V_{IP3,2}^2} + \frac{G_{v1}^2 \cdot G_{v2}^2}{V_{IP3,3}^2} \dots \quad (7.6)$$

Where the voltage gain and IIP3 voltage for each stage is used to calculate the equivalent input referred IP3 for the receiver. With this value and the overall noise figure of the receiver, the spurious free dynamic range can be estimated (SFDR) [9]:

$$SFDR = \frac{2(P_{IP3} + 174dBm - NF - 10\log(bandwidth))}{3} + SNR_{min} \quad (7.7)$$

Where the terms above are all in dB or dBm. The result for this receiver breakdown is 33.3 dB. This does not meet the interference requirements set for this receiver, but does not yet take into account the filtering in the receiver. It refers to the SFDR of the receiver if both interfering tones are in the signal bandwidth.

To see the effect of the Q-enhanced filtering, signals at defined offsets must be used to calculate the total inter-modulation distortion induced at each stage. At this point only the IP3 voltage of each stage has been specified thus; one has to work back from the definition of IP3 to estimate what the inter-modulation levels of each stage will be. A

receiver stage is defined as having the following relationship between input and output signals (ignoring non-linear terms greater than 3):

$$V_{out} = \alpha_1 V_{in} + \alpha_2 V_{in}^2 + \alpha_3 V_{in}^3 \quad (7.8)$$

It can be shown [9] that the IP3 is given by:

$$V_{IP3} = \sqrt{\frac{4}{3} \left| \frac{\alpha_1}{\alpha_3} \right|} \quad (7.9)$$

So to find the total inter-modulation values for each block, equation 7.9 can be rearranged as:

$$|\alpha_3| = \frac{4}{3} \frac{|\alpha_1|}{V_{IP3}^2} \quad (7.10)$$

At this point it is necessary to define the input signal used to test the SFDR of the receiver. The input signal can be expressed as:

$$V_{in} = V_1 \cos(\omega_1 t) + V_2 \cos(\omega_2 t) \quad (7.11)$$

For the purpose of the SFDR test, the amplitude of the two interfering tones is equal so the term V_0 can be used for the signal amplitude ($V_1 = V_2 = V_0$). The frequencies of the interfering signals are such that the third order inter-modulation lands directly in band. This can be expressed as:

$$2\omega_1 - \omega_2 = \omega_0 \quad (7.12)$$

Where ω_0 is the center of the channel and ω_1 and ω_2 are the interfering signal frequencies. Expressing the interfering frequencies as a delta from the channel frequency is more convenient, giving the following relationship:

$$2(\omega_0 - \Delta\omega_1) - (\omega_0 - \Delta\omega_2) = \omega_0 \longrightarrow 2 \cdot \Delta\omega_1 = \Delta\omega_2 \quad (7.13)$$

Thus when testing the interference performance at a 10 MHz offset, interferers at 10 MHz and 20 MHz offsets are applied, allowing the inter-modulation product to fall exactly in the center of the channel. Now equation 7.11 is inserted into 7.8 and expanded to find the term that contains $2\omega_1 - \omega_2$. The result can be shown to be [9]:

$$V_{out}(2\omega_1 - \omega_2) = \frac{3}{4} \alpha_3 V_1 V_2 \quad (7.14)$$

With equation 7.14 the amount of inter-modulation product that is generated at each receiver stage can be found. This is important when there is filtering present in certain stages, because the simple IP3 cascade calculations would not account for the reduction in signal level due to the filtering (see explanation in Chapter 3).

It is easier to analyze the inter-modulation products in tabular format, using a spreadsheet. Table 7-3 shows an example calculation. The input power is shown at the top of the table. This input power is used to calculate the voltage at the input of the first block (ATTEN) (1). The signals are then cascaded using the block gain specification, G_v . In the case of Filter 1 and Filter 2, the gains are recalculated with the specified Q and frequency offset by using equation 6.3. The α_3 parameter is calculated using the specified IIP3, the gain (α_1) and equation 7.10. Using this value, along with the input

voltages at each block input, the amount of inter-modulation at that stage can be calculated using equation 7.14. This value is then referred to the input of the receiver using the cumulative gain up to that block. All the input referred inter-modulation terms (expressed as a power, with respect to the 50 Ohm input impedance) can then be added to give a total overall inter-modulation power. The SFDR is found by changing P_{in} until the $P_{intermod}$ term is equal to the noise level at the sensitivity point. The SFDR is then defined as $P_{in} - P_{min}$.

Table 7-3: Inter-modulation Calculation

				Pin = -60.0 dBm		Intermod	
	Input IP3	G_v (α_1)	α_3	Vin (10 MHz)	Vin (20 MHz)	Vin ($2\omega_1 - \omega_2$)	Input referred
	(V)	(V/V)	(V/V ³)	(dB-V)	(dB-V)	(dB-V)	dBm
Atten 1	2	0.8	0.3	-77.0	-77.0	-244.6	-230.1
LNA / Filter 1	0.3	@ ω_0 : 21.1 @ ω_1 : 12.9 @ ω_2 : 7.6	312.9	-78.5	-78.5	-188.2	-200.1
Filter 2	0.04	@ ω_0 : 35.5 @ ω_1 : 3.3 @ ω_2 : 1.7	29568	-52.3	-54.6	-72.3	-115.3
Buffer	0.3	0.9	13.2	-28.9	-36.5	-74.3	-116.3
Sampler	0.2	0.8	28.0	-29.9	-37.5	-70.8	-111.3
PGA	0.5	14.1	75.3	-31.4	-39.0	-66.7	-130.2
ADC	5	1.0	0.1	-8.4	-16.0	-60.7	-124.2
total input referred $P_{intermod}$							-105.8
P_{min} due to inter-mod							-105.8
SFDR at 10 MHz offset							45.8 dB

Note that one simplification used above is that all the interference sources have been added in power. In reality these signals are not completely uncorrelated since they essentially come from the same input signals. It would be difficult to determine the exact interaction between all the phase and inter-modulation distortion products; however they

likely add constructively in some cases and destructively in others. It seems a reasonable estimation to consider them as independent, uncorrelated sources. For the above example, it would matter little since there is one very dominant source of inter-modulation distortion: Filter 2.

7.3.3 Image Rejection

Images will appear at offsets of $f_s/2$ from the signal frequency, in the case of this particular receiver is approximately 25 MHz. Images will be attenuated by Filter 1 and Filter 2 responses. By multiplying the relative attenuation of Filter 1 and 2 at the image offsets, the image rejection of this receiver can be predicted. Figure 7-3 shows the response of the filters and the attenuation of the image signals. In this example, the center frequency is chosen as 922 MHz. Since the filters are tuned for each channel in the band, the image rejection will remain constant for all channels. Also, because the sampling frequency was chosen to be greater than 2 times the width of the band of interest (in this case the 900 MHz ISM band is 22 MHz wide), the first image lands outside of the signal band. This fact is desirable as it would be expected that similar devices on different channels could be a significant source of interference. In this example, the first image frequency is attenuated by approximately 38 dB. Images that lie further out could be even further attenuated by filtering of the input match, antenna or external band-selection filtering.

Figure 7-3: Image Rejection of Receiver Implementation

7.4 Detailed Implementation

In this section the circuit design of each block is discussed in more detail. In most cases, the circuits are not in themselves novel and the design techniques to optimize such circuits are well known, thus some of the more detailed design methodology is not discussed here. However the structures that allow the sub-sampling receiver architecture to work are a focus of this discussion.

7.4.1 Programmable Bias Levels

Throughout the receiver design, bias levels are made programmable. Figure 7-4 shows an example simplified current bias circuit that has 8 bit programmable current. This allows for further optimization of the receiver performance when the circuits are being tested in the lab. Also, the attenuators used in the receiver use a programmable bias, referenced with a resistor so that their attenuation can be varied by a digital control as well. The automatic gain control loop can then be made part of the digital portion of the

receiver with a controller that adjusts the attenuators and PGA and possibly other bias controls in order to optimize the signal level for reception.

Figure 7-4: Programmable Bias Used Throughout Test Chips

7.4.2 Input Match and Attenuator

An input attenuator is required for the receiver front end for two reasons. First, while tuning the second filter to a Q of 500, any input signals must be attenuated sufficiently so that they do not interfere with the tuning operation. Second, the AGC algorithm for the receiver requires that the gain before the second filter be adjusted to ensure the signal is within its linear range. Both attenuators consist of series and shunt NMOS switches. Attenuator 1 is shown in Figure 7-5. The shunt switch connects the inputs of the LNA to AC ground (V_b) while the series switches (V_a) provide further attenuation by creating a high impedance path to the input of the LNA.

Figure 7-5: Circuit Diagram of Input Attenuator

7.4.3 LNA / Filter 1

The LNA is a standard cascode, common-source amplifier with a tuned output load. Matching for the LNA is done off-chip through the use of a filter-balun and single-ended input match. A lower cost but also lower performance alternative for matching the LNA would be to use a balun-less differential input match. The LNA uses inductor degeneration to improve the input matching [49].

Figure 7-6: LNA and Q-enhanced Filter

Both Filter 1 and Filter 2 make use of off-chip inductors that provide higher Q than the on-chip inductors, especially so at 900 MHz. Given that g_m must be set close to g_0 , the higher Q of the off-chip inductors results in a lower required g_m . Typical off-chip, low cost inductors have a Q of approximately 25 at these frequencies. A 5nH inductor with that Q gives an approximate equivalent parallel resistance of 700 Ohms. To reach oscillating condition ($g_m = g_0$) would require a transconductance of 1.4 mS. If transistors $M_{3a,b}$ are sized with a W/L of $20\mu\text{m}/0.18\mu\text{m}$, this would require a current less than 50 μA .

