

Copyright © 2000, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

**ADAPTIVE DIGITAL SIGNAL PROCESSING
ALGORITHMS FOR IMAGE-REJECTION
MIXER SELF-CALIBRATION**

by

Gabriel M. Desjardins

Memorandum No. UCB/ERL M00/65

9 August 2000

cover

**ADAPTIVE DIGITAL SIGNAL PROCESSING
ALGORITHMS FOR IMAGE-REJECTION
MIXER SELF-CALIBRATION**

by

Gabriel M. Desjardins

Memorandum No. UCB/ERL M00/65

9 August 2000

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

Adaptive Digital Signal Processing Algorithms for Image-Rejection Mixer Self-Calibration

By Gabriel M. Desjardins

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences,
University of California at Berkeley, in partial satisfaction of the requirements for the
degree of **Master of Science, Plan II.**

Approval for the Report and Comprehensive Examination:

Committee:

Professor P.R. Gray
Research Advisor

(Date)

* * * * *

Professor B. Nikolic
Second Reader

(Date)

For the purpose of this report, the following information is provided:

Information is provided in

the following manner:

The following information is provided in the following manner:

The following information is provided in the following manner:

Information is provided

Information is provided

Information is provided

Information is provided

Information is provided

*Adaptive Digital Signal
Processing Algorithms for
Image-Rejection Mixer Self-
Calibration*

by Gabriel M. Desjardins

Adaptive Digital Signal Processing Algorithms for Image-Rejection Mixer Self-Calibration

by

Gabriel M. Desjardins

This report details Digital Signal Processing (DSP) methods for optimizing the performance of a self-calibrating image-rejection mixer, as well as a custom VLSI implementation of one such algorithm. The mixer is part of a Double-Conversion Wide-Band IF receiver designed to support the Digital European Cordless Telephone (DECT) and the Global System for Mobile Communications (GSM) Standards, which have carrier frequencies of 1.9 GHz. Different numerical methods were evaluated according to computational complexity, convergence time and convergence accuracy. We developed a variable step-size adaptive algorithm to remove inherent mismatches in the mixer circuit. Various DSP architectures were also simulated with the minimization of circuit area and power consumption in mind. Simulation showed that the optimization process arrived at a phase mismatch of 0.0129 degrees and a gain mismatch of 0.0114% for an ADC output noise level of -76 dBV, as specified in the receiver frequency plan. This allows us to achieve a mean image-rejection ratio of 78 dB across a wide frequency range irrespective of variations with operating temperature. The mean convergence time of the algorithm is 590 μ s. The DSP system occupies less than 0.08 mm² of chip area in a 0.25 μ m CMOS process and has a pre-layout power consumption of 13 μ W for a 1.0V supply.

Acknowledgments

Regardless of how it seems, this project could not have been completed without the help of many other people. I would like to thank my research advisor, Professor Paul Gray, for his support throughout this project. I would also like to thank Professor Bora Nikolic for reading countless drafts of my thesis and Dr. Chris Rudell for answering an equally countless number of technical questions.

My time at Berkeley has not been spent entirely on schoolwork and research. But I owe a debt of gratitude to Don Baron, Andy Klein, Robert Wood and Brian Otis for reminding me to focus on academics and for leading by example with their supreme motivation for all things scholastic. I must also thank Johan Vanderhaegen, who solved all of my problems with Simulink and probably saved me two months of simulations. The Cal Ice Hockey team (is there really any other kind of hockey?) also helped me to retain my Canadian heritage, for which I will be eternally grateful.

Lastly, and most importantly, I would like to thank my parents for their years of love and devotion. They have always wanted the best for me and I dedicate this work to them.

Table of Contents

Introduction 1

1.1	Wireless Communications	1
1.2	Analog Receiver Section	2
1.3	Thesis Organization	3

Receiver Architectures 4

2.1	General Considerations	4
2.2	Receiver Architectures	5
2.2.1	Super-Heterodyne Receiver	5
2.2.1.1	The Image Frequency Problem	6
2.2.2	Direct Conversion Receiver	7
2.2.3	Image-Reject Receivers	8
2.2.3.1	Image-Rejection Architectures	8
2.2.3.2	Double Conversion Wide-band IF Receiver	9

Image-Rejection Mixers 12

3.1	Introduction	12
3.2	Image-Rejection Mixers	13
3.3	Image-Rejection in the Wide-Band IF receiver.	14
3.3.1	Non-Idealities of the Weaver Method	16
3.3.2	Self-Calibrating Image-Rejection Mixer	17
3.3.3	Performance Specifications	19

Image-Rejection Algorithms 21

4.1	Image-Rejection Mixer	21
4.2	Spectral Estimation	22
4.2.1	Computation of the Energy Spectrum	23
4.2.2	Discrete Fourier Transform	24
4.2.3	Estimation of Mixer Output Spectrum	24
4.3	Performance in the Presence of Noise	28
4.3.1	Analogy to Binary Digital Communications Receiver	28
4.3.2	Receiver Noise Specifications	29
4.3.3	Estimation of Mixer Output Spectrum	29
4.3.4	Algorithm Refinements	31
4.4	Image-Rejection Algorithms	32
4.4.1	Computation Time	33
4.4.1.1	Linear Search	33

4.4.1.2	Higher-Order Methods.....	34
4.4.2	Algorithm Complexity	36
4.4.2.1	Higher-Order Methods.....	36
4.4.2.2	Linear Search	37
4.4.2.3	Gear-Shifting	37
4.5	Time-Domain Simulations	41
4.5.1	Bessel Filter Group Delay	42

DSP Architecture 44

5.1	Custom DSP Implementation	44
5.2	DSP System Architecture	44
5.2.1	Discrete Fourier Transform Block Architecture	45
5.2.1.1	Arithmetic Refinements	46
5.2.1.2	Figures of Merit	47
5.2.2	Finite State Machine Block Architecture	48
5.2.3	Combined DSP Architecture.....	49

Conclusion 50

6.1	Conclusion	50
6.2	Future Work	51

References 52

1

Introduction

1.1 Wireless Communications

The wireless communications market has grown substantially during the last decade [1]. Recent advances in wireless technology have reduced the size and cost of mobile radios while improving performance. The increasing level of integration in wireless circuits has led to many of these improvements. However, increased integration is dependent on the development of novel transceiver architectures that allow the designer to eliminate many large discrete electronic components and to combine multiple circuit blocks on a single chip.

At the same time, numerous wireless standards have been introduced which dictate the performance specifications of the hardware in wireless devices. Hardware requirements differ substantially between wireless communications applications. As of 1998, for example, no cellular phone manufacturer in North America had a mass-production dual-mode (AMPS/CDMA) handset that met the requirements for both standards [2], despite the existence of numerous handsets that met specifications for

one of the standards. The implementation of multiple cellular standards in a single architecture also requires novel transceiver design approaches.

1.2 Analog Receiver Section

In most current designs, the analog part of a receiver uses multiple packaged and discrete integrated circuits which may have been implemented in CMOS, Gallium Arsenide, Silicon Germanium or Bipolar fabrication processes. Many higher-frequency discrete components are designed for a given frequency, while baseband components are designed for a particular bandwidth and modulation scheme specific to one standard. It is difficult to meet the requirements of multiple standards using discrete components and simultaneously reduce the size of a receiver.

Future wireless receivers will integrate all of their components on a single chip using an inexpensive process, such as CMOS. This has obvious advantages in terms of cost and size; integrated radios will also provide more functionality in terms of supporting multiple RF standards and optimizing power consumption and radio performance. One key to integration and supporting multiple standards is the effective reuse of hardware by each standard [3].

A multi-standard receiver was proposed in [4] which uses an image-rejection mixer to provide sufficient image rejection. In past, such designs have focused on careful circuit layout and matching in order to achieve maximum performance. This design features a different approach - in order to allow the system to support multiple standards, the mixer has the capacity for self-calibration, a concept which has been used in analog-to-digital converters[13][14][15]. In this case, the self-calibration procedure can be used to optimize the system performance over a wide frequency

range, such as across the entire channel bandwidth of a given standard. This improves the system performance for each standard and allows hardware reuse.

1.3 Thesis Organization

Chapter 2 provides an overview of different receiver architectures. The relative merits of these architectures are discussed with respect to integration and supporting multiple wireless standards. Chapter 3 discusses different image-reject mixers and provides a detailed description of the self-calibrating image-rejection mixer system used in the Dual-Conversion Wide-Band IF receiver. Chapter 4 details possible Digital Signal Processing (DSP) algorithms for self-calibration and discusses the derivation of and simulation results for their performance in the context of noise. Chapter 5 details the custom VLSI implementation of these DSP algorithms in a standard CMOS process. Chapter 6 concludes the thesis and suggests future study related to this project.

2

Receiver Architectures

2.1 General Considerations

The wireless communications environment imposes severe constraints on radio transmitter and receiver design. Only a limited spectrum is allocated to each user, which translates to a limited rate of information and requires signal coding, compression and efficient modulation.

The small available bandwidth also affects RF design since the receiver must be able to process a desired wireless channel while sufficiently rejecting neighboring interferers. For example, in DCS-1800, a band-pass filter must have a pass-band 200kHz wide and provide 60 dB of attenuation 500 kHz from the center of a desired channel. The filter's Quality Factor (Q) is on the order of 10^5 , which cannot be achieved without using a Surface Acoustic Wave (SAW) filter, or similar technology.

However, insertion loss increases with Q, which increases the noise figure of successive stages by the loss factor. If a band-pass filter has an insertion loss of 3 dB and the Low-Noise Amplifier that follows has a noise figure of 2 dB, then the

combined noise figure of the two cascaded devices rises to 5 dB, which reduces the receiver's overall sensitivity. Receiver design incorporates numerous trade-offs to simultaneously maximize receiver sensitivity, selectivity and the level of integration.

2.2 Receiver Architectures

The superheterodyne is the most common RF receiver architecture. In order to provide good channel selectivity and receiver sensitivity, the superheterodyne requires high-Q filters, which makes it difficult to integrate an entire receiver and to support multiple standards. The direct-conversion receiver eliminates the need for several high-Q filters, but it suffers from significant DC offset problems. Image-rejection receivers offer greater flexibility in supporting multiple standards and lend themselves to increased integration, but their performance is limited by inherent circuit mismatches. The performance of these three receiver architectures will be discussed with respect to the level of integration, ability to implement multiple standards and ease of circuit implementation.