In the case of the first filter, the LNA is the input transconductance. The output load of the LNA is Filter 1. Filter 1 is frequency tuned precisely using the method described in Chapter 6. The Q of Filter 1 is set only slightly higher than the nominal Q to a Q of 50. This provides some band selection and attenuation of far out interferers. The Q cannot be made much higher at this stage because the overall linearity of the receiver would suffer (recall that the dynamic range of the filter is proportional to $1/Q^2$). It would be very

difficult to maintain a higher Q value for larger input signals as the first attenuator can only realistically achieve an attenuation of 20-30 dB.

Q tuning of Filter 1 is more difficult because of the lower Q setting (and therefore higher bandwidth). The tuning algorithm suggested in Chapter 6 would not be suitable because in order to see a significant drop in the noise level at the output of Filter 1, the receiver will need the ability to tune to a very wide frequency range. It would also be more difficult to achieve the necessary NNR for Q tuning since the filter is essentially part of the LNA. Instead, the first filter can be set at a fixed bias value in the Q -enhancement as process and temperature variation could be accounted for in the design, also because of the low Q setting.

7.4.4 Filter 2

Before reaching the second filter, the signal goes through a second attenuator. This attenuator provides further AGC range to the receiver allowing larger signals to be received, while maintaining the high Q setting of Filter 2.

The second filter uses a smaller input transconductance than the first, g_{mi2} as a low noise figure is not required after the gain of the first filter. Also, the overall voltage gain will be much higher because the Q is set to 500 in Filter 2. This translates to an output impedance 20 dB higher than that of Filter 1. The gain difference is balanced somewhat by a reduction in the input transconductance.

A Q of 500 in Filter 2 results in a 3dB bandwidth of 1.8 MHz. This gives the receiver partial channel selection (Note: the final channel selection and demodulation would be done digitally in the baseband processor). The signal spectrum at the output of Filter 2

will have greatly reduced interference amplitudes and therefore the subsequent receiver stages will not require as high a dynamic range. This is of particular importance for the sub-sampling mixer.

7.4.5 Frequency Control DAC

Rather than the varactors of a conventional RF tank circuit, the Q-enhanced filters use a bank of switched capacitors to control center frequency. Capacitors and switches provide higher Q, larger tuning range and better noise rejection. The capacitors are vertical capacitors, created using the fringing capacitance between adjacent metal lines. Eliminating the requirements for a special capacitor structures in the technology allows the receiver to be implemented in a standard digital CMOS technology. These capacitors have been shown to have higher capacitance density and better matching than their parallel plate equivalent [14].

The circuit diagram of the capacitor DAC is shown in Figure 7-7. PMOS transistors are used to switch one side of the binary weighted capacitor from high impedance to Vdd. The switch resistance while turned on will determine what the overall Q of the capacitor bank will be and thus the switches must be made large enough so that they do not reduce the overall tank circuit Q significantly. On the other hand, a larger switch size means that the parasitic capacitance seen by the tank circuit, even when the switch is off, is larger. This reduces the tunability of the tank circuit by reducing the on versus off capacitance.

Figure 7-7: Frequency DAC

The minimum step size will be determined by the ability of the system to deal with an offset from the ideal filter center frequency. Filter 2 is high Q and can add to the EVM seen by the receiver if significant distortion is seen in band. This can be exacerbated if the filter is off-centered significantly. For this design an offset target of 300 kHz was chosen, although it should be noted that the digital baseband implementation in the receiver would ultimately determine what this maximum offset could be. An alternative to reducing the minimum frequency step would be to increase the bandwidth of the filter in order to reduce the increase in EVM. The minimum step size is achieved by the use of small (approximately 5 fF) capacitors.

It is desirable to have the capacitors as far from the substrate as possible to minimize parasitic capacitance to the substrate or nwell. Although metal 6 is the top metal, metal 4 and 5 are chosen since the minimum design rules are smaller. The capacitors found in Figure 7-9 are surrounded by a guard ring to improve the matching of capacitor banks. The ring is left floating in order to minimize any parasitic capacitance. The capacitors

are connected in one direction by metal 3 and in the other by metal 5. The layout of the small capacitors can be viewed on the left side of Figure 7-9.

The switch transistors add a parasitic capacitance to the bottom plate of the capacitor of approximately 2.5 fF, meaning the tunability of the capacitors is about 66% (the difference between the on and off capacitance). To achieve this, the PMOS transistors were sized with a W/L of $3\text{ }\mu\text{m} / 0.18\text{ }\mu\text{m}$.

A competing requirement for the switches in the capacitor bank is the required Q of the filter tank circuit. By choosing off-chip inductors, the Q of the tank circuit can be made as high as 30, assuming the inductor dominates the tank circuit losses. That means that the Q of the capacitance bank should be at least this high or higher. With the switch sizes chosen above and a 5 fF capacitor, the Q of the capacitor can be calculated to be approximately 43. Since the Qs add in a manner similar to parallel resistors, this will reduce the inductor Q of 30 down to approximately 18. The expense of this lower Q will be essentially a higher current consumption in the Q-enhanced filters as well as higher noise, although the filter should not be a dominant noise source. The only way to counter this effect would be to increase the switch size and thereby reduce the tuning range.

An improvement can be made to the circuit shown in Figure 7-7. Figure 7-8 illustrates a different switch configuration whereby a single PMOS switch is placed between the two capacitors. The center of this transistor is a virtual ground node, and turning this switch on and off has the same effect of connecting and disconnecting the capacitance from the tank circuit. The benefit gained here is that the switch resistance is half as much as it was in the previous design while the parasitic capacitance remains the same. The two extra

switches to V_{DD} shown are required to ensure that the bottom plates of the capacitors are biased at the proper voltage. These switches are of minimum size so the overall parasitic capacitance is still lower with this design. The test chip designed for the RF front-end used the design shown in Figure 7-7, but the second design would provide a better tuning range and tank circuit Q tradeoff.

Figure 7-8: Improved Switch Configuration for Frequency DAC

Larger banks of capacitors are used to switch between broad frequency ranges. These capacitors have a higher density than the small fringing capacitors designed above since they use all six layers of metal. These capacitors are designed so that there is sufficient overlap between the ranges to account for the fact that the two capacitor types are not precise multiples of each other. The larger unit capacitor uses metal layers 1 through 5 for the inter-digitated fingers and metal 6 for interconnection. This capacitor is approximately 50 fF and is used in banks of 8. To minimize the effect of etching irregularities, the fingers in both unit capacitors are designed such that the parallel edges of the two plates are not in the same direction for the entire capacitor.

Figure 7-9: Unit Capacitors for Frequency DAC

7.4.6 Sub-sampling Mixer

A sampler or sub-sampling mixer is a simple sample and hold circuit. The difference is that in this case, careful attention is paid to selecting an appropriate sampling capacitor and matching it to the overall system design, given the relationships derived in Chapter 5. The sampling frequency (f_s) is chosen such that the offset between the incoming RF signal and the closest harmonic of f_s is exactly $\frac{1}{4}f_s$. This means the output of the sampler will be a discrete-time signal centered at $\frac{1}{4}f_s$. This simplifies the conversion to baseband in the digital portion of the receiver. A limitation of previous implementations is the low RF isolation when the sampling switch is turned off [1]. At RF frequencies the routing and source to drain capacitance of the sampling switch can cause the switch off impedance to be lower than desired and the RF signal will disturb the sampled voltage. To overcome this limitation, a second sampling switch is added to provide feed forward input cancellation during hold mode as shown in Figure 7-10. The primary sampling switch consists of transistors M1-M4, where M1 and M3 sample the differential input signal and transistors M2 and M4 are used to reduce charge injection [15]. For the

system architecture presented above, the high Q RF filtering preceding the sub-sampling mixer reduces the linearity requirements and renders the proposed structure practical. The sampling capacitors C_H is sized at approximately 600 fF, including the routing and parasitic capacitances.

The secondary sampling switch consists of transistors M5-M8, which are identically sized to those in the primary sampling switch. The differential input signal to the secondary switch is reversed as compared to the primary switch and the clock signal (dclk) is pulled low so that the secondary switch is always turned off. During track mode the clock signal (clk) is turned on and the input signal is placed on the hold capacitors (C_H). During hold mode the clock signal is turned off and both the primary and secondary sampling switches are disabled. Since the two switches have opposite polarity inputs, any differential input signals that pass through the sampling switches will be cancelled. Combined with differential circuitry, the RF isolation and common mode rejection is limited by component matching.

Figure 7-10: Circuit Diagram of Sub-Sampling Mixer

To achieve maximum RF isolation, it is essential that the signal paths through the primary and secondary sampling switch have identical parasitic capacitance and impedances. This is accomplished using: (1) inter-digitated fingers and a common centroid layout for the sampling switches and their interconnecting signals, (2) vertical capacitors using interdigitated fingers in metal layers 4 and 5 with symmetric edges to minimize edge effects and (3) matched clock drivers on sampling switches.

Figure 7-11 has a single-ended version of the sub-sampling mixer with the output buffer of Filter 2 and the output buffer to the PGA included, although the implemented circuit is pseudo-differential. The buffer used to drive the programmable gain amplifier (PGA) is a

pseudo-differential PMOS source follower. The signal at this point is at a much lower frequency than the sampler input, so the power required to drive the PGA input is not high, even though in maximum gain mode the PGA input capacitance is large.

The tank circuit from Filter 2 is directly connected to the source follower input, NM₁. The benefit of this is that the capacitance due to the source follower input is resonated out by inductors in Filter 2. In other words, the input capacitance of the source follower NM₁ is essentially part of the Filter 2 tank circuit. The source follower NM₁, combined with the sub-sampling mixer switch NM₃ form a low pass pole with the sampling capacitor, C_H and circuit parasitic capacitance. This pole should be sufficiently higher than the RF frequency to allow for process variation. To a first order this pole can be approximated with the following equation:

$$\omega_{3dB} = \frac{1}{\left(\frac{1}{g_{m1}} + R_{on3} \right) \cdot C_{samp}} \quad (7.15)$$

where g_{m1} is the transconductance of the input source follower, R_{on3} is the on resistance of the sampling switch and C_{samp} is the total capacitance seen at the node V_s . This capacitance includes parasitics in the sampling switches and coupling capacitor (C_C) as well as the input capacitance of the output source-follower, PM₂.

Figure 7-11: Sub-sampling Mixer with Input and Output Buffers

The sampler noise figure is given by equation 5.30. For this design, the sampling frequency is approximately 45 MHz and the sampling capacitor is 600 fF. Therefore the noise figure of the sub-sampling mixer is approximately 27 dB. Note that this includes the effects of the input buffer whose noise is also sub-sampled by the mixer. The noise of the source follower is equivalent to the noise of a resistor with the same output impedance [15] so the overall noise of the input buffer, switch and capacitor will still obey the kT/C relationship. All that is required to size the switch and source follower is to ensure that the combined bandwidth is enough to meet the gain requirements for the sub-sampling mixer. The noise of the output source follower is of much less concern since it is not sub-sampled and it is sized smaller to drive the PGA input capacitance at a much lower frequency.

The clocks used to drive the sub-sampling mixer came from an off-chip signal source. A self-biased inverter was used to create a square wave signal from an AC coupled

sinusoidal input. The same clock input was then used to create non-overlapping clocks for the PGA and ADC following the sub-sampling mixer.

7.4.7 Programmable Gain Amplifier

Since the output of the sampler is a discrete-time signal centered around $\frac{1}{4}$ of the sampling frequency, it is straight forward to design the remainder of the receiver using switched capacitor circuits. For the PGA, a capacitive-reset gain circuit was designed using a switch capacitor configuration shown in Figure 7-12 [51]. The clock signals (1 & 2) are standard, non-overlapping switched-capacitor clocks. In the first phase, the input signal is sampled and amplified by the ratio of the input capacitor C_1 over the integrating capacitor C_2 . In the second phase the input capacitor is reset while the output signal is held at the same level by switching the capacitor C_2 to the op-amp input and disconnecting C_3 from ground. The size of C_3 is not that important since it is only used to hold the preceding voltage level during the second clock phase. The gain has 4 bits of programmability with gain ranging from 0 dB to 23 dB. The gain of the PGA is equal to C_1/C_2 and is adjusted by varying the input capacitance of the PGA using logic to derive the clock signals for a bank of capacitors.

The benefit of a gain circuit like this is that the output does not need to slew from zero to the full amplitude of the signal, but remains at the last output signal level. This benefit is perhaps not as critical for the sub-sampling receiver since a signal centered at $f_s/4$ will often switch from full amplitude to zero between consecutive samples.

Figure 7-12: PGA Circuit Diagram

Both the PGA and ADC design are entirely switched-capacitor based, meaning that the output impedance of the op-amps does not need to be low. They will drive only capacitive loads so the output impedance can be high, as long as there is sufficient transconductance in the op-amp to meet the bandwidth requirements. This variety of op-amps are sometimes called operational transconductance amplifiers (OTAs). The gain accuracy of the PGA is not that critical so an opamp gain of 20 dB greater than the maximum PGA gain is more than sufficient. The unity gain bandwidth of the op-amp should be approximately 5 times the sampling frequency as a rule of thumb to allow the op-amp to settle in all situations. A single-stage, folded cascade op-amp is sufficient for this situation. Since the opamp is fully differential, a switched capacitor common feedback circuit is included to maintain the correct DC bias.

The noise generated by the PGA is dominated by two sources, the switched capacitor noise (kT/C noise) and the equivalent noise of the opamp. The switching of capacitors C_1 and C_2 add equal amounts of noise to the circuit since the gain of the PGA is C_1/C_2 and

the noise at the capacitor C_2 is multiplied by the gain in order to obtain the total input referred noise. C_2 is sized as small as possible to keep the size of C_1 reasonable and as a result, C_3 is made the same size as C_2 . The equivalent output noise of the op-amp should be divided by the gain of the opamp. The result is that the op-amp noise is more of a factor at the lower gain setting, where the noise specification on the PGA is not as stringent. The overall noise of the PGA can be estimated by the following equation:

$$v_{ni}^2 = 3 \cdot \frac{kT}{C_1} + v_{nopamp}^2 \cdot \frac{C_2}{C_1} \quad (7.16)$$

where v_{ni} is the equivalent input noise of the PGA and v_{nopamp} is the output referred op-amp noise. The first term accounts for the kT/C noise of capacitors C_1 , C_2 and C_3 while the second term accounts for the input referred op-amp noise. As with the sub-sampling mixer, the noise is reduced by the over-sampling ratio in the PGA. This noise reduction will be realized in the digital baseband portion of the receiver where the filtering and decimation will remove any noise not falling in the signal band. Thus, equation 7.16 can be modified to include the over-sampling ratio (OSR) as:

$$v_{ni}^2 = 3 \cdot \frac{kT}{C_1} \cdot \frac{1}{OSR} + v_{nopamp}^2 \cdot \frac{C_2}{C_1} \cdot \frac{1}{OSR} \quad (7.17)$$

The PGA makes use of the same lateral fringing capacitances used in the frequency DAC, shown in Figure 7-9 on the right hand side. The capacitor is approximately 50 fF and this sets the value of C_2 and C_3 . C_1 is then roughly 750 fF and the input referred noise due to the capacitors is equal to approximately 14 μ Vrms. This means that the input referred op-amp noise can be as high as 20 μ Vrms and the PGA will still meet the

block and level requirements laid out in Table 7-1. This is easy to achieve in an op-amp design, particularly since the signal is a band-pass signal and the $1/f$ noise of the op-amp transistors will not be a significant contributor to the noise seen in the signal band.

The linearity of the PGA is of much less concern because of the filtering done by Filter 2. Because it is a closed-loop amplifier it is not difficult to achieve high linearity and resulting IIP3. All switches that are not virtual ground or bias are made CMOS to ensure good linearity with large signal inputs. The IIP3 budgeted for the PGA in Table 7-2 is easily accommodated with this design.

7.4.8 ADC

The signal entering the ADC has not been completely filtered for close interference and thus it must still be over-sampled at this point. The remainder of the channel selection filtering is done in the digital domain, which places a requirement on the ADC to handle the bandwidth and dynamic range dictated by the interference requirements. Although a Nyquist rate ADC (e.g. pipeline or flash) could be used to do the data conversion, it would be wasteful in terms of the conversion efficiency. A Nyquist rate ADC converts all frequencies up to the Nyquist frequency ($f_s/2$) and the SNR is constant over that spectrum. Because the signal must remain over-sampled, a $\Sigma\Delta$ ADC is a practical choice. The $\Sigma\Delta$ noise shaping makes use of the over-sampling to increase the dynamic range of the ADC. The noise spectrum of a $\Sigma\Delta$ ADC is not flat over the Nyquist band, but is “shaped” so that more noise appears at frequencies outside of the signal band. A digital filter is then used to remove this noise.

The target for interference rejection of the receiver at 10 MHz offset is 40 dB. By the time an interfering signal at that offset reaches the ADC, Filter 2 will have attenuated it by 20 dB. Similarly, at 5 MHz offset it is attenuated by 15 dB. This would indicate that with sufficient receiver gain, the ADC should only require 20 dB of dynamic range. In practice much more margin is required to allow the receiver to operate properly. The Analog Gain Control (AGC) steps will be coarse and thus the signal will need to be placed higher than the ADC noise floor by some margin. Also, the signal itself may be much larger in some cases even if there is no interference and the ADC will need to be able to handle large signals.

The ADC for this receiver should target a dynamic range of 45 dB. This should allow some flexibility in the AGC design as well as allowing large signals to be converted.

$\Sigma\Delta$ ADC Theoretical Operation

The $\Sigma\Delta$ ADC is a natural choice for data conversion in a sub-sampling receiver. Due to image rejection requirements in the sub-sampling down conversion, the discrete time IF signal will be a bandpass signal centered at one quarter of the sampling frequency ($f_s/4$).

The critical block of the $\Sigma\Delta$ ADC is the filter. In combination with the feedback loop, the filter provides noise shaping, creating a notch in the noise transfer function corresponding to the signal band. The “noise” that is shaped by the $\Sigma\Delta$ ADC is actually the error caused by the fact that one is digitizing an analog signal. This noise is commonly referred to as “quantization noise” [55].

The basic block diagram of a $\Sigma\Delta$ ADC is shown in Figure 7-13a. This figure shows a single bit, discrete time $\Sigma\Delta$ ADC. The input signal is $X(z)$, the output is $Y(z)$ and the loop filter transfer function is $H(z)$. The result is the addition of quantization noise to the signal. A common approximation for this phenomenon is to assume the quantization noise is a white noise source ($Q(z)$) [15]. This approximation does not work well in situations where the input signal is a DC level or a pure tonal signal. In the case of a radio receiver there is generally a lot of randomness to the signal, both from noise and from the nature of the digitally modulated signal. For this reason the white noise approximation is a good one.