2.2.1 Super-Heterodyne Receiver.

In the superheterodyne architecture (Figure 1), the signal band is translated to a lower frequency, which relaxes the Q requirements for the channel selection filter.

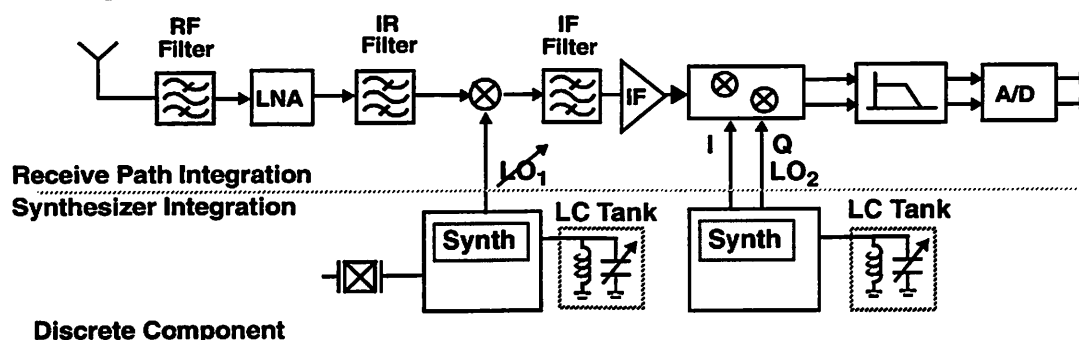


Figure 1. Conventional Super-Heterodyne Receiver. [Courtesy J. Rudell]

The translation of the carrier signal to a lower frequency is accomplished using a mixer, which can be viewed as an analog multiplier. If a signal centered at ω_{RF} is to be translated to a frequency ω_{IF} , it is mixed with a sinusoid $\omega_{LO} = \omega_{RF} - \omega_{IF}$ generated by a local oscillator, which yields a signal at ω_{IF} , and a signal at $2\omega_{RF} - \omega_{IF}$, which is removed through low-pass filtering. The lower frequency is known as the intermediate frequency (IF).

2.2.1.1 The Image Frequency Problem

The frequency translation process creates unwanted signal artifacts, which can interfere with the desired signal. It is helpful to consider a simple time-domain representation of the mixing process. Figure 2 shows a down-conversion mixer:




Figure 2. Simplified model of a single mixer. [Courtesy J. Rudell]

The down-conversion mixing process generates two output signals, one at a desired IF, and a second at a higher undesired frequency. The mixer also translates an undesired signal to the IF. If the local oscillator is at frequency $\omega_{LO} = \omega_{RF} - \omega_{IF}$, then both the signal band at ω_{RF} and the signal band at $\omega_{RF} - 2\omega_{IF}$ are translated to the IF. The signal at $\omega_{RF} - 2\omega_{IF}$ is known as the image frequency.

The image frequency problem is significant because while wireless standards restrict the output power of their own users, they have no control over adjacent bands. The power of the image signal can be much higher than that of the desired signal,

which forces the wireless designer to include some form of image rejection in a receiver architecture.

The most common method of image suppression in a superheterodyne architecture is to place an image-reject filter immediately before the mixing stage in the receive path (Figure 1). This filter must have low loss in the signal band, and large attenuation in the image band. Provided the receiver's IF is large enough, the filter's Q will be small enough for a practical filter implementation; however, the receiver's IF must be small in order for an IF channel selection filter to be practical. This results in a significant trade-off in the design of superheterodyne receivers since a low IF allows great suppression of adjacent channel interferers, while a high IF gives substantial image rejection[5].

2.2.2 Direct Conversion Receiver

The “homodyne” or “direct-conversion” receiver architecture eliminates many discrete components in the receive signal path. In this approach, the RF signal is translated directly to baseband using a single mixer stage (Figure 3). This removes the IF stage from the receiver, and eliminates the need for image rejection. Energy from undesired channels is removed with on-chip low-pass filtering.




Figure 3. Direct Conversion Architecture. [Courtesy J. Rudell]

There are problems associated with the direct conversion architecture despite its greater degree of integration than a superheterodyne system. Since the Local Oscillator (LO) is at the same frequency as the RF carrier, the receiver may have LO leakage directly to the mixer input. The LO signal may also leak to the antenna, and be transmitted and “re-received”, which results in self-mixing and gives a time-varying DC offset at the mixer output[6]. Inherent circuit offsets and this time-varying DC offset reduce the receiver’s dynamic range.

2.2.3 Image-Reject Receivers

The use of integrated image-reject filters limits the flexibility of the heterodyne receiver architecture. In contrast, image-reject receivers suppress image frequencies by processing the signal and its image differently. Using several phase shifts, it is possible to generate a replica of the input signal that contains the desired signal and the negation of the image signal. The image signal may then be cancelled by a simple summation operation at baseband.

2.2.3.1 Image-Rejection Architectures

Image-reject receivers can typically be divided into single-conversion and double-conversion architectures. The single-conversion architecture is based on the Hartley modulator[7] and uses a single quadrature mixing stage followed by a pair of quadrature phase shifters. The main drawback of this design is that it is difficult to build low-loss quadrature phase shifters at high frequencies.

An alternative approach is based on the double-conversion Weaver architecture [8], which uses a second quadrature mixing stage to eliminate the phase shifters in the Hartley architecture. An image-rejection architecture that lends itself to the Weaver

architecture and facilitates full receiver integration is the wide-band double conversion architecture.

2.2.3.2 Double Conversion Wide-band IF Receiver

The receiver system proposed in [4](Figure 4) translates the entire channel band from RF to IF using a single mixer. A low-pass filter at IF removes up-converted frequency components, and passes the channel band to the second mixing stage. The channels are then translated to baseband using a tunable channel-select frequency synthesizer. A variable-gain baseband filtering network is used to remove unwanted channels.

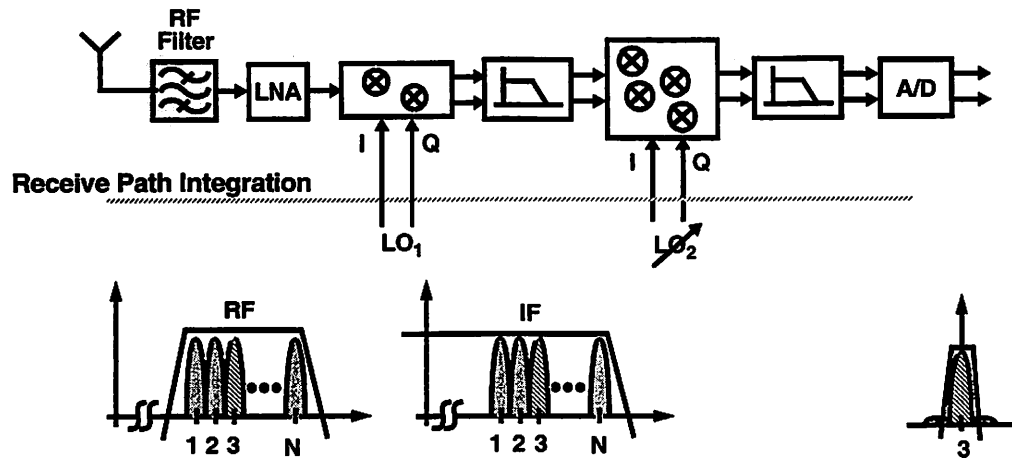


Figure 4. Wide-Band IF Double Conversion Receiver [Courtesy J. Rudell]

This is similar to the superheterodyne approach in that it uses several stages of frequency translation. Unlike a conventional superheterodyne receiver, the first mixer translates the entire receive band, which results in a large signal bandwidth at IF. Channel selection is performed at a second lower-frequency mixer with a tunable local oscillator. As with direct conversion, channel filtering is done at baseband, where digital filters can be used to provide multi-standard receiver features.

The superheterodyne receiver can provide superior performance, but not without the use of discrete component filters, which increases the size of the receiver and tailors it to a particular standard. The small size and increased flexibility of future portable transceivers requires integrated solutions. Both the Direct-Conversion and Wide-Band IF receivers perform channel filtering at baseband, where it is possible to implement an integrated programmable channel filter with multi-standard capabilities.

Compared to the Direct-Conversion approach, the Wide-Band IF architecture offers several potential advantages. The most significant is that channel tuning takes place at the second mixing stage and does not use the first LO, which can then be implemented using a Phase-Locked Loop (PLL) and a fixed-frequency crystal oscillator. This reduces the VCO's contribution to phase noise and allows it to be implemented using low-Q on-chip components. In addition, since channel tuning is performed by the IF oscillator, the PLL's divider ratio can be reduced. This reduces the contributions of the reference oscillator, phase detector and divider circuit to the phase noise of the frequency synthesizer. A lower divider ratio also reduces spurious tones generated by the PLL. The wide-band IF architecture is also insensitive to LO re-transmission and self-mixing since there is no local oscillator at the same frequency as the RF carrier.

The wide-band IF system facilitates high integration, but several non-idealities limit overall receiver performance. The first local oscillator is fixed in frequency, and channel selection is performed at the second mixing stage. This increases the relative frequency tuning range required of the IF synthesizer. The IF channel select filter is also removed, which makes adjacent channel interference more of a concern, and increases the dynamic range requirement of receiver blocks. The harmonics of the second LO may also down-convert unfiltered interferers from IF to baseband. The most

significant non-ideality is the inherent LO phase mismatch and conversion gain mismatch between the I and Q channels of the mixers. The phase and gain mismatches will be discussed in Chapter 3, with particular emphasis on design methodologies that can be used to eliminate circuit mismatches.

3

Image-Rejection Mixers

3.1 Introduction

Some issues associated with integrated receiver architectures were described in Chapter 2. Receivers other than the direct-conversion receiver architecture must provide some means of image suppression. In a conventional superheterodyne receiver which uses discrete components, the image is suppressed by a series of discrete filters. In the interest of both circuit cost and size it is desirable to attenuate the image band using a minimum of off-chip discrete components.

This chapter describes several image-rejection methods. The image-reject mixer architecture that is at the core of the Double-Conversion Wide-Band IF receiver architecture is then described and analyzed. The non-idealities of this mixer will be discussed along with methods to overcome them.

3.2 Image-Rejection Mixers

Image-rejection mixer designs have typically attempted to increase the Image-Rejection Ratio (IRR) by generating low-mismatch quadrature signals[9] or with symmetry between the I and Q signal paths. For example, each set of I and Q mixers can be made mirror images of each other, common-centroid layout techniques can be used, and local oscillator I and Q signal traces can have perfect symmetry. The highest image-rejection ratio reported using these methods is 44dB[10]. Usually, the image-rejection ratio achievable using such a mixer is approximately 30-40dB using good layout practices[11].

45dB of image-suppression is adequate for receivers with low to moderate selectivity requirements for short-range applications, but is far less than that required in a long-range heterodyne receiver with more demanding selectivity requirements. In past, higher image-rejection ratios, typically on the order of 35-50dB, have been achieved by external tuning or laser trimming[4]. Narrowband image-rejection of 80dB has recently been reported for hand-tuning of the local oscillator phase offset[12]; this method offers 60 dB of image-rejection over a 400 MHz bandwidth. A recent analog signal processing approach achieves 58 dB of narrowband image-rejection [28] consumes 170 mW and converges in 500 μ s.

Clearly, trimming can improve image-suppression, but it adds significant cost to large-scale production. It also does not guarantee adequate image-rejection over time and temperature variations. Therefore, trimming techniques together with improved layout techniques restrict this type of single-sideband mixer to applications that have only moderate image rejection and blocking requirements. An alternative monolithic solution involves the implementation of an on-chip high-Q band-pass noise

filter after the LNA, which may be difficult to do using the available silicon components in a given process.

The need for a practical monolithic solution for an image-rejection mixer that can perform internal gain and phase calibration is obvious. Self-calibration techniques have been used extensively in analog-to-digital converters [13][14][15] to correct for non-idealities in baseband components. An adaptive method has been proposed for a low-IF receiver [16] and an analogous method for the image-reject mixer in the Wide-Band-IF receiver is desirable.

3.3 Image-Rejection in the Wide-Band IF receiver.

This section describes a self-calibrating Weaver mixer configuration for the Double-Conversion Wide-Band IF receiver, which conforms to the DCS1800 standard and has a 400MHz IF. An analysis of the effects of mismatch on image-suppression will be presented, with particular emphasis on the requirements of the phase and gain tuning circuits that will be used in this system.

The Double-Conversion Wide-Band IF receiver performs a two step frequency translation of the carrier to baseband, which creates a need for some form of image-rejection. The dual-conversion approach facilitates implementation of an image-rejection mixer based on the Weaver method[8]. Six mixers are used implement the image-rejection function and modulate the carrier to baseband quadrature channels, as shown in Figure 5.

The idea in image-reject architectures is to use complex phase shifts to allow the signal and its image to be processed differently; this amounts to changing a cosine

into a sine and vice-versa. A frequency-domain representation of the image-rejection mixer's behavior is shown in Figure 6.




Figure 5. Weaver-based Wide-band IF mixer configuration [Courtesy J. Rudell]

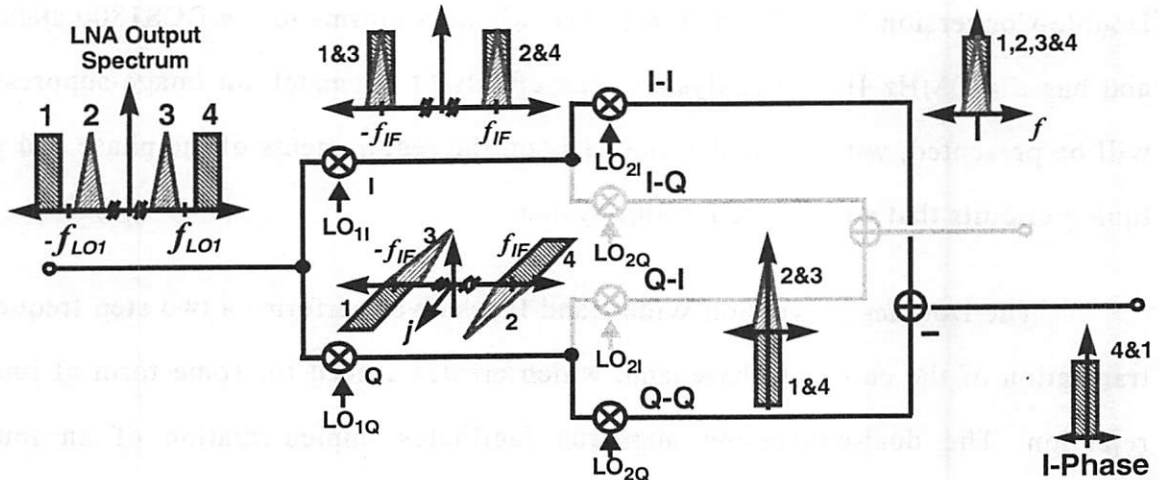


Figure 6. Frequency-Domain representation of Weaver-based Image-Reject Mixer [Courtesy J. Rudell]

The RF carrier is translated to the baseband frequency by two pairs of quadrature mixers, which exploit the relationship between the image and the desired signal. By summing the baseband channels, the image frequencies cancel while the

desired band adds constructively for both the I and Q channels. This is shown in the time-domain in Figure 7.

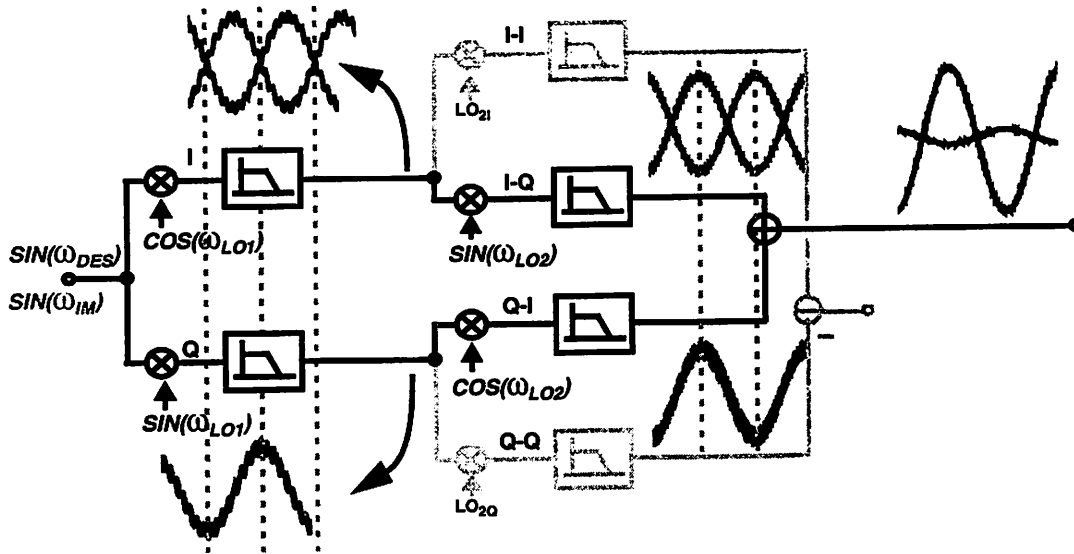


Figure 7. Time-Domain representation of Weaver-based Image-Reject Mixer [Courtesy J. Rudell]

The Weaver architecture has several advantages compared to other image-rejection topologies. The second quadrature mixing stage eliminates the need for phase-shifters to generate the correct phase between the image and desired bands. Low-pass filtering of up-converted terms also gives very wide-band image-rejection.

3.3.1 Non-Idealities of the Weaver Method

For perfect matching, the Weaver architecture provides infinite image-rejection. In practice, however, there is a phase mismatch between quadrature signals and a gain mismatch between mixer signal paths. The magnitude of the image attenuation in the wide-band IF architecture is a function of the mismatch between the I and Q phases of the first and second local oscillators and the gain matching between

signal paths. The image rejection as a function of the mismatch is derived in [4] and given by:

$$IRR = \frac{1 + (1 + \Delta A)^2 + 2(1 + \Delta A)\cos(\phi_1 + \phi_2)}{1 + (1 + \Delta A)^2 - 2(1 + \Delta A)\cos(\phi_1 - \phi_2)} \quad (\text{Eq 3.1})$$

Where ϕ_1 and ϕ_2 represent the deviation of the local oscillators from quadrature in the first and second LOs, respectively, while ΔA is the aggregate gain error between the I and Q signal paths. A plot of Equation 3.1 is given in Figure 8:




Figure 8. Image suppression as a function of the total phase and gain mismatch. [Courtesy J. Rudell]

For a sufficiently high IF, the image-rejection may be performed by both the RF front-end filter and the image-rejection mixer. We desire at least 60dB of image-rejection from the image-reject mixer using this approach. In order to provide an IRR of 60dB with 0.1% gain mismatch, the local oscillator phase errors can be no greater than 0.1 degrees; with a 0.1 degree phase mismatch, the gain error between signal paths must be less than 0.1%.

3.3.2 Self-Calibrating Image-Rejection Mixer

Since the DCS1800 standard has substantial blocking requirements and thus effectively sets the receiver's image-rejection requirements, it will be used to develop

a system specification for the self-calibrating mixer that is described in this section. The fundamental concept of this self-calibration technique is that a test calibration tone at the image frequency is injected into the mixer. By observing the magnitude of the mixer's output tone, we gain information about its inherent phase and gain mismatch.

The magnitude of the output baseband tone can be observed using a digital signal processor (DSP). The calibration tone may be applied to every channel when the mixer is first powered. The mixer's inherent phase and gain mismatch vary with time and temperature, so the mixer may also be calibrated periodically (e.g. - between received frames) to ensure optimal performance.

A detailed block diagram of the adaptive image rejection mixer is shown in Figure 9. Again the six mixer Weaver configuration is shown with its I and Q baseband outputs. The higher frequency oscillator (LO1) generates a fixed 90° phase shift between the mixer's I and Q LO input ports.