Figure 7-13: Single Bit, Discrete Time $\Sigma\Delta$ ADC Block Diagram

By analyzing the diagram, the signal transfer function (STF) and noise transfer functions (NTF) in the Z domain will be can be derived as:

$$STF(z) = \frac{Y(z)}{X(z)} = \frac{H(z)}{1 + H(z)} = \frac{N(z)}{D(z) + N(z)} \quad (7.18)$$

$$NTF(z) = \frac{Y(z)}{Q(z)} = \frac{1}{1 + H(z)} = \frac{D(z)}{D(z) + N(z)} \quad (7.19)$$

Where $H(z)$ has been represented as $N(z)/D(z)$. The feedback structure causes the poles of $H(z)$ to become zeros in the NTF. So if the poles of $H(z)$ are chosen so that they land in the signal band, they will create notches in the noise seen in the output signal $Y(z)$. The simplest example of this is the first order low-pass modulator. The first order (low-pass) loop filter is a single integrator shown in Figure 7-14. The STF and NTF with this loop filter respectively are:

$$STF(z) = z^{-1} \quad (7.20)$$

$$NTF(z) = 1 - z^{-1} \quad (7.21)$$

This causes the signal at the output to be delayed by one period and the noise will have a zero at DC ($z = e^{j\omega T} = 1$ when $\omega = 0$). Similarly, higher order filters can be used to create more zeroes in the NTF. For low pass ADCs these zeros are all placed at DC for first and second order filters, while for orders higher than 3, complex zero pairs can create notches at other frequencies [56]. For the general low-pass ADC where all zeros are at DC, the NTF will have the form:

$$NTF(z) = (1 - z^{-1})^L \quad (7.22)$$

where L is the order of the modulator. Several different values of L are plotted in Figure 7-15 where the x-axis is denoted as the ratio of the input frequency to the sampling frequency, f_s .

Figure 7-14: First Order Loop Filter **Figure 7-15: NTF for Different Orders of Low-pass ADCs**

7.4.8.1.1 Lowpass to Bandpass transformation

To create a bandpass $\Sigma\Delta$ ADC, the loop filter, $H(z)$ will simply have poles at the frequency of interest. The simplest form of this involves a transformation from the lowpass version where $-z^2$ is substituted for z . For example, the transformation from the first order lowpass would be:

$$H_{BP}(z) = H_{LP}(-z^2) = \frac{-z^{-2}}{1+z^{-2}} \quad (7.23)$$

This results in signal and noise transfer functions:

$$STF(z) = -z^{-2} \quad (7.24)$$

$$NTF(z) = 1 + z^{-2} \quad (7.25)$$

which produces zeros in the noise transfer function where $z = \pm j$ or in other words, the frequency is one quarter of the sampling frequency ($f_s/4$). The resulting ADC would be

considered a second order bandpass modulator. A more general transformation can be used to create zeros at any point in the sampled frequency domain between 0 and $f_s/2$, but in the case of the sub-sampling receiver, the optimum NTF is as shown in equation 7.25. As with the low-pass version, higher order band-pass loop filters can be constructed. The NTF for different orders of bandpass ADCs is shown in Figure 7-16.

Figure 7-16: NTF for Different Orders of Band-pass ADCs

7.4.8.1.2 Theoretical Maximum SNR for low pass $\Sigma\Delta$ ADC

$\Sigma\Delta$ ADCs are often compared in terms of their maximum SNR, which is often quoted in units of “bits” using the fact that each bit of resolution is approximately equivalent to 6 dB of SNR [15]. With a $\Sigma\Delta$ ADC the actual bits are not obtained until after a decimation filter. In the case of the sub-sampling receiver this decimation filter is implemented as part of the digital baseband portion of the receiver. The dynamic range of the ADC on the other hand should be considered as the range of input amplitudes over which the ADC maintains a certain minimum SNR [53]. This definition is more useful for a receiver, but it is dependent on the specific receiver specifications because the minimum SNR is

needed. In this receiver example the minimum SNR is 0 dB and the signal gain of the ADC is 0 dB, so the maximum SNR of the ADC in this case is a reasonable estimate of the usable dynamic range of the ADC. The dynamic range is likely higher because the SNR does not drop immediately to zero when compression first occurs, but it is difficult to determine what the effect of distortion induced tones will have on the signal at high levels of compression.

The theoretical maximum SNR for an ADC can be derived for any order L and any number of bits, n. Up until this point, it has been assumed that the quantizer is 1-bit and consists of 1 comparator. Also, the band-pass ADC is twice the order of the equivalent low-pass ADC so the term L here refers to the low-pass equivalent order. The resulting maximum SNR can be shown to be [57]:

$$SNR_{\max} = \frac{3 \cdot 2^{(2n-1)} \cdot (2L+1)}{\pi^{2L} OSR^{(2L+1)}} \quad (7.26)$$

and this relationship is plotted below in Figure 7-17. The plot for L=0 refers to the case where there is no loop filter at all and the increase in SNR for higher sampling rates is due solely to over-sampling. Also plotted in Figure 7-17 is the intersecting point where the over-sampling ratio of this design lies. If a first order ADC was used, the maximum SNR would almost be enough but a safer choice is a second order ADC where an over-sampling ratio of 18.75 gives us a maximum SNR of almost 60 dB.

Figure 7-17: Maximum Theoretical SNR for 1-bit ADCs

7.4.8.2 Pseudo Two Path ADC

A standard second order low-pass architecture, converted to a band-pass architecture, is shown in Figure 7-18 [54]. The signal and noise transfer functions are as follows:

$$STF(z) = -z^{-4} \quad (7.27)$$

$$NTF(z) = (1 + z^{-2})^2 \quad (7.28)$$

Figure 7-18: Block Diagram of 4th Order Bandpass ADC

The architecture chosen for the ADC was a pseudo-two path structure, similar to that devised in [20] and is shown in Figure 7-19. The main difference between this and the block diagram shown in Figure 7-18 is that I and Q signals have been split apart into two

signal paths. They are not yet converted to DC, but are separated in the time domain between discrete samples.

The pseudo two-path structure in [20] had the benefit of reusing the same integrators for both I and Q paths, thereby limiting mismatch between the paths and effectively creating a band-pass modulator, centered at $f_s/4$. The only difference in this implementation is that the input modulator at $f_s/4$ is actually implemented in the sub-sampling mixer. This is accomplished by switching the sense of the clk and $dclk$ signals every two clock cycles which requires fewer switches and clock phases. Each path shown in Figure 7-19 is a 2nd order low-pass loop filter. The actual resulting ADC block diagram is as shown above in Figure 7-18. Although the ADC integrators could be implemented directly as band-pass integrators, as shown in equation 7.23, this transfer function is difficult to implement and requires additional op-amps. The pseudo two-path architecture is implemented with integrators that are low-pass and the capacitance ratios are very simple. The resulting loop filter requires only two opamps, although there are twice as many output capacitors as a low-pass ADC since the I and Q paths remain separate.

Figure 7-19: Conceptual Block Diagram of Pseudo-Two Path ADC

The circuit diagram of the integrator used for the ADC is shown in Figure 7-20. In the first clock phase, the input is sampled on the capacitor C_{in} and in the second clock phase

Figure 7-20: ADC Integrator Circuit Diagram

The noise in the ADC will come from the kT/C noise of the switched capacitors and the integrator op-amp noise as well as the quantization noise left over after decimation in the digital baseband circuitry. The kT/C and op-amp noise is referred to as the circuit noise, while the quantization noise is a result only of the digitization process. In other words, even with perfectly ideal circuits, the quantization noise is still present. The circuit noise can be analyzed in a very similar way to the PGA. In this case the noise of the first integrator is the sole concern since the second integrator noise is actually shaped by the filtering of the first integrator. The quantization noise is ultimately determined by the

choice of the voltages v_{ref+} and v_{ref-} which roughly determine the maximum input power of the ADC. Using voltages of 1 V and 600 mV (analog ground is defined as 800 mV) yields a peak signal swing differentially of 400 mV. Converted to an rms quantity and reducing the value by the maximum SNR of 60 gives an rms noise voltage of 250 μ V. The kT/C and op-amp noise can easily be kept well under this number. An input and integrating capacitor of 200 fF and 400 fF respectively and an op-amp noise of 20 μ Vrms yields a total circuit noise of less than 30 μ Vrms. Thus the ADC's contribution to the receiver noise is heavily dominated by the quantization noise floor. For this reason the reference voltages are made programmable in the design. It is then possible to tradeoff the maximum signal the ADC can handle for a lower ADC noise floor. A full schematic simulation result of the designed ADC is shown below in Figure 7-21.

Figure 7-21: Simulated ADC Response with a 40 MHz Sampling Clock

7.5 Experimental Results

Two test chips were implemented in 0.18 μm CMOS technology. Micro-photographs of the test chips are shown in Figure 7-22 and Figure 7-23 and the blocks included in each test chip are shown in Figure 7-24.

Figure 7-22: Microphotograph of Test Chip 1

Figure 7-23: Microphotograph of Test Chip 2

The first chip is composed of the RF front-end and includes the blocks up to the second filter (including the output buffer) and has an active area of 1.6 mm^2 . The second test chip was implemented with the sub-sampling mixer, PGA and ADC with an active area of 0.22 mm^2 . The chips were packaged in a chip-on-board package and a custom test circuit board was designed for each chip. To facilitate driving the RF signal at the output of the second filter between the two chips, an output source-follower was biased off-chip with an external resistor.

Figure 7-24: Receiver Blocks Included in Test Chips

7.5.1 Test chip #1: RF front-end

The board design for the first test chip included a single-ended to differential balun for matching to the receiver.

7.5.1.1 Frequency DAC results

A limitation was quickly realized with the first test chip design. The Q of the frequency control DAC was found to be much lower than anticipated. It was discovered that the reason was due to the layout of the unit capacitors and switches and the resulting routing between the two. Although the switch resistance was low enough to give a good Q for the frequency DAC, the routing resistance added too much in series with each capacitor, causing the overall capacitor Q to be much lower than expected. This was exacerbated by the fact that very large capacitor banks were added as a contingency for tuning range. These extra capacitor banks were connected with a very narrow route of metal so that they could be “fused” out if required. The end result was that with these extra capacitor banks connected, even if they were turned off, the Q of the overall tank circuit was much lower than expected. The low Q of the parasitic capacitance on these capacitors was a large contributor to on the overall tank circuit Q. It was found through measurement that the Q of the filter tank circuits without enhancement was approximately 11. With this filter design, Filter 2 required a bias current of 1 mA to achieve the desired enhanced Q of 500. This current was much more than expected for Filter 2.

An experiment was done to remove some of the large “contingency” capacitor banks using a laser. Also, the frequency was tuned to the high end of the frequency range to reduce the effect of the low Q capacitor banks by switching off the majority of them. The

result was that the Q of the tank circuit without enhancement was over 17. This is much closer to the expected Q of 18 predicted by the frequency DAC design. Filter 2 was again used to achieve an enhanced Q of 500 and the required bias current in the filter was now only 360 μA .

A second problem was realized when testing the frequency DAC of this test chip. The step size requirement can be met with a unit capacitor of 5 fF quite easily, but uncertainty in matching and routing capacitance makes it very difficult to maintain that step size if the capacitor banks are sized using binary weighting, as shown in Figure 7-7. The step size was plotted for the first 8 bits of the frequency DAC codes and is depicted in Figure 7-25. The non-linearities at code boundaries are much larger than the required step size of 300 kHz. To fix this problem, the bank of small capacitors should be designed as a thermometer encoded bank. If all 8 bits were done this way, the digital decoding block and routing would be very large, so the first 2 bits could be binary weighted as they were in this test chip, with the final 6 bits thermometer encoded. This would result in 64 capacitors that are each comprised of 4 unit capacitors and each have their own control bit. This improvement would both eliminate the non-linearities seen in Figure 7-25 and it would improve the capacitor Q by reducing the routing required between the switch and the capacitor unit. The larger capacitor banks are designed such that there is a significant overlap in between the total capacitance of the entire small capacitor bank and the smallest step size of the large capacitor bank. This ensures that any difference or change in the two capacitors will not cause any frequencies within the tuning range to be missed.

Figure 7-25: Frequency DAC Step Size

To summarize, the following changes and improvements are necessary to make the frequency control DAC and resulting receiver front-end functional:

- Reduce the series routing resistance in the capacitor bank by using wider and shorter metal routes
- Use the improved switch design of Figure 7-8 to reduce the on resistance of the switch
- Use thermometer encoding for 6 out of the first 8 bits in the frequency DAC

With the first two changes, the capacitor bank should be able to maintain a tank circuit Q of at least 18 over the entire tuning range since the dominant loss in the capacitor bank will be the series switch.

7.5.1.2 Input Match

The maximum filter oscillation frequency was 1.06 GHz with 5.7 nH inductors. The input was later changed to 920 MHz for full receiver testing, but due to of the issues with

the frequency control DAC described above, the input match and performance of the RF front-end was optimized closer to the upper end of the tuning range. This allowed testing the receiver at the desired power consumption level for the filters. The input match was therefore optimized for approximately 1.04 GHz. The result of the input match of the receiver is shown in Figure 7-26 and the input return loss was lower than 10 dB over a 40 MHz band. This was accomplished using only the single-ended matching components at the antenna side of the balun. A broader band match would likely be possible by working with the differential match between the balun and the chip and would perhaps reduce the losses seen by the balun due to the fact that the balun does not see 200 Ohms at the differential port.

Figure 7-26: Receiver Input Match

7.5.1.3 LNA bias and Receiver Sensitivity

Although the input match varied slightly with the bias current in the LNA, it was possible to find the optimal power consumption in the LNA by varying the bias and measuring the noise figure of the RF front-end. It was found that increasing the bias current of the LNA past 1 mA provided little benefit to the overall receiver noise figure. This is as expected given the system and LNA design outlined earlier.

Figure 7-27: Receiver Noise Figure versus LNA Bias

7.5.1.4 RF Front-end Transfer function and Power Consumption

The remainder of the circuit blocks are biased as expected from design, with some adjustments and optimizations similar to that which was done with the LNA. Both filters were tuned manually to the desired Q. The overall gain of the RF front-end was measured using a network analyzer and is shown in Figure 7-28. Given the block and level breakdown completed above, a gain of 55 dB for the receiver front-end is expected. The test setup had losses in the input (due to the cables, connector and balun) that would

amount to approximately 6 dB and the output path of the test board, including the output buffer, balun, connectors and cables, would account for approximately 10 dB. Thus, the results presented in Figure 7-28 would show an in-channel gain of 53 dB, which is very close to the expected value.

Figure 7-28: Composite Transfer Function of Receiver RF Front-end

The total current breakdown after LNA bias optimization and Q tuning of the filters is shown in Table 7-4. The total operating current was just under 3 mA and the supply voltage used was 1.6 V. This equates to a power consumption of 4.7 mW.

Table 7-4: RF Front-end Current Consumption Breakdown

LNA	1.0 mA
Filter 1	120 μ A
G_{mi2}	370 μ A
Filter 2	360 μ A
Rfamp	500 μ A
Bias & DACs	410 μ A
Total	2.94 mA

7.5.1.5 *Linearity and Dynamic Range*

The linearity of the RF front-end was tested by injecting two tones at specified offsets into the receiver. First, the two tones were injected at the center of the bandwidth of the filter with a separation of 200 kHz. These tones and the inter-modulation products see little or no filtering from the receiver. The power of the injected tones was varied from close to the input sensitivity level up to a point where the 3rd order inter-modulation products (IM3) were very large. The resulting data is plotted in Figure 7-29. The output power is uncalibrated, but the input power in the plot is calibrated to include the assumed losses in the test setup. The input sensitivity of the receiver is -106 dBm and this corresponds to a noise figure of 8 dB. Note that this is also the point shown in Figure 7-27 where the LNA bias is set to 1 mA. The output noise that determines this sensitivity is shown as a straight line across Figure 7-29. The SFDR can be found for this data by finding the input point at which the IM3 power line crosses the “P_{noise}” line and subtracting the input sensitivity point from this. This works out to 27.8 dB for in-channel interferers. This is less than was expected from the block and level breakdown where 33 dB was expected. The likely cause is a higher than expected level of non-linearity in Filter 2 where the high Q-enhancement is done. Since the initial block and level budget had some margin, the design is sufficient despite the shortfall. This number can be improved by enabling the degeneration options designed into Filter 2, but the cost is increased power consumption. By considering interference at specified offsets, the performance criteria for this receiver can be met without this modification.

Figure 7-29: In-band IM3 plot

The same test is repeated at various frequency offsets with the results outlined below in Table 7-5. As expected, the SFDR increases rapidly with the offset frequency of the interfering signals. Although the quoted input IP3 in-band would seem very low, it is more significant to focus on the IP3 and SFDR at frequency offsets where there will be interference that can cause inter-modulation.

Table 7-5: RF front-end Linearity at Different Frequency Offsets

Frequency Offset (MHz)		IIP3	SFDR
F ₁	F ₂		
0	0	-65.0	27.8
2	4	-44.8	41.3
4	8	-36.7	46.7
20	40	-14.5	61.5

The in-band input IP3 was also tested at the maximum attenuation of the first input attenuator. When the first attenuator is biased so that the attenuation is maximum and the

LNA bias is turned down to zero, the input signal is attenuated by 57 dB. This brings the IIP3 of the receiver up to -8 dBm. Hence for interferences at 2 and 4 MHz away, the receiver can accommodate signals over 40 dB higher than the input signal, over a range of input signals from -106 dBm to -49 dBm (i.e. the interfering signals are at -9 dBm). This naturally assumes that the automatic gain control (AGC) algorithm can always settle to the optimal gain setting. In reality it will be necessary to include a margin in the algorithm to place the incoming signal above the noise floor, resulting in less dynamic range for higher input signals.

7.5.2 Test-chip #2: Sub-sampling Mixer, PGA and ADC

The second test chip included the sub-sampling mixer, PGA and ADC blocks. The input to the sub-sampling mixer included the source follower and bias as shown in Figure 7-11. The input to this test chip was matched on-chip with a resistor. This allows a reasonable approximation as to the voltage at the input of the sub-sampling mixer, even though parasitic capacitance in the board and chip bonding can alter the impedance to a certain extent.

Two versions of the PGA/ADC were designed using two different op-amps. One opamp used was a two stage op-amp with greater than 60 dB of gain. Slow settling time limited this op-amp to a maximum frequency of about 35 MHz. The second implementation used a single stage op-amp as described above that was able to operate up to 60 MHz, but due to a design error in the clock distribution the input modulator could not be used. Therefore the second version of the design could be demonstrated in low-pass mode only. Even with the lower gain of a single stage op-amp, it is possible to come very close to the

theoretical maximum SNR for the ADC. Test input and output points were made accessible between the sub-sampling mixer and PGA and between the PGA and ADC. The input paths for the test board used are shown in Figure 7-30. For the input of the sub-sampling mixer, an RF balun was used with a 4:1 impedance ratio. The resistor match is then placed on the differential side of the balun. A large, low frequency balun was required for the input to the PGA and ADC that had a sufficiently low minimum frequency so that the PGA and ADC could be tested in low-pass mode. It is prudent at this point to recall that the input modulator of the sub-sampling mixer is required in order for the chip to be used in band-pass mode. In each test input, the balun was used to provide the correct DC bias at the test block input. In the case of the PGA and ADC, the bias was not necessary if the signal was driven from the preceding stage.

Figure 7-30: Test Chip #2 Simplified Test Board Schematic**7.5.2.1 Sub-sampling Mixer Isolation**

The sub-sampling mixer was tested on its own to evaluate the linearity and the RF isolation. A screen capture in Figure 7-31 from a spectrum analyzer shows the difference in output power between the sub-sampling mixer being in the “off” state and the “on” state. The isolation is over 50 dB, implying that the matched input paths are effective at cancelling the RF signal that is routed to the output through parasitic paths.