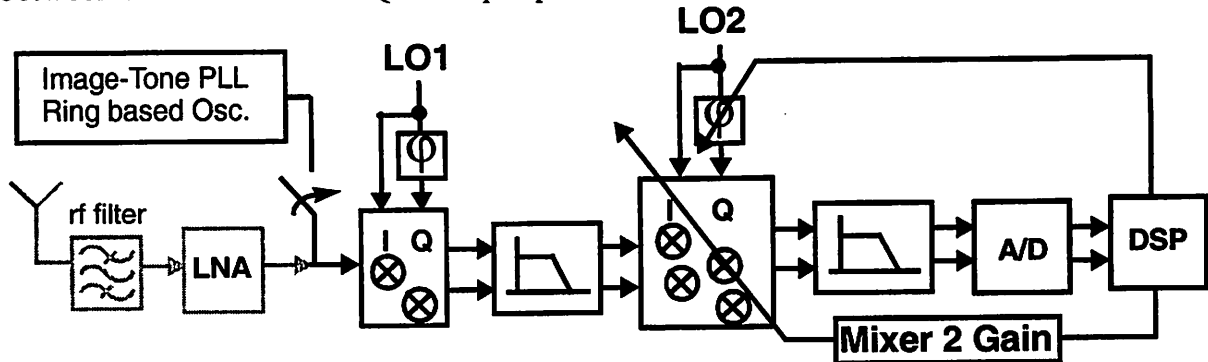


Figure 9. Block level conceptual diagram of the adaptive image rejection mixer. The total phase error is corrected using one tunable generator. [Courtesy J. Rudell]

The second LO section corrects the phase error attributed to ϕ_1 and ϕ_2 . The gain mismatch is then corrected by modulating the gain in two of the four mixers in the second conversion stage. This is represented in Figure 9 by a single tunable gain block. The gain and phase tuning processes must occur independently of each other. When the

minimum image response is located for a given channel, the gain and phase control values are stored in memory and the tuning process continues with successive channels.

3.3.3 Performance Specifications

It is necessary to determine performance specifications for the individual tuning blocks required to achieve a given image-rejection ratio. For a heterodyne system with a 400MHz IF, the DCS1800 standard requires 60dB of image-suppression. From Equation 3.1, it is clear that the image-suppression depends on the resolution of both mixer phase and gain tuning. The required phase tuning resolution depends on the minimum gain tuning resolution and vice-versa.

To determine the required phase tuning resolution it was necessary to make a practical estimate of gain resolution and the maximum gain tuning range required. This was previously determined through simulation and experimental measurements for a prototype DECT IR mixer[4].

For good layout, the gain mismatch between signal paths was found to be +/- 1%; to ensure sufficient gain tuning range, the system was designed to correct as much as a 5% gain error between two channels. In practice, the gain tuning was achieved by modulating the tail current through two Gilbert cell mixers fed by LO2. From simulation, it was determined that a minimum gain tuning resolution of 0.001 between signal paths was possible[3]. This defines both the number of tuning bits required for gain control and the minimum resolution in phase tuning:

Max. ΔA required: 0.05

Min. ΔA required: 0.001

Total resolution of gain control: $0.05/0.001 = 50$ levels or ~ 6 bits of control.

Figure 8 shows the mixer's image-suppression as a function of phase mismatch plotted with contours of the gain mismatch ΔA . The shaded area shows the target performance for a gain mismatch of 0.001, and 60dB of image rejection. A resolution of 0.05° was used to allow higher ratios than 60dB to be achieved. Since tuning must be performed for all channels, a resolution better than 0.1° must be supplied for the entire range of LO2 frequencies.

For this system the total phase mismatch in LO1 (ϕ_1) was specified to be less than 0.5° , while the total error within the signal path (ϕ_2) can be made less than 0.5° . If we specify that the total error in LO2 should be less than 2° , the total range of the phase tuner is 3° . This gives the number of bits required to implement the digital phase tuner:

Maximum range of tuning: 3°

Minimum resolution in phase: 0.05°

Total number of bits required: $3^\circ/0.05^\circ = 150$ levels or 8bits.

With both the phase and gain tuning resolutions defined, a method must be developed to detect when the tuning process has approached the maximum image-rejection. Digital signal processing algorithms for performing this optimization process will be described in Chapter 4.

Image-Rejection Algorithms

4.1 Image-Rejection Mixer

The image-rejection mixer can be modeled as shown:

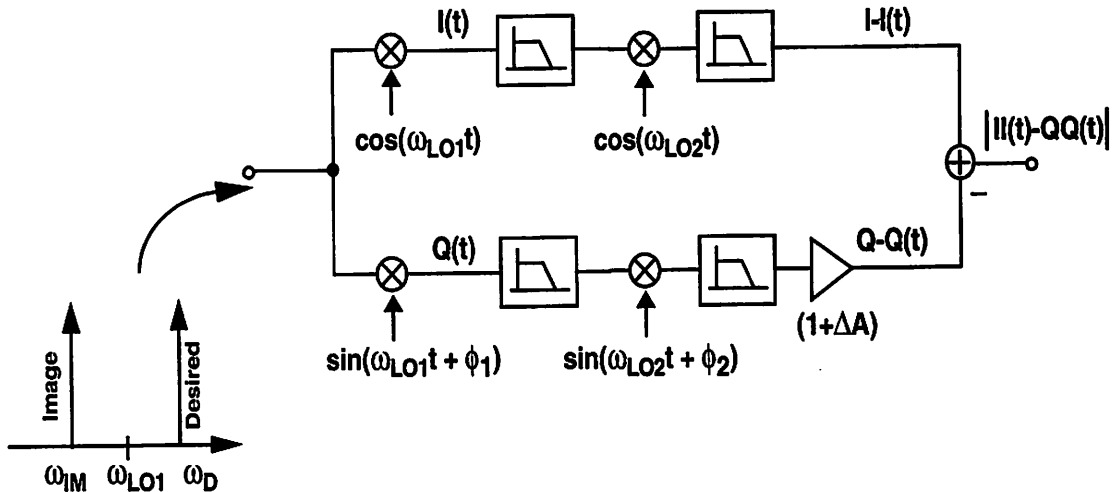


figure 10. Model used to analyze the image-rejection performance as a function of LO phase and gain path matching.[Courtesy J. Rudell]

The entire phase mismatch between LO_1I and LO_1Q is lumped into LO_1Q without loss of generality. The adaptive phase and gain offset values are fed into LO_2Q and the variable gain blocks, respectively.

From [4], for an input of the form $\cos(\omega_{im}t) + j\sin(\omega_{im}t)$, the signals $II(t)$ and $QQ(t)$ are given by:

$$II(t) = \frac{I}{4}(\cos(\omega_{IF}t) - j\sin(\omega_{IF}t)) \quad (\text{Eq 4.1})$$

$$QQ(t) = \frac{I}{4}(\cos(\omega_{IF}t + \epsilon I + \epsilon 2) - j\sin(\omega_{IF}t + \epsilon I - \epsilon 2)) \quad (\text{Eq 4.2})$$

And the time-average of the magnitude squared of the mixer's output signal is:

$$|H(\Delta A, \Delta \theta)|^2 = \frac{I}{16}(I + (I + \Delta A)^2 - 2(I + \Delta A)\cos(\Delta \theta)) \quad (\text{Eq 4.3})$$

This assumes perfect filtering and no LO or RF feed-through in any of the mixers. Clearly the resulting output signal contains only a single frequency. We wish to minimize the power in this signal, so it is imperative that we be able to determine its magnitude. In our sampled system, we will make use of discrete spectral estimation techniques which are described in subsequent sections.

It should be clear that $X(t)$ may be minimized by adjusting the values of ΔA and $\Delta \phi$. Algorithms for generating ΔA and $\Delta \phi$ are also described in subsequent sections, as is the expected performance of these algorithms in the context of noise.

4.2 Spectral Estimation

The basic problem that must be considered is the estimation of the power density spectrum of a signal using only a finite observation interval. For statistically stationary signals, a power spectrum estimate can be made more accurate simply by taking a longer data record [17]. However, for non-stationary signals, we cannot select an arbitrarily long sequence for our estimate, since we are limited by the rate at which the signal statistics vary over time. Ultimately, our goal is to select as short an observation period as possible that still allows us to resolve closely-spaced spectra into their individual spectral components.

4.2.1 Computation of the Energy Spectrum

Consider first the computation of the spectrum of a deterministic signal from a finite sequence. The sequence $x(n)$ results from sampling a continuous-time signal $x_a(t)$ at a sampling rate $1/T_s$. If the Fourier transform of $x_a(t)$ exists, then the quantity $|X_a(f)|^2$, where $X_a(f)$ is the Fourier transform of $x_a(t)$, represents the distribution of signal energy as a function of frequency and is called the energy density spectrum of the signal:

$$S(f) = |X_a(f)|^2 \quad (\text{Eq 4.4})$$

Suppose that we compute the energy spectrum of $x_a(t)$ from samples taken every T_s seconds. The Fourier Transform of the sampled signal is given by:

$$X(f) = \sum_{k=-\infty}^{\infty} x(k) e^{-j2\pi f k T_s} \quad (\text{Eq 4.5})$$

Setting $f = F/(1/T_s)$, and assuming no aliasing, we may express $X(f)$ as:

$$X\left(\frac{F}{F_s}\right) = F_s \sum_{k=-\infty}^{\infty} X_a(F - kF_s) = F_s X_a(F) \quad (\text{Eq 4.6})$$

Thus the voltage spectrum of the discrete signal is identical to the voltage spectrum of the continuous signal, and the energy spectrum of the discrete signal is:

$$S_{xx}\left(\frac{F}{F_s}\right) = \left|X\left(\frac{F}{F_s}\right)\right|^2 = F_s^2 |X_a(F)|^2 \quad (\text{Eq 4.7})$$

In reality, only a finite-duration sequence $x(n)$, $0 \leq n \leq N-1$, is available to estimate the spectrum of the signal. This is equivalent to multiplying $x(n)$ by a rectangular window of length N , or in the frequency domain:

$$X(f) = X(f) \bullet W(f) \quad (\text{Eq 4.8})$$

The convolution of $X(f)$ with the window function smooths its spectrum, provided that the spectrum of $W(f)$ is relatively narrow compared to $X(f)$ - which implies that the

window must be sufficiently long. Regardless of the length of $W(f)$, this convolution results in leakage of spectral energy into bands where there is in fact no energy.

Normally, we would reduce this leakage by using a window with a smooth time-domain cutoff, but as we saw in Chapter 3, our received sequence contains only a single tone (and noise.) Spreading is of limited concern when there is only a single tone. For simplicity, we use a rectangular window, regardless of its leakage properties.

4.2.2 Discrete Fourier Transform

The energy spectral density $S(f)$ can be estimated using the Discrete Fourier Transform (DFT). Given N data points, the computation yields samples of the spectral density at the frequencies $f_k = k/N$:

$$P_{xx}\left(\frac{k}{N}\right) = \frac{1}{N} \left| \sum_{n=0}^{N-1} x(n) e^{-j2\pi n(k/N)} \right|^2 \quad k = 0, 1, \dots, N-1 \quad (\text{Eq 4.9})$$

Note that we are only concerned about the magnitude of a single tone and not the entire received spectrum. Thus we only need to calculate a single sample of $S(f)$ and not the entire DFT. This reduces the computational complexity of spectral estimation substantially.

4.2.3 Estimation of Mixer Output Spectrum

From Section 4.1, the signal output of the mixer is:

$$X(t) = II(t) - QQ(t) = \frac{1}{4}(\cos(\omega_{IF}t) - j\sin(\omega_{IF}t)) - \frac{1}{4}(\cos(\omega_{IF}t + \epsilon_1 + \epsilon_2) - j\sin(\omega_{IF}t + \epsilon_1 - \epsilon_2)) \quad (\text{Eq 4.10})$$

The signal's spectral content can be estimated using an N -point DFT. The value of a given bin of the DFT is:

$$Y(n) = \sum_{k=0}^{N-1} X(k) e^{\frac{-j2\pi nk}{N}} \quad (\text{Eq 4.11})$$

We would like to determine the magnitude of the signal content at 100 kHz. Since the sampling frequency $f_s = 1/T_s = 300$ kHz is exactly three times this frequency, it is convenient to use a 3-point DFT since the spectrum of the 100 kHz signal falls within $Y(1)$ of the DFT. The DFT bin value is calculated by multiplying the signal in the digital domain by 100 kHz quadrature signals. The DFT recovers the 100 kHz spectral content; because of the structure of the calibration architecture, it will later be convenient to work with the magnitude squared of this quantity. For a 3V calibration tone, this is:

$$|H(\Delta A, \Delta \theta)|^2 = \frac{9}{16}(I + (I + \Delta A)^2 - 2(I + \Delta A)\cos(\Delta \theta)) \quad (\text{Eq 4.12})$$

Figure 11 and Figure 12 show plots of $|H|^2$ versus phase mismatch for different values of ΔA and $|H|^2$ versus gain mismatch for different values of $\Delta \theta$, respectively.

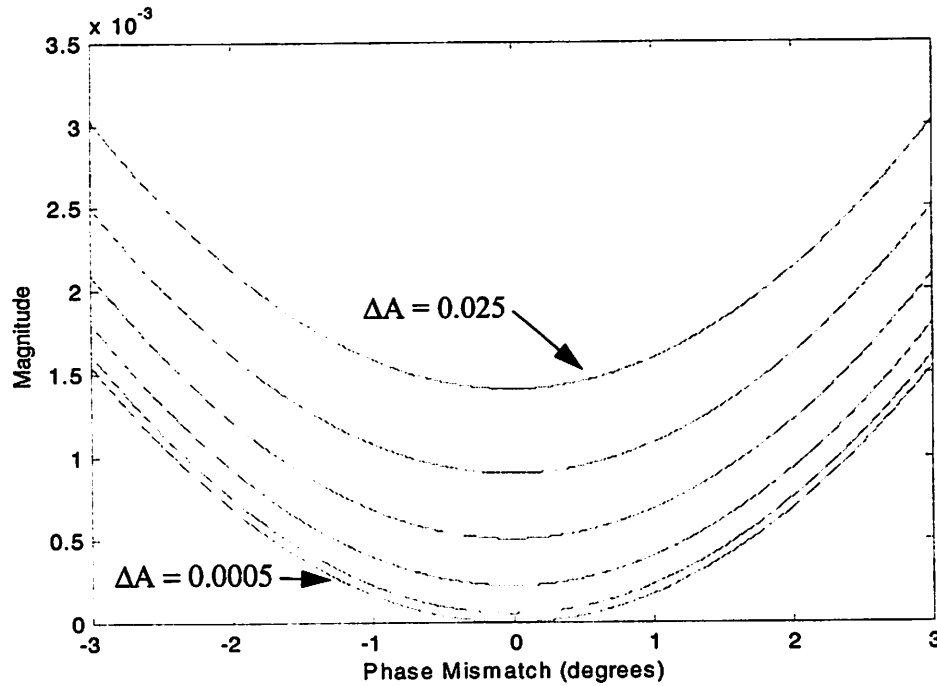


Figure 11. Magnitude of $|H|^2$ vs. Phase Mismatch for varying ΔA (0.0005, 0.005, 0.01, etc...)

Clearly the magnitude of the mixer's output signal reaches a minimum when there is no phase mismatch between quadrature oscillators or gain mismatch along quadrature




Figure 12. Magnitude of $|H|^2$ vs. Gain Mismatch for varying $\Delta\phi$ (0.05, 0.25, 0.45, etc...) channels. The shape of the mixer transfer characteristic is important in developing a minimization algorithm. From Figure 11, it can be seen that the minimum output level is several orders of magnitude lower than the output level at a phase offset of one degree, for example, which allows for easier detection of the minimum. We also note that $|H|^2$ is concave up for the phase offset range of interest, and that it has only one minimum.

The derivatives of $|H|^2$ with respect to gain and phase mismatch are shown in Figure 13 and Figure 14, respectively. Both the derivatives are linear in the range of interest and have only a single zero. These properties allow a fairly simple class of algorithms to be used to locate the zero in the derivative and thus locate the optimal operating point of the mixer.




Figure 13. Two-point derivative of $|H|^2$ with respect to gain mismatch




Figure 14. Two-point derivative of $|H|^2$ with respect to phase mismatch

4.3 Performance in the Presence of Noise

In an ideal system, only a minimum number of signal samples is required to determine a signal's magnitude. However, the performance of a communications receiver is fundamentally limited by its noise figure and the corresponding noise floor power level at each of the receiver blocks. It is thus important to determine the behavior of the minimization process in the context of noise. The most important measure of noise performance in this system is how it affects estimates of the signal at consecutive phase mismatch steps. In other words, if the signal is estimated at two different phase mismatch values, it is important to know how likely it is that noise will corrupt each signal sufficiently that one signal will incorrectly appear larger in magnitude.

4.3.1 Analogy to Binary Digital Communications Receiver

In order to determine the error performance of the calibration algorithm, it is convenient to model the signal processing architecture as a digital communications receiver. Rather than attempting to detect positive and negative voltages that represent information in a bit sequence, the tuning algorithm must estimate the difference between successive signal estimates, which can be viewed as the amplitudes of received bits. The receiver must determine if these "received bits" are negative or positive in magnitude.

The effect of noise is to corrupt the difference between two signal estimates; if a negative difference is made positive by noise or vice-versa, an error occurs. This is exactly analogous to the corruption of a Polar NRZ signal [18] by noise in a binary baseband communications system. The error probability of the system is determined by

the error function and is a function of the signal-to-noise ratio (SNR) of the difference signal. For simulation purposes, the noise in this system was assumed to be Additive White Gaussian Noise (AWGN) and was modeled by a random Gaussian source that was added to the output signal of the mixer at baseband.

4.3.2 Receiver Noise Specifications

Since this algorithm will be used to optimize the Wide-Band IF Double-Conversion receiver [4], the actual noise performance of the receiver determines its usefulness. From the receiver specifications [19], the noise voltage level at the output of the ADC is 153 μV , which corresponds to -76.3 dBV or -46.3 dBm. The tuning algorithms must function in noise levels above this figure.

4.3.3 Estimation of Mixer Output Spectrum

As the gain and phase mismatch of the mixer are adjusted, the mixer output signal varies. In order for any tuning algorithm to work correctly, different signal levels must be reliably distinguishable from each other by spectral estimation. Assume a 1V input sinusoid - then for a change in phase offset, $\Delta\phi$, from $\Delta\phi_0$ to $\Delta\phi_1$, with gain offset, ΔA , held constant, the expected change in the estimated signal energy $|H|^2$ is:

$$\frac{9}{16}(1 + (1 + \Delta A)^2 + 2(1 + \Delta A)\cos(\Delta\theta_1)) - \frac{9}{16}(1 + (1 + \Delta A)^2 + 2(1 + \Delta A)\cos(\Delta\theta_0)) = \frac{9}{8}(1 + \Delta A)(\cos(\Delta\theta_1) - \cos(\Delta\theta_0)) \quad (\text{Eq 4.13})$$

Similarly, from a change in gain offset, ΔA , from ΔA_0 to ΔA_1 , with $\Delta\phi$ held constant, the change in $|H|^2$ is:

$$\frac{9}{16}(1 + (1 + \Delta A_1)^2 + 2(1 + \Delta A_1)\cos(\Delta\theta)) - \frac{9}{16}(1 + (1 + \Delta A_0)^2 + 2(1 + \Delta A_0)\cos(\Delta\theta)) = \frac{9}{16}(\Delta A_1 - \Delta A_0)(\Delta A_1 + \Delta A_0 + 2 - 2\cos(\Delta\theta)) \quad (\text{Eq 4.14})$$

Both the previous and current energy spectrum estimates are corrupted by AWGN with variance σ^2 . Since a 3-point DFT was used for spectral estimation, the noise power is divided evenly across the three DFT bins, so the actual noise power that corrupts the signal at 100 kHz has variance $\sigma^2/3$.