Figure 7-31: Sub-sampling Mixer Isolation**7.5.2.2 Sub-sampling Mixer Gain and Noise Figure**

The gain of the sub-sampling mixer was extracted from the test results by comparing the relative output level at the ADC when a signal is input at the RF port. A signal of -15 dBm input to the ADC test port, accounting for losses, yields a relative output at the ADC of -30 dB. The same signal input to the RF port yields a relative output at the ADC

of -30.9 dB. Given that the gain of the PGA is 20 dB during this test and the input losses to the sub-sampling mixer are estimated at 15 dB, the gain of the sub-sampling mixer is estimated to be -5.9 dB. This is a crude estimate, but gives an idea of what the voltage gain is. It is lower than desired from the block and level diagram, partly because of the use of two source followers in the design. The voltage gain could be improved by using an additional op-amp in the output path rather than a source follower, or possibly integrating the sub-sampling mixer into the input stage of the PGA. Also, it is expected that an integrated solution placing Filter 2 with the first source follower would give less voltage loss.

The sub-sampling mixer noise was estimated by comparing the output spectrum at the ADC with and without the sub-sampling mixer and PGA enabled. The noise is expected to be higher for both cases because a lower sampling frequency was used due to the design error with the single-stage op-amp design. The sampling frequency used in this test was 20 MHz and the two noise spectra are shown in Figure 7-32. The ADC noise spectrum would rise much later if the correct sampling frequency of 45 MHz were used, but the noise density due to the sub-sampling mixer can still be extracted by looking at the spectrum below 400 kHz. This spectral density can then be integrated to a bandwidth of 1.2 MHz to get an idea of the noise introduced by the sub-sampling mixer. This noise works out to an equivalent RF input level of -78 dBm or 26 μ Vrms. Note that this again takes into account the estimated 15 dB of loss in the test board and matching. The noise figure of the sub-sampling mixer then works out to 34 dB. Although this is higher than

the design criteria, the lower sampling frequency accounts for 7 dB of this increase and thus the noise figure is very close to what was expected.

Figure 7-32: Noise Spectrum with Signal at RF Input Compared to the ADC only with a 25 MHz Sampling Frequency

7.5.2.3 PGA gain

The gain of the PGA was measured in a relative manner by comparing the output signal between the different programmable gain settings. Figure 7-33 shows a comparison of the theoretical gain as determined by the designed capacitor ratios and the actual measured gain. The discrepancy is due to the finite gain of the op-amp in the PGA, in this case the single stage op-amp. The low gain of the op-amp causes a gain error in the closed loop PGA. The error is not that critical in the overall receiver design, although the maximum gain of the measured PGA is closer to 20 dB rather than the expected 23 dB.

Figure 7-33: PGA Theoretical and Measured Gain

7.5.2.4 ADC Signal to Noise Ratio

The ADC was evaluated using the single-stage op-amp version of the test chip which allow the sampling frequency to be as high as 60 MHz. As expected the output spectrum shows that the quantization noise of the ADC is dominant. This can be seen by the fact that the spectrum is not flat, but shows the Noise Transfer Function (NTF) down to frequencies well below the signal band. Most of the ADC testing was done using single or 2-tone sinusoidal signals. When testing the ADC alone, the output spectrum reveals idle tones. Figure 7-34 shows the ADC spectrum with a 2-tone input signal. The idle tones are the lower spikes both in the signal band and all the way up to the Nyquist frequency of $f_s/2$. These tones are due to the fact that the input signal does not have enough noise or random nature [55]. Compare this to Figure 7-35 where the signal was input at the PGA and the result was that more noise was seen by the ADC input. The idle tones are much reduced and the noise spectrum looks much closer to the ideal.

Figure 7-34: ADC Spectrum with 2-tone, 0dBm Input Signal to ADC

Figure 7-35: ADC Spectrum with PGA Input at High Gain

The SNR of the ADC was captured over a wide range of input signals and is shown in Figure 7-36. Also shown is the Signal to Noise plus Distortion Ratio (SNDR). This ratio includes distortion due to tones created by the input signal itself. The SNDR is lower than the SNR for many input levels, but much of this is due to the idle tones discussed above. Furthermore, as discussed earlier, in terms of a radio receiver, the actual SNR or SNDR of an ADC is not that important. The key parameter is the dynamic range of the ADC which is actually the range of input signals over which the SNR is greater than the $\text{SNR}_{\min}$ dictated by a particular radio standard. As can be seen in Figure 7-36, the dynamic range is well over the required 45 dB and if the SNR plot were extended to lower signal levels it would work out to over 60 dB.

Figure 7-36: ADC SNR and SNDR

7.5.2.5 Complete Test chip #2 results

The overall gain of the blocks in test chip #2 is lower than intended from the original block and level design laid out in Table 7-1. This does not have a significant effect on the overall receiver noise figure when the maximum gain setting is used. If the gain of the receiver is lowered through the use of Automatic Gain Control (AGC), it is likely that the noise from the analog baseband blocks (e.g. the ADC) will begin to dominate. The lower gain of the sub-sampling mixer and PGA will mean that the ADC noise will be dominant sooner (i.e. at a higher AGC setting) than it would have been if the gain target was met. This may not be an issue for the receiver, but will depend on the details of the digital demodulation and AGC algorithm. A higher gain in the PGA and less loss in the

sub-sampling mixer will enable the receiver to be a more general purpose design and will certainly improve the dynamic range of the receiver at lower gain settings.

A possible way to improve the output noise when dominated by the ADC is to lower the ADC reference voltage. The quantization noise level in the ADC is set by the voltage level used as a feedback reference. Reducing the reference voltage level is equivalent to adding voltage gain to the signal in the ADC. The downside to this is that the large signal limit of the ADC follows the reference voltage as well and the ADC cannot handle as large a signal.

The large signal distortion of the blocks in test chip #2 is dominated by the sub-sampling mixer, as predicted. The PGA and ADC inter-modulation distortion were tested and found to be sufficiently high for the sub-sampling mixer to be the dominant source, as expected from Table 7-2. The input IP3 of the PGA was measured at 1.8 V. It is more difficult to get an IP3 number for the ADC since the nature of the Sigma-Delta modulator is such that tones are not generated at multiples of the input signal frequency until the signal is very large. On the other hand, Figure 7-36 shows that the SNDR is higher than required for very large input signal levels (0 dBm corresponds to 223 mVrms) so one can assume that the input IP3 is sufficiently high. When all 3 blocks are put together, the IP3 is lower and is dominated by the sub-sampling mixer. Figure 7-37 shows the ADC output spectrum when a two-tone, -20 dBm signal (22.3 mVrms each tone) is input to the sub-sampling mixer. The third order tone can be seen at approximately -65 dB, while the output level is -20 dB, relative to the ADC reference voltage. Extrapolating these results out gives an input IP3 for the sampler that is over 200 mVrms; well above the target set

out in Table 7-2. Also seen in Figure 7-37 is the image of the input signal at a relative amplitude of approximately 40 dB. Mismatch in the ADC I and Q paths, essentially just the I and Q integrating capacitors, is responsible for this image. This image poses no threat to the receiver because it is self generated and will be guaranteed to be much lower than the input signal itself, in this case 20 dB. More careful layout and capacitor matching could potentially make this as high as 30 dB if it were necessary.

Figure 7-37: ADC Spectrum with -20 dBm (22.3 mVrms) Sub-sampling Mixer Input

Total current consumption for test chip #2 was 1.87 mA from a 1.6 V supply voltage. Using the noise figure and IP3 found above, dominated by the sub-sampling mixer, the SFDR for this test chip works out to 55.8 dB. This should be sufficient for use with the RF front-end presented in the first test chip, particularly since much of the interference

will be filtered by the time the signal reaches the sub-sampling mixer. The extra dynamic range allows the AGC algorithm to more readily choose a gain setting that falls within the proper range of the analog baseband circuitry.

7.5.3 Combined Receiver Results

Operating the two test chips together is a trivial task once both have been shown to be functional. As stated earlier, the only difference in circuit design when integrating all blocks together would be the removal of the output buffer of test chip #1. The combined results of the entire receiver are dominated by the performance of the RF front-end, as expected in the system analysis. The sensitivity and instantaneous dynamic range or SFDR is essentially the same for the full receiver as it is for the RF front-end. Although not implemented in this project, the AGC algorithm would likely deal with the noise limitations of the analog baseband circuitry, particularly the quantization noise of the ADC. Table 7-6 summarizes the key parameters of the receiver. The overall power consumption is under 8 mW and the current consumption is less than 5 mA. This power consumption leaves enough margin in the specification of 7.5 mA to accommodate the synthesizer and digital baseband processor which were not included in this measurement.

Table 7-6: Summary of Receiver Key Parameters

Parameter	Measure
Frequency	900 – 920 MHz
Noise Figure	8 dB
Receiver Sensitivity	-106 dBm
SFDR	
2 MHz offset	41 dB
4 MHz offset	47 dB
Image Rejection	38 dB
Current / Power Consumption	4.81 mA / 7.67 mW

A random bit-pattern BPSK signal was used to show the potential of the analog receiver for use in a digital radio. The signal was coded differentially, meaning the actual bit or symbol is represented by the difference between two adjacent phase symbols. The signal was input at -90 dBm to the input of the receiver. Note that since the spread spectrum coding gain in this example is not realized, the minimum sensitivity of the receiver will be much higher than -106 dBm. Figure 7-38 shows a result where 6700 bits are captured and 40 bits are in error, leading to a bit error rate of 0.6%. This result would indicate the signal is very near the sensitivity point of the receiver, although the signal processing to recover the bits in this example is almost certainly not optimized. Some signal distortion is obvious due to the shape of the constellation and this is probably due to the band-limiting of the Q-enhanced filters, particularly Filter 2. The signal quality could be improved if a digital filter was constructed that compensated for this distortion. Despite this limitation, the example proves the potential of this analog front-end for use in a complete digital radio.

Figure 7-38: Example Received Symbol Constellation

7.6 Conclusions

The receiver design examined in this section clearly demonstrates how the sub-sampling receiver architecture can be used to enable low power operation. The combination of high Q filtering for anti-aliasing and sub-sampling combine to make this receiver achieve good performance at low power consumption. Furthermore, the architecture is not limiting in the bandwidth or modulation methodology as in the case of some architectures. The benefits of the sub-sampling architecture are realized in the implementation of a full receiver which has a current consumption of less than 5 mA and has a high dynamic range and good sensitivity.

Table 7-7 summarizes some of the specific examples and compares the power consumption to the receiver described and implemented in Chapter 7. It is always difficult to make a direct comparison between receivers, particularly when the data rates or bandwidth of the receivers and to some extent the RF frequency used can be different. However, the sub-sampling receiver does not restrict the bandwidth of a receiver implementation like that of an oscillator-based receiver (such as the super-regenerative receiver). The sensitivity and SFDR of the example in Chapter 7 would meet the Bluetooth standard and performs better than receivers presented in literature [48].

Table 7-7: Comparison of Receiver Power Consumption

Receiver	Architecture	Power Consumption	Notes
[46]	Direct Conversion	91 mW	Multi-band, includes 4 GHz
[48]	Low-IF	39 mW	Bluetooth
[44]	Super-regenerative	5 mW	Limited bandwidth, limited modulation method
This work	Sub-Sampling	9.17 mW	1.5 mW included for synthesizer; flexible in frequency and modulation method

Chapter 8

CONCLUSIONS AND RECOMMENDATIONS

The system designer of a radio must consider many criteria when selecting an architecture for use in a low power wireless device. In recent decades the trend for radio transceivers has been toward a high level of silicon integration for both power and cost motives. This has lead to the use of new receiver architectures to both facilitate this level of integration and to exploit it. The power consumption of the analog portion of the radio remains a concern for the system designer and much effort is put forward to reduce it. Furthermore, the power consumption of the receiver often dominates the overall consumption of a battery powered wireless device.

This thesis has explored the analog portion of a radio receiver with a focus on the goal of achieving low power consumption. The sub-sampling receiver architecture presented is a sound choice to achieve low power in a digital radio when compared to other common receiver architectures. In comparison to the direct conversion or Low-IF architecture, the sub-sampling architecture has eliminated the need for an RF synthesizer and quadrature signal paths at high frequency. These architecture components account for a large portion of the analog power consumption of the receiver and thus the sub-sampling receiver presents a significant power savings opportunity. Chapter 4 gives a comparison of some of the most common architectures with specific examples chosen from references.

The nature of sub-sampling allows more design flexibility, since the RF frequency does not determine the design of the frequency synthesizer and signal conversion stages as it does with the direct conversion or Low-IF architectures. The comparison in Table 7-7 shows that the sub-sampling receiver can operate at much lower power than direct conversion or Low-IF without sacrificing the performance.

The sub-sampling receiver is fundamentally limited in use by the anti-aliasing requirements due to image folding. The fundamental relationship between the Q of the band-pass anti-aliasing filter and the image rejection of the sub-sampling receiver was shown in Chapter 5. The implementation of the sub-sampling receiver in Chapter 7 is an example of how this relationship functions. To make this receiver realizable, Q -enhanced filters are utilized for anti-aliasing. The high Q desired for the sub-sampling receiver in turn dictates that the filters are tuned directly as described in Chapter 6.

The architecture presented in this thesis can be extended to many different wireless systems and its functionality possibly expanded upon. The use of a low frequency synthesizer in the sub-sampling architecture means that a large range of RF frequencies could be converted with very little modification to the overall circuit architecture. Multi-standard radios, where several frequency bands and/or wireless standards are addressed with a single radio transceiver would be good candidates for implementing a sub-sampling architecture with directly tuned Q -enhanced filters. Also, much higher frequency radios, well into the microwave frequency ranges could use the sub-sampling architecture as a secondary conversion stage, further extending the possibilities for this architecture.

The contributions set out at the beginning of this thesis have been covered and are summarized here:

- Low Power Receiver architecture: The sub-sampling receiver architecture with high Q anti-aliasing filtering presented has the potential to decrease power consumption compared to conventional architectures. The implemented receiver has less than $\frac{1}{4}$ of the power consumption of the best reported Bluetooth receiver.
- Anti-aliasing requirements for efficient Sub-sampling Receiver implementation: The fundamental requirements of anti-aliasing for a sub-sampling receiver dictate that the Q/n ratio is constant for a given required image rejection. This fundamental relationship governs the design of a sub-sampling receiver.
- Q-enhanced filter for Sub-sampling Receiver anti-aliasing: The Q-enhanced filter with direct digital tuning allows the sub-sampling receiver designer to use a high Q/n ratio. This is the key to implementing a sub-sampling receiver with a low frequency synthesizer and therefore low power consumption.
- Circuit design contributions related to a Sub-sampling Receiver: A novel frequency control for LC filters was shown in the frequency DAC design of the example receiver. The sub-sampling mixer design was made specifically for interfacing with the Q-enhanced anti-aliasing filter and exhibits high RF isolation in hold mode.
- Direct digital tuning methodology: The concepts outlined in Chapter 6 may be suitable for directly tuning a bandpass filter and allows the filter to operate at an

accurate high Q setting. The method also makes use of the sub-sampling receiver frequency plan to find and tune the center frequency of the anti-aliasing filter accurately.

8.1 Possibilities for Future Work

The research in this thesis leads to other interesting work involving low power sub-sampling receiver architectures. Some suggestions for further areas of research are:

- Full implementation of a 900 MHz radio: Design of a complimentary digital baseband and demodulator as well as an appropriate processor would complete a system on a chip. Implementation of filter tuning algorithms could be completed, yielding a self contained complete radio receiver.
- A sub-sampling receiver targeted for 2.4 GHz: The popularity of this frequency band is due to the world-wide unlicensed spectrum available. The sub-sampling architecture presented in this thesis would be well suited to implementation at this frequency and would compare quite favourably in power consumption to existing radio implementations. For example, a sampling rate of 120 MHz and a Q of 200 in a Q-enhanced anti-aliasing filter would be reasonable starting points, yielding a Q/n ratio of 10. The discrete-time baseband could be run at $\frac{1}{2}$ the sampling rate to save power by splitting the I and Q samples between two paths.
- Higher Frequency radios: Other radio bands may benefit with the sub-sampling architecture as all or part of the total receiver architecture. Different passive and active filtering structures could enable the use of sub-sampling in higher frequency bands, or

the sub-sampling could be used as a second down conversion stage in a microwave frequency radio

REFERENCES

- [1] Shen, D., Hwang, C., Lusignan B., "A 900 MHz RF Front-End with Integrated Discrete-Time Filtering," *IEEE Journal of Solid-State Circuits*, Volume: 31, Issue: 12, Dec 1996.
- [2] Kuhn, W.B., Yanduru, N.K., Wyszynski, S., "Q-Enhanced LC Bandpass Filters for Integrated Wireless Applications," *IEEE Transactions on Microwave Theory and Techniques*, vol. 46, no. 12, Dec. 1998.
- [3] Kuhn, W.B., "Design of Integrated, Low Power, Radio Receivers in BiCMOS Technologies," PhD Dissertation, Virginia Polytechnic Institute and State University, Blacksburg, VA 1995. (<http://scholar.lib.vt.edu/theses/public/etd-81197-164118/etd-title.html>)
- [4] Arkesteijn, V.J., Klumperink, E.A.M., Nauta, B., "Jitter Requirements of the Sampling Clock in Software Radio Receivers" *IEEE Transactions on Circuits and Systems – II: Express Briefs*, Vol. 53, No. 2, Feb. 2006.
- [5] DeVries, C., Mason, R., "A 0.18 μ m CMOS, High Q-enhanced Bandpass Filter with Direct Digital Tuning," *IEEE Custom Integrated Circuits Conference Proceedings*, 2002.
- [6] H. Ahmed, C. Devries, R. Mason, "RF, Q-enhanced Bandpass Filters in Standard 0.18 μ m CMOS with Direct Digital Tuning," *Proceedings of the IEEE International Symposium on Circuits and Systems ISCAS2003*, Bangkok, Thailand, May 2003.
- [7] H. Ahmed, C. Devries, R. Mason, "A Digitally Tuned 1.1 GHz, Subharmonic Injection Locked VCO in 0.18 μ m CMOS with Direct Digital Tuning.," 29th European Solid-State Circuits Conference ESSCIRC2003, Estoril, Portugal, September 16-18, 2003.
- [8] Abou-Allam, E., Manku, T., El-Masry, E.I., "Q-enhanced 1.9GHz tuned CMOS RF amplifier," *Electronics Letters*, vol. 32, n. 5, Feb. 1996.
- [9] Razavi, B., *Design of Analog CMOS Integrated Circuits*, McGraw-Hill, 2001.
- [10] Lee, T., *The Design of CMOS Radio-Frequency Integrated Circuits*, Cambridge University Press, Cambridge, U.K., 1998.
- [11] Herold, E.W., "The Operation of Frequency Converters and Mixers for Superheterodyne Reception", *Proceedings of the IRE*, vol. 30, Issue 2, Feb. 1942.
- [12] Holland, M.G., Claiborne, L.T., "Practical Surface Acoustic Wave Devices," *Proceedings of the IEEE*, Vol. 62, Issue 5, 1974.
- [13] Landau, H.J., "Sampling, Data Transmission, and the Nyquist Rate," *Proceedings of the IEEE*, Vol.552, Issue 10, 1967.
- [14] Aparicio, R., Hajimiri, A., "Capacity Limits and Matching Properties of Integrated Capacitors," *IEEE Journal of Solid-State Circuits*, Volume: 37, Issue: 3, Mar 2002.
- [15] Huehne, K., Lovelace, D., Ovalle, P., "A low power 900 MHz PLL frequency synthesizer for ISM applications using 0.25 μ m BiCMOS," *Devices, Circuits and*