If the magnitude of the noise is greater than the difference between the two energy estimates, then detection will be in error. For a change in phase offset with gain offset held constant, we have:

$$2N_o \geq \frac{9}{8}(I + \Delta A)(\cos(\Delta\theta_1) - \cos(\Delta\theta_0)) \quad (\text{Eq 4.15})$$

The probability of detection error is given by:

$$P(\text{error}) = Q\left(\frac{9(I + \Delta A)(\cos(\Delta\theta_1) - \cos(\Delta\theta_0))}{16(\sigma^2/3)}\right) \quad (\text{Eq 4.16})$$

where Q is the Gaussian Error Function. For an error rate of less than 0.001, for example, the argument of Q must be approximately 3.1 or greater. The noise level that permits this is:

$$\sigma < \sqrt{\frac{27(I + \Delta A)(\cos(\Delta\theta_1) - \cos(\Delta\theta_0))}{49.6}} V/(RtHz) \quad (\text{Eq 4.17})$$

For the smallest possible tuning steps - i.e.: $\Delta A = 0.001$, $\Delta\theta_1 = 0.01$, $\Delta\theta_0 = 0.005$ - we have $\sigma < 79\mu V/RtHz = -49$ dBm. From Section 4.3.2, the expected ADC output noise level is 153 μV ; the phase tuning algorithm does not meet the noise requirements of the system for a simple three-sample DFT. For a 12-sample DFT, however, the noise power is divided by a factor of four, which gives:

$$\sigma < \sqrt{\frac{27(I + \Delta A)(\cos(\Delta\theta_1) - \cos(\Delta\theta_0))}{12.4}} V/(RtHz) \quad (\text{Eq 4.18})$$

This gives $\sigma < 158\mu V/RtHz = -46$ dBm, which meets ADC noise requirements. The algorithm's performance may be further improved using the method described in Section 4.3.4.

We may perform a similar analysis for a change in gain offset with phase offset held constant. From Eq. 4.14, we have:

$$2N_o \geq \frac{9}{16}(\Delta A_I - \Delta A_\theta)(\Delta A_I + \Delta A_\theta + 2 - 2\cos(\Delta\theta)) \quad (\text{Eq 4.19})$$

The probability of detection error is given by:

$$P(\text{error}) = Q\left(\frac{9(\Delta A_I - \Delta A_\theta)(\Delta A_I + \Delta A_\theta + 2 - 2\cos(\Delta\theta))}{32(\sigma^2/3)}\right) \quad (\text{Eq 4.20})$$

where Q is the Gaussian Error Function. Again, we require an argument greater than 3.1 to guarantee an error rate of less than 0.001:

$$\sigma < \sqrt{\frac{27(\Delta A_I - \Delta A_\theta)(\Delta A_I + \Delta A_\theta + 2 - 2\cos(\Delta\theta))}{99.2}} V/(RtHz) \quad (\text{Eq 4.21})$$

For the smallest possible tuning steps - i.e.: $\Delta\theta = 0.005$, $\Delta A_I = 0.0015$, $\Delta A_\theta = 0.0005$ - we have $\sigma < 750\mu V/RtHz = -29$ dBm, which is greater than the expected ADC output noise level, so the gain tuning algorithm meets the noise requirements of the system.

4.3.4 Algorithm Refinements

As in a communications receiver, it is possible to improve noise performance using different coding schemes. We have already calculated the noise performance of the system for three samples of the mixer output. If we use 30 samples in our signal estimation process, then the noise performance improves somewhat since the sampled noise is divided by a factor of 10. The noise performance of these algorithms is shown in Figure 15.

It is clear that substantial performance gains can be achieved by using more than the minimum number of samples for each signal estimate.




Figure 15. Noise performance for different DFT lengths

4.4 Image-Rejection Algorithms

The range of minimization algorithms that can be used for image-rejection is limited both by the amount of time available for spectral estimation and the available computation power. Clearly, it is desirable to minimize completion time and computational complexity simultaneously.

The completion time of the minimization algorithm is a function of the number of signal estimates required, as well as the number of samples required for each signal estimate. It is not difficult to obtain an estimate of the mixer output signal that is sufficient to distinguish a phase mismatch of one degree from a two degree phase mismatch. For very coarse phase and gain tuning, relatively few samples are

required to obtain signal estimates that are not easily corrupted by noise. For finer phase tuning (e.g. - 0.005 degree steps), the difference between signal estimates at successive phase tuning values is much smaller, and estimates are more susceptible to noise; more samples are required to distinguish successive estimates from each other. In order to reduce the number of samples required to estimate the mixer's output signal, a minimal number of signal estimates should be made using the mixer's smallest step-size.

4.4.1 Computation Time

The time required to locate the minimum operating point is significant in evaluating the performance of the image-rejection mixer. It is desirable to perform the optimization process between received frames in a TDMA system, or across all of the RF channels at start-up. Minimizing the convergence time of the self-calibration process is thus important.

4.4.1.1 Linear Search

Consider a linear search with the smallest possible tuning step-size, 0.005 degrees. Since the inherent phase and gain mismatches of the mixer are distributed normally about 0, the linear search begins at 0 degrees phase offset (or 0 gain offset). The phase tuning range is from -3 degrees to 3 degrees. In the absence of noise, the maximum number of tuning steps necessary to reach the optimal operating point is 600 (i.e. - + or - 3 degrees). When noise is added to the system, a linear search may take more tuning steps since the search decision process can be corrupted by errors.

4.4.1.2 Higher-Order Methods

Higher-order numerical methods, such as Newton's Method, perform better than the linear searches since they have a higher order of convergence. The main drawback of higher-order methods is that they require more estimates of a function per iteration, since they typically require that the function's derivative be known. In this system, calculation of a derivative at a point, as is required by Newton's Method, is unreliable; this causes Newton's Method to diverge, which degrades performance compared to an untuned circuit.

Figure 16 shows the mean number of iterations required to locate the minimum operating point of the image-rejection mixer for a phase mismatch with zero-mean and a standard deviation of 0.3 degrees.




Figure 16. Mean number of iterations to locate minimum operating point. (10000 trials)

Newton's Method was halted when successive iterations fell outside of the mixer's phase tuning range, while the linear search was limited to 2500 iterations per trial.




Figure 17. Mean phase mismatch at operating point. (10000 trials)

For lower noise levels, Newton's Method reaches an operating point quite rapidly, usually requiring only two or three function estimates. The linear search algorithm converges far more slowly. Beyond a certain noise level, however, Newton's Method no longer converges, while the linear search algorithm converges in all trials.

The convergence behavior, or lack thereof, of the two search algorithms can be seen in Figure 16. For lower noise power levels, Newton's Method and the Linear Search perform identically, that is, they both locate the minimum operating point to within the smallest possible precision. But the performance of Newton's Method degrades rapidly. A phase mismatch of no more than 0.2 degrees is required to achieve 45 dB of image-rejection for typical gain mismatch values, which makes the variant of

Newton's Method proposed here unsuitable for the noise levels expected at the output of the ADC. The linear search algorithm, despite requiring more time to converge, meets tuning requirements for noise levels up to -30 dBm.

4.4.2 Algorithm Complexity

Higher-order methods are typically much more computationally-intensive than a linear search. In order to simplify the digital signal processing hardware required to control the mixer, it is desirable to limit the number of multiplication and division operations per clock cycle. This reduces the time-slice required on a shared processor, for example; alternatively, a simpler computation method also results in more efficient custom digital hardware.

4.4.2.1 Higher-Order Methods

Consider the form of Newton's Method:

$$f_{n+1}(x) = f_n(x_0) + f'(x_0)(x - x_0) = 0$$

Newton's Method has 3rd-order convergence, but requires an accurate estimate of a function's instantaneous derivative in order to perform the minimization process. For variable phase mismatch, the simplest estimate of a derivative uses two points:

$$\frac{\partial H^2(\theta)}{\partial \theta} = \frac{H^2(\theta + \Delta\theta) - H^2(\theta)}{\Delta\theta} \approx \frac{9}{2\Delta\theta}(1 + \Delta A)(\cos(\theta + \Delta\theta) - \cos(\theta))$$

This method does not provide the most accurate estimate of the function's derivative, compared to multipoint rules like Simpson's Rule, which require even more computation time. This method is also easily corruptible by noise for small $\Delta\theta$. If the derivative is not calculated correctly, then successive iterations of Newton's Method will overshoot or undershoot their intended location, thereby reducing the order of the method and increasing the number of iterations required for convergence. Any method

that requires estimation of the function's derivative at a point, such as the Secant Method, becomes a lower-order method in the context of noise. In addition, successive iterations of Newton's Method require that division operations be performed, which further reduces the efficiency of the method.

4.4.2.2 Linear Search

A linear search is simple computationally, but it has a very slow convergence rate relative to other methods. Its most significant advantage over higher-order methods is that it is not as dependent on the accuracy of derivative estimates in order to converge accurately. Instead, it requires only that the sign of the derivative be known. (In the absence of noise, the linear search is inferior to higher-order methods.) Adaptive techniques can improve the convergence rate of a linear search.

4.4.2.3 Gear-Shifting

In conventional adaptive filtering algorithms, a fixed step-size is used, which is a trade-off between convergence time and steady-state error. A large step-size results in fast convergence but large steady-state error, while the opposite is true for a small step-size. To overcome this problem, time-varying step sizes may be used[21]. This approach, also known as gear-shifting, attempts to sense the distance from the correct operating point and adapt the step-size accordingly.

The problem is then to determine the optimal sequence of step-sizes and some method of sensing the location of the correct operating point. Some methods for optimizing the step-size for an adaptive filter are described in [22]. In this system, we may adjust our linear search step-size to any value from 0.005 to 3 degrees. The factors that govern convergence time and steady-state error are the number of different step sizes used and their values, as well as the values of the thresholds used to decide when

to change the step-size. The optimal values for these parameters cannot be easily determined since they are highly dependent on noise power levels, but near-optimal values can be found by trial-and-error.

It is evident that gear-shifting does improve performance compared to a fixed step-size. Figure 18 shows the mean convergence time of a gear-shifting linear search for a somewhat arbitrary set of parameters (step sizes: 0.64, 0.16, 0.04, 0.005 degrees; switching thresholds: 800, 90, 3, 0.08 mV, with a limit of 14, 14, 14 and 18 steps per step size). Figure 19 shows the mean convergence time for gain tuning (step sizes: 0.036, 0.006, 0.001 degrees; switching thresholds: 10, 0.2, 0.009 mV, with a limit of 14, 14, and 20 steps per step size).




Figure 18. Mean number of iterations to locate minimum phase operating point. (10000 trials)




Figure 19. Mean number of iterations to locate minimum gain operating point. (1000 trials)

Newton's Method at lower noise levels. Because of the difficulty in accurately calculating the function's derivative, Newton's Method begins to diverge around -50 dBm, while the gear-shifting linear search maintains its convergence rate well beyond the level required for this system, increasing slightly as noise levels increase. For a different set of step-size parameters, the convergence time of the linear search could be made approximately constant for the range of noise values shown in Figure 18 and Figure 19.

Gear-shifting also improves the overall noise performance of the mixer. The mean phase mismatch and gain mismatch for different noise power levels are shown in Figure 20 and Figure 21. The gear-shifting linear search converges as rapidly as




Figure 20. Mean phase mismatch. (10000 trials)




Figure 21. Mean gain mismatch. (2000 trials)

Newton's Method, and also converges to a lower gain or phase error at higher noise power levels than the linear search, since it takes fewer iterations and is thus less likely to be corrupted by a large noise voltage. It is possible to further optimize the set of step-size parameters and achieve better noise performance.

4.5 Time-Domain Simulations

The mixer was given an inherent phase mismatch of -0.1925 degrees and a gain mismatch of -5.6% in simulation. The time-domain behavior of the self-calibration process is shown in Figure 22:




Figure 22. Mixer calibration process

The phase tuning process is first performed, followed by the gain tuning. The image-rejection ratio remains approximately constant as the phase is changed, as we would expect for a gain offset greater than 1%. Once the optimal phase offset is reached, any change in the gain offset results in large change in the image-rejection ratio. The Image-Rejection Ratio begins at approximately 38 dB and increases to almost 80 dB by the time the process completes. If the optimal phase offset is reached, the calibration process results in an IRR in excess of 60 dB well before it reaches its final gain tuning value.

The tuning process converges quickly enough that it can be used in conjunction with a GSM receiver. In GSM, the receiver is active during only one out of every eight TDMA time slots. The length of each time slot is 576 μ s [24], so 4.03 ms are available for each stage of the tuning process. From Figure 22, we can see that the optimal phase and gain values are reached in less than 0.4 ms, so the calibration process can easily be interleaved with received GSM frames. The mean number of both phase and gain tuning steps required for convergence is 11, so the average convergence time of the calibration algorithm is 0.22 ms.

4.5.1 Bessel Filter Group Delay

In order to suppress mixing harmonics, a fourth-order low-pass Bessel filter was placed between the output of the mixer and the input of the ADC. During the calibration process, the phase and gain offsets to the mixer are periodically adjusted. This is akin to applying an impulse to the system, which appears at the output of the mixer, and is thus applied to the Bessel Filter. Because of the filter's finite group delay, its output is distorted for some time period following the impulse and can not be used for signal estimation. Thus after it updates the phase or gain offset value, the

calibration algorithm must pause for the length of the filter's group delay before taking another sample from the DSP output.

For a 100 kHz signal, the filter's group delay was calculated to be 15 μ s [23], which means that the calibration process must pause for 5 samples following a change in either the phase or the gain offset. This results in a mean convergence time of 0.59 ms. The phase delay profile of the Bessel Filter is shown in Figure 23. The maximum phase delay occurs for a 20 kHz input signal, but since the DFT suppresses spectral content away from 100 kHz, we are not concerned about the worst-case delay.




Figure 23. Group delay profile for Bessel Filter

5

DSP Architecture

5.1 Custom DSP Implementation

The goal of this work was to develop a custom DSP implementation to interface to the test mixer chip developed in [4] that would consume a negligible amount of battery energy over the battery's lifetime. For example, 0.0001% of a typical cellular handset battery's energy[25] corresponds to approximately 1mW of power consumption at a duty cycle of 1/4000 over a five-hour period. It is also desirable that the DSP circuit implementation occupy less than 0.1 mm², which is approximately the size of the additional digital-to-analog converters used to allow on-chip phase and gain tuning.

5.2 DSP System Architecture

The system consists of three components: an analog mixer, a Discrete Fourier Transform (DFT) block and a Finite State Machine (FSM) block. The DFT and FSM blocks are implemented within the DSP structure. The overall system structure is shown in Figure 24:




Figure 24. Block Diagram of mixer and DSP system

The DFT block estimates the mixer output signal's spectral content, and the DSP algorithms are implemented in the FSM. In order to facilitate the development of a custom DSP ASIC, the system was implemented in simulation using fixed-point arithmetic.

5.2.1 Discrete Fourier Transform Block Architecture

We examine the form of the 3-point ($N=3$) Discrete Fourier Transform:

$$Y(n) = \sum_{k=0}^{N-1} X(k) e^{-j \frac{2\pi n k}{N}} \quad (\text{Eq 5.1})$$

For real $X(k)$, we may implement $Y(n)$ as two parallel Finite Impulse Response (FIR) filters. We wish to obtain the magnitude of $Y(n)$ squared, so we may take the outputs of the real and complex FIR filters and square them and add them to arrive at our required value. One possible architecture to perform this function is shown in Figure 25:

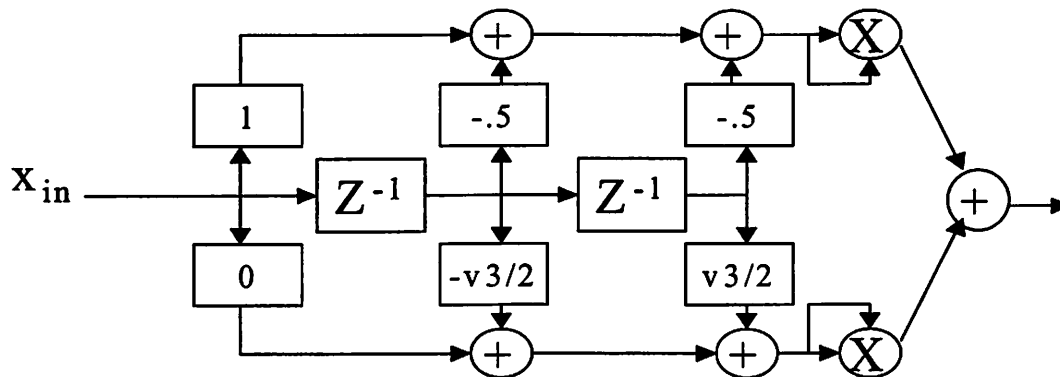


Figure 25. Block Diagram of 3-point DFT Architecture

We may represent all of the rational-valued filter taps using 2 bits of precision, but the representation of the irrational filter taps poses a challenge. In order to represent $\sqrt{3}/2$ accurate to 3 decimal places, we need 13 bits of precision. For a 13-bit signed input from the ADC, the input to the squaring multiplier at the bottom of Figure 25 is 27 bits wide. The output of this multiplier is then 54 bits wide, and the final summation is 55 bits wide.

The DFT architecture may be improved by making a subtle change to the datapath, as shown in Figure 26:




Figure 26. Block Diagram of modified 3-point DFT Architecture

The $\sqrt{3}$ term may be moved past the squaring multiplier, where it appears as $(\sqrt{3})^2 = 3$. Now the set of filter taps $\{-0.5, 0, 0.5, 1\}$ may be represented using only 2 bits of precision. The inputs to the squaring multipliers are only 16 bits, and the summation output is 39 bits wide. For ease of simulation, we truncate the DFT output to 32 bits (so that it may be represented as an integer in Matlab.) This architectural re-arrangement reduced the DFT output width by 23 bits (and reduced the internal wire widths by similar factors) and eliminated the effects of using finite precision filter taps.

5.2.1.1 Arithmetic Refinements

Regardless of the architectural improvements, the majority of the delay through the DFT block is still in the squaring process (16-bit x 16-bit multiplication). The DFT block was

implemented using Synopsys Module Compiler, which implements squaring circuits using a special multiplier that is smaller and faster than a standard multiplier [26], but which does not support Booth encoding.

Indeed, there are few substantial architectural tradeoffs in the DFT block. Operator merging reduces delay slightly in the DFT structure; Carry-save adders offer a tradeoff between area and delay, but their effect is limited compared to the specifications of the squaring multiplier. A Transpose Form FIR filter structure is not suitable to this design because it does not allow sharing of the delay line as in the Direct Form filter. The optimal DFT architecture uses a direct form FIR filter with operator merging; Booth encoding is unimplemented. The relative power consumption, delay and area of different architectures are shown in Table 1:

Table 1: Architectural Trade-offs for DFT Block

FIR Form	Operator Merging	Carry-Save	Area (μm^2)	Delay (ns)	Normalized Power
Transpose	Y	N	123102	24.1	1.93
Direct	N	N	68373	24.9	3.33
Direct	Y	N	65142	24.2	2.84
Direct	Y	Y	65340	24.1	2.84
Modified	Y	N	61218	25.1	1.07
Modified	Y	Y	60363	24.7	1.00

5.2.1.2 Figures of Merit

The target power consumption for the DFT block was 1 mW. The block was implemented in a 1.0V 0.25 μm CMOS process using a custom design flow[27]. A set of test vectors derived from a simulation with arbitrary gain and phase mismatch parameters was used to determine the block's power consumption. The average power consumption over the tuning interval was 11.6 μW ; the peak power consumption was 10.2 mW.

Since the power consumption was found to be much lower than the target figure, and the simulated delay figure of 25ns is much less than that required for operation at 300 kHz, the design was optimized for area. The lowest chip area was obtained for the modified datapath with operator-merging and carry-save adders. The total transistor count for the design was approximately 19500 transistors, and the circuit area was approximately $60000 \mu\text{m}^2$.