- Systems, 2000. Proceedings of the 2000 Third IEEE International Caracas Conference on*, March 2000.
- [16] Xintian, Shi, et. al., "A Low-Jitter and Low-Power CMOS PLL for Clock Multiplication," *Solid-State Circuits Conference, 2006. ESSCIRC 2006. Proceedings of the 32nd European*, Sept. 2006.
 - [17] Raha, P., "A 0.6-4.2V low-power configurable PLL architecture for 6 GHz-300 MHz applications in a 90 nm CMOS process," *VLSI Circuits, 2004. Digest of Technical Papers. 2004 Symposium on*, June 2004.
 - [18] O'Neil, Peter V., *Advanced Engineering Mathematics – 3rd addition*, Jan 1995.
 - [19] Johns, D.A., Martin, K., *Analog Integrated Circuit Design*, John Wiley & Sons Inc., 1997
 - [20] Song, B., "A Fourth Order Bandpass Delta-Sigma Modulator with Reduced Number of Op amps," *IEEE Journal of Solid-State Circuits*, Volume: 30, Issue: 12, Dec 1995.
 - [21] Parssinen, A., *Direct Conversion Receivers in Wide-Band Systems*, Kluwer Academic Publishers, 2001.
 - [22] Proakis, J.G., *Digital Communications – Fourth Edition*, McGraw-Hill, 2001.
 - [23] Leon-Garcia, A., *Probability and Random Processes for Electrical Engineering – Second Edition*, Addison Wesley, 1994.
 - [24] Rogers, J., Plett, C., Dai, F., *Integrated Circuit Design for High-Speed Frequency Synthesis*, Artech House, 2006.
 - [25] Namgoong, W., Meng, T.H., "Direct-conversion RF Receiver Design," *IEEE Transactions on Communications*, Vol. 49, no.3, March 2001.
 - [26] Crols, J., Steyaert, M.S.J., "Low-IF Topologies for High-Performance Analog Front Ends of Fully Integrated Receivers," *IEEE transactions on Circuits and Systems – II: Analog and Digital Signal Processing*, vol. 45, no. 3, March 1998.
 - [27] Shahani, A.R., "Radio-Frequency Conversion and Synthesis (For a 115 Milliwatt GPS Receiver)", Ph.D. dissertation, Stanford University, March 1999.
 - [28] Ahmed, H., DeVries, C., Mason, R., "A Digitally Tuned 1.1 GHz Subharmonic Injection-Locked VCO in 0.18 μ m CMOS," *European Solid-State Circuits Conference Proceedings (ESSCIRC)*, 2003.
 - [29] DeVries, C., Mason, R., "A 0.18 μ m CMOS 900 MHz Receiver front-end Using RF Q-enhanced Filters," *International Symposium on Circuits and Systems Proceedings (ISCAS)*, 2004.
 - [30] DeVries, C., Mason, R., "An RF Sub-sampling Mixer, PGA and $\Sigma\Delta$ ADC for Conversion at 900 MHz," *International Symposium on Circuits and Systems Proceedings (ISCAS)*, 2004.
 - [31] Nauta, B. "A CMOS Transconductance-C Filter Technique at Very High Frequencies," *IEEE Journal of Solid State Circuits*, vol. 27, no. 2, pp. 142-153, Feb. 1992.
 - [32] Li, Dandan, Tsividis, Y., "A 1.9 GHz Active LC Filter with On-Chip Automatic Tuning," *International Solid State Circuits Conference Proceedings (ISSCC)*, pp 368-369, 2001.

- [33] Valdes-Garcia, A., Silva-Martinez, M., Sanchez-Sinencio, E., "RF Bandpass Filter Design Using Capacitive Degeneration," *Midwest Symposium on Circuits and Systems* (Proceedings of), 2005.
- [34] Pekau, H., Haslett, J.W., "A 2.4 GHz CMOS Sub-Sampling Mixer with Integrated Filtering," *IEEE Journal of Solid-State Circuits*, Vol. 40, No. 11, November 2005.
- [35] Hajimiri, A., Lee, T.H., *The Design of Low Noise Oscillators*, Kluwer Academic Publishers, 1999.
- [36] Shaeffer, D.K., "The Design and Implementation of Low Power CMOS Radio Receivers", Ph.D. dissertation, Stanford University, December 1998.
- [37] Capodiferro, M, Avitabile, G., Lofu, N., "A direct wireless signals receivers architecture with noise sources modeling", *European Conference on Wireless Technology* (Proceedings of), 2005.
- [38] Yuce, M.R., Liu, W., "Design and implementation of a multirate sub-sampling front-end in software radio systems", *IEEE Radio and Wireless conference* (Proceedings of), 2004.
- [39] Yuce, M.R., Liu, W., "A Low-Power Multirate Differential PSK Receiver for Space Applications," *IEEE Transactions on Vehicular Technology*, Vol. 54, No. 6, November 2005.
- [40] Kuhn, W.B., Yanduru, N.K., Wyszynski, S., "Q-Enhanced LC Bandpass Filters for Integrated Wireless Applications," *IEEE Transactions on Microwave Theory and Techniques*, vol. 46, no. 12, Dec. 1998.
- [41] Yu, Wei, Leung, B.H., "Noise Analysis for Sampling Mixers Using Stochastic Differential Equations," *IEEE Transactions on Circuits and Systems – II: Analog and Digital Signal Processing*, Vol. 46, No. 6, June 1999.
- [42] Dunham, W., Euler: The Master of Us All, MAA, 1999.
- [43] Frink, F.W., "The Basic Principles of Super-Regenerative Reception", *Proceedings of the Institute of Radio Engineers*, vol. 26, no. 1, January 1938.
- [44] Chen, J.Y., Flynn, M.P., Hayes, J.P., "A 3.6mW 2.4GHz Multi-Channel Super-Regenerative Receiver in 130nm CMOS", *Custom Integrated Circuits Conference Proceedings*, 2005.
- [45] Popplewell, P.H.R., et. al., "An Injection-Locked 5.2 GHz SoC Transceiver with On-Chip Antenna for Self-Powered RFID and Medical Sensor Applications", *IEEE Conference on VLSI Circuits, Proceedings of*, 2007.
- [46] Kaczman, D.L., et. al., "A Single-Chip Tri-Band (2100, 1900, 850/800 MHz) WCDMA/HSDPA Cellular Transceiver", *IEEE Journal of Solid-State Circuits*, Vol. 41, No. 5, May 2006.
- [47] Darabi, H., et. al., "A Dual-Mode 802.11b/Bluetooth Radio in 0.35-um CMOS", *IEEE Journal of Solid-State Circuits*, Vol. 40, No. 3, March 2005.
- [48] Marholev, B., et. al., "A Single-Chip Bluetooth EDR Device in 0.13um CMOS", *International Solid State Circuits Conference Proceedings (ISSCC)*, pp 558-559, 2007.
- [49] Lee, T., *The Design of CMOS Radio-Frequency Integrated Circuits*, Cambridge

University Press, Cambridge, U.K., 1998.

- [50] Wang, C.K., Castello, R., Gray, P.R., "A Scalable, High Performance Switched-Capacitor Filter," *IEEE Journal of Solid-State Circuits*, Vol. sc. 21, No. 1, Feb 1986.
- [51] Martin, K., Osoolak, L., Lee, Y.S., Temes, G.S., "A Differential Switched-Capacitor Amplifier," *IEEE Journal of Solid-State Circuits*, Vol. SC-22, no. 1, Feb 1987.
- [52] Boser, B.E., Wooley, B.A., "The Design of Sigma-Delta Modulation Analog-to-Digital Converters," *IEEE Journal of Solid-State Circuits*, Vol. SC-23, no. 6, Dec 1988.
- [53] Cherry, J.A., Snelgrove, W.M., Weinan Gao, "On the design of a fourth-order continuous-time LC delta-sigma modulator for UHF A/D conversion," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, Vol. 47, Issue 6, June 2000.
- [54] Ong, A.K., Wooley, B.A., "A Two-Path Bandpass $\Sigma\Delta$ Modulator for Digital IF Extraction at 20 MHz", *IEEE Journal of Solid-State Circuits*, Vol. 32, No. 12, Dec 1997.
- [55] Norsworthy, S.A., Schreier, R., Temes, G.C., *Delta-Sigma Data Converters – Theory, Design, and Simulation*, IEEE press, 1997.
- [56] Ferguson, P.F. Jr., Ganesan, A., Adams, R.W., "One bit higher order sigma-delta A/D converters", *IEEE International Symposium on Circuits and Systems*, Proceedings of, vol. 2, pp. 890-893, May 1990.
- [57] Brigati, S., et. al., "Design considerations on very high resolution sigma-delta modulators", *Proceedings of the 40th Midwest Symposium on Circuits and Systems*, Vol. 1, Aug. 1997.
- [58] Shaumann, R., Van Valkenburg, M.E., *Design of Analog Filters*, Oxford University Press, 2001.
- [59] 802.15.4-2003 PHY (Zigbee) specifications
- [60] Riley, T.A.D., Filiol, N.M., Du, Q., Kostamovaara, J., "Techniques for In-Band Phase Noise Reduction in $\Delta \Sigma$ Synthesizers," *IEEE Transactions on Circuits and Systems – II: Express Briefs*, Vol. 50, No. 11, Nov. 2003.
- [61] D. B. Leeson, "A simple model of feedback oscillator noise spectrum," *Proc. IEEE*, vol. 54, pp. 329–330, Feb. 1966.