5.2.2 Finite State Machine Block Architecture

The variable step-size linear search algorithm was implemented using a finite state machine. The stateflow is shown in Figure 27:

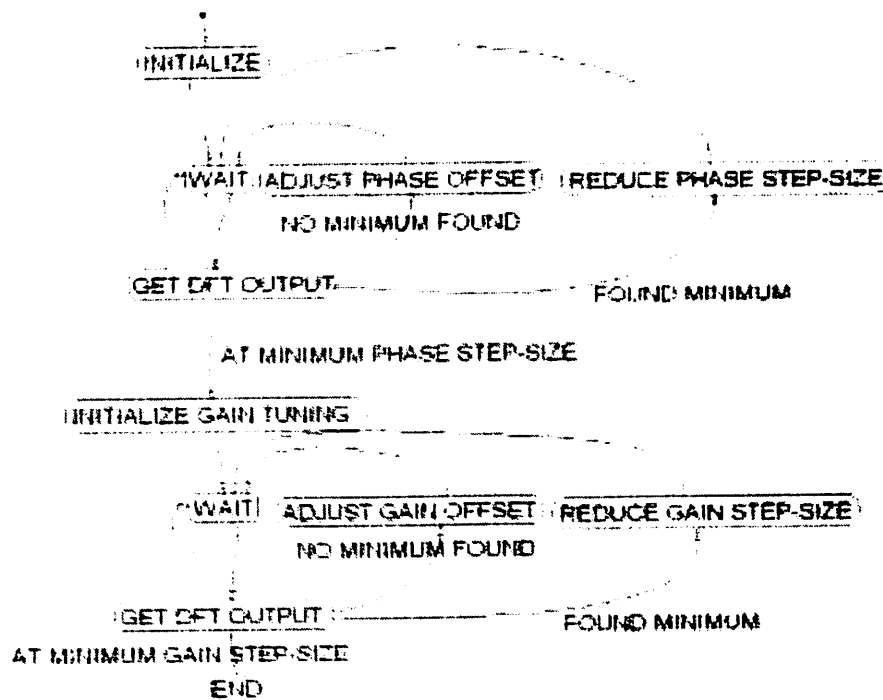


Figure 27. Stateflow diagram of Finite State Machine block

A layout was generated for this state machine in a 1.0V 0.25 μm CMOS process using a custom design flow. The layout consumed an unconnected area of $19914 \mu\text{m}^2$ and had a power consumption of approximately $1 \mu\text{W}$. The phase-tune and gain-tune sections of the state machine are effectively the same at the stateflow level. The state machine's power

consumption is extremely low because of its simplicity, but primarily because its clock rate is only 300 kHz, compared to many MHz in typical processor designs.

5.2.3 Combined DSP Architecture

The FSM and DFT blocks were not integrated in a single layout because of incompatibility between design flows, but we may estimate the performance of the two blocks in combination. The overall DSP power consumption is approximately the sum of the power consumption figures of the two blocks, which is slightly less than 13 μ W. The total area of the design is on the order of 0.08 mm². The mean convergence time is 590 μ s.

The mean total energy used during a calibration process is less than 0.3 nJ. The calibration will likely be performed once per second during the operation of the receiver, giving a duty cycle of less than 0.025%. Over the course of five hours of talk time on a cellular handset, for example, this accounts for less than 0.007% of the energy of a typical handset battery[25], which reduces its life by a negligible amount. A lack of other DSP solutions to the problem of image-reject mixer I and Q calibration makes it difficult to draw comparisons with this DSP architecture, though a recent analog signal processing approach [28] consumes 170 mW and converges in 500 μ s.

6

Conclusion

6.1 Conclusion

This work developed and analyzed digital signal processing algorithms to perform adaptive calibration of an image-rejection mixer. The performance of first-order and higher-order numerical methods was explored and a variable step-size linear search algorithm was developed to perform calibration of both the phase and gain mismatches of the image-rejection mixer.

Simplified models of the image-rejection mixer and the signal processing hardware were developed. Simulation of the complete system showed that the linear search algorithm exhibited sufficiently low convergence times that the process could interleave with TDMA frames during a data transmission.

For noise power levels specified by the receiver's frequency plan, it was shown through simulation that the adaptive algorithms achieved a mean phase mismatch of 0.00129 degrees and a mean gain mismatch of 0.0114%. This corresponds to an image-rejection ratio of 78 dB, which can be achieved across a broad range of frequencies and temperature variations.

The DSP system occupies approximately 0.08 mm^2 of chip area, and consumes slightly less than $13 \text{ }\mu\text{W}$ of power while it is operating. The mean convergence time is $590 \text{ }\mu\text{s}$, and the mean total energy used during a calibration process is less than 0.3 nJ . The calibration will likely be performed once per second during the operation of the receiver, giving a duty cycle of less than 0.025% . Over five hours, this accounts for less than 0.007% of the energy of a typical cellular handset battery[25], which reduces its life by a negligible amount.

6.2 Future Work

Since the DFT is computationally intensive, it is desirable to develop some variant of the Least-Mean Squares (LMS) adaptive algorithm that will simplify arithmetic operations in hardware. It is not known if an LMS algorithm will converge more slowly than a DFT-based algorithm.

The self-calibration technique described in this thesis has been demonstrated only for a TDMA system, and not for spread-spectrum systems. Further study might determine if an image tone could be modulated and fed into a CDMA signal such that the self-calibration process could take place during the course of data transmission. It would be advantageous to develop a general calibration technique that is independent of wireless standards.

References

- [1] P. Gray and R. Meyer, "Future Directions of Silicon ICs for RF Personal Communications," *Custom Integrated Circuits Conf.*, pp. 83-90, May 1995.
- [2] M. Coyne and G. Desjardins, "RF Performance of Dual-Mode Handsets," *Technical Report for Sony of Canada*, Aug. 1998.
- [3] J.C. Rudell, PhD Thesis, 2000
- [4] J.C. Rudell, J.J. Ou, T.B. Cho, G. Chien, F. Brianti, J.A. Weldon, and P.R. Gray, "A 1.9 GHz Wide-Band IF Double Conversion CMOS Receiver for Cordless Telephone Applications," *IEEE J. of Solid-State Circuits*, vol. 32, pp. 2071-2088, Dec. 1997.
- [5] B. Razavi, "RF Microelectronics," Upper Saddle River, NJ: Prentice Hall PTR, 1998, Chapter 5.
- [6] T. H. Lee, *The Design of CMOS Radio Frequency Integrated Circuits*, 1st ed., Cambridge University Press, New York, Chapter 18.
- [7] R. Hartley, "Modulation System," US Patent 1,666,206, April 1928.
- [8] D.K. Weaver, "A Third Method of Generation and Detection of Single-Sideband Signals," *Proc. IRE*, Vol. 44, pp. 1703-1705, Dec. 1956.
- [9] Navid S., Behbahani F., Fotowat A., Hajimiri A., Gaethke R., and Delurio M., Put "Level-Locked Loop, A Technique for Broadband Quadrature Signal Generation", in *Custom Integrated Circuits Conf.*, Santa Clara, CA., May 1997, pp. 411-414.
- [10] J.R. Long, M.C. Maliepaard, "A 1V 900MHz Image-Reject Downconverter in 0.5um CMOS", in *Custom Integrated Circuits Conf.*, San Diego, CA, May 1999, pp. 665-8.
- [11] S. Wu, B. Razavi, "A 900-MHz/1.8-GHz CMOS Receiver for Dual-Band Applications," *IEEE J. of Solid-State Circuits*, vol. 33, pp.2178-2185, Dec. 1998.
- [12] J. Maligeorgos, J.R. Long, "A 2V 5.1-5.8 GHz Image-Reject Receiver with Wide Dynamic Range," in *Int. Solid-State Circuits Conf.*, Feb. 2000, pp.322-3.

- [13] I.E. Opris, L.D. Lewicki, B.C. Wong, "A single-ended 12-bit 20 Msample/s self-calibrating pipeline A/D converter," *IEEE J. of Solid-State Circuits*, vol. 33, pp.1898-903, December 1998.
- [14] U.K. Moon, B.S. Song, "Background digital calibration techniques for pipelined ADCs," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 44, Feb. 1997.
- [15] M.K. Mayes, S.W. Chin, "Monolithic low-power 16 b 1 MSample/s self-calibrating pipeline ADC," in *Int. Solid-State Circuits Conf.*, Feb. 1996, pp. 312-3.
- [16] J.P.F. Glas, "Digital I/Q Imbalance Compensation in a Low-IF Receiver," in *IEEE Globecom*, Sydney, Australia, Nov. 1998, pp. 1461-6.
- [17] J.G. Proakis, D.G. Manolakis, "Digital signal processing: principles, algorithms and applications," Upper Saddle River, NJ: Prentice Hall, 1996, Chapter 12.
- [18] S.S. Haykin, "An Introduction to Analog and Digital Communications," John Wiley and Sons, 1989, pp. 197-8
- [19] J. Rudell, J. Weldon, J. Ou, L. Lin, P. Gray, "An Integrated GSM/DECT Receiver: Design Specifications," University of California Berkeley ERL Memo, 1997. (ERL Memo # UCB/ERL M97/82)
- [20] S.B. Wicker, "Error Control Systems for Digital Communication and Storage," Upper Saddle River, NJ: Prentice Hall, 1995, Chapter 4.
- [21] R.H. Kwong, E.W. Johnston, "A variable step size LMS algorithm," *IEEE Transactions on Signal Processing*, vol. 40, pp. 1633-42, July 1992.
- [22] G. Keratiotis, L. Lind, "Optimum variable step-size sequence for LMS adaptive filters," *IEE Proceedings-Vision, Image and Signal Processing*, vol. 146, pp. 1-6, Feb. 1999.
- [23] J.C. Rudell, Private Communications
- [24] T.S. Rappaport, "Wireless Communications: Principles and Practice," Prentice Hall, 1996.
- [25] Nokia Extended NiMH (900 mAh) Battery BMS-2S, Nokia Battery Specifications
- [26] Module Compiler User Guide, version 2000.05, Synopsys
- [27] W.R. Davis, A Design Flow for Single-Chip Radios, BWRC Retreat Presentation, Summer 1999
- [28] R. Montemayor, B. Razavi, A Self-Calibrating 900-MHz CMOS Image-Reject Receiver, ESSCIRC 2000, pp. 292-295, Sept. 2000.