Ultra-Low Power Wake-Up Receivers for Wireless Sensor Networks

Nathan Pletcher
Jan M. Rabaey

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2008-59
http://www.eecs.berkeley.edu/Pubs/TechRpts/2008/EECS-2008-59.html

May 20, 2008

Acknowledgement

The authors wish to acknowledge the contributions of Simone Gambini, as well as the students, faculty, and sponsors of the Berkeley Wireless Research Center. This research was supported by the National Science Foundation Infrastructure Grant No. 0403427, the California Energy Commission, and wafer fabrication donation from STMicroelectronics.

Ultra-Low Power Wake-Up Receivers for Wireless Sensor Networks

by

Nathan Michael Pletcher

B.S. (Case Western Reserve University) 2002
M.S. (University of California, Berkeley) 2004

A dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Engineering - Electrical Engineering and Computer Sciences

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, BERKELEY

Dissertation Committee:

Professor Jan Rabaey, Chair
Professor Ali Niknejad
Professor Paul Wright

Spring 2008

The dissertation of Nathan Pletcher is approved by:

Chair Date

Date

Date

University of California, Berkeley

Spring 2008

Ultra-Low Power Wake-Up Receivers for Wireless Sensor Networks

by

Nathan Michael Pletcher

Abstract

Ultra-Low Power Wake-Up Receivers for Wireless Sensor Networks

by

Nathan Michael Pletcher

Doctor of Philosophy

in

Engineering - Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Jan M. Rabaey, Chair

The realization of truly ubiquitous wireless sensor networks (WSN) demands ultra-low power wireless communication capability. Because the radio transceiver consumes power whenever it is active, it most efficient to leave the receiver off and wake it up asynchronously only when needed. A dedicated wake-up receiver can continuously monitor the channel, listening for a wake-up signal from other nodes and activating the main receiver upon detection. By maximizing the node sleep time without compromising network latency, the use of a wake-up receiver can improve overall network performance. Wake-up receivers are also applicable in asymmetric links such as "active" RFID, where the tag listens in standby mode until queried by a reader.

In order to be practical, the power consumption of the wake-up receiver must be minimized while still preserving adequate sensitivity to detect the wake-up signal. This thesis explores the specific requirements and challenges for the design of a dedicated wake-up receiver, leading to the design of two prototype receivers implemented in 90 nm CMOS technology and incorporating RF-MEMS resonators. The first prototype combines all required blocks in a low power test system, including a simple RF front-end and mixed-signal baseband. The final wake-up receiver design uses a novel "uncertain-IF" architecture to achieve a sensitivity of -72 dBm at 2 GHz while consuming just 52 µW from a 0.5 V supply. The power consumption is nearly an order-of-magnitude below previously published receiver designs for WSN.

Professor Jan Rabaey
Dissertation Committee Chair

Contents

List of Figures

LIST OF FIGURES

List of Tables

Acknowledgment

As with any research project, this work would not have been possible without the support and assistance of many other individuals. First, I would like to thank my advisor, Professor Jan Rabaey, for his support and guidance throughout my time in Berkeley. His vision and ability to think big have been both inspiring and challenging. On a larger scale, the creation of the Berkeley Wireless Research Center by Jan and Professor Bob Brodersen has provided an interdisciplinary environment that fosters all sorts of productive research.

I also thank Professor Ali Niknejad and Professor Paul Wright for reading this thesis and providing helpful feedback. My interactions and discussions with Paul's group over the years have helped to shape and direct the larger context of this research.

The close collaboration with BWRC member companies has been invaluable during this entire research project. Thanks to Rich Ruby, Mike Frank, and Lori Callaghan at Avago Technologies for taking an interest in our research at the BWRC and providing resonators. I am also grateful to ST Microelectronics for CMOS foundry service donation.

The BWRC is fortunate to employ a very talented and dedicated staff. Special thanks is due to Tom Boot, whose tireless efforts to beautify and maintain the center make it such a comfortable and productive place to work. Kevin Zimmerman, Brian Richards, and Brad Krebs keep the computer systems and tools running smoothly and play an integral role in the successful fabrication of chips at the BWRC. Thanks to Sue Mellers, who fights a constant battle to keep the lab neat and clean. As the gatekeeper of Jan's calendar, the completion of my degree would truly have been impossible without the helpful assistance of Brenda Farrell. Thank you all.

The students of BWRC, my co-workers of the last 6 years, have been a constant source of helpful discussion and distractions. Naming them all would take a page by itself, but from the early days I particularly enjoyed working with the PicoRadioRF group of Brian Otis, Yuen-Hui Chee, Richard Lu, and Simone Gambini. Brian Otis got me up to speed early on and was always up for surfing. Special thanks to Simone, the only one left now, for being my co-author on much of this work and for being a great basketball teammate. I have also enjoyed the friendship of the students in Jan's group, especially Michael Mark, Jesse Richmond, and Louis Alarcón, with whom I have shared many cups of coffee.

Finally, I want to thank my parents, Ken and Justine, and my brother Todd for their support and encouragement over the years. Thanks for all your advice and for listening to my updates every week, only occasionally inquiring about when I plan to graduate.

Last but of course not least, I thank my wife Andrea for her constant support, patience, and willingness to move across the country so that I could pursue this degree. After a great 6 years in Berkeley, I look forward to sharing many more adventures with her. Thank you.

Nathan M. Pletcher

Berkeley, California

Chapter 1

Introduction

The vision of wireless sensor networks (WSN) is ubiquitous wireless, with large networks of wirelessly connected nodes enabling a wide variety of compelling applications. As just one example out of many, WSN are being used to monitor energy consumption in residential buildings with fine-grained sensing capability [1]. The use of WSN enables real-time pricing and adaptive energy usage without user intervention.

The PicoRadio project [2] was begun at the University of California, Berkeley by Professor Jan Rabaey to comprehensively address the challenges in implementing WSN on a large scale, from high-level routing to physical layer electronics. The goal of the project is ubiquitous wireless that disappears into the environment with seamless connectivity and without regular maintenance.

1.1 WSN Implementation Requirements

In order to make these networks a reality, the node hardware and implementation should be optimized for three characteristics:

- **Low cost:** The utility of the network depends on high density and ubiquity, which means large numbers of nodes. In order to make large-scale deployments economically feasible, nodes must be very low cost.
- **Small size:** For the same reasons, the size of modules must be small so that the network is unobtrusive.
- **Low power:** For large networks with many nodes, battery replacement is difficult, expensive, or even impossible. Nodes must be able to function for long periods, ideally up to 10 years, without running out of power.

Each of these three factors are somewhat intertwined. For example, electronic components are already so small that overall module size is limited by power supply or energy storage requirements. For this reason, reducing power consumption of the elctronics is an effective way to shrink size as well. Another example is that highly integrated circuits with few external components can simultaneously reduce both size and cost.

One of the most compelling reasons to reduce power consumption is to enable the use of new power supply technologies like energy harvesting [3] and low cost printable batteries [4]. These early-stage developing technologies cannot supply much power, so any means of reducing power requirements will

hasten the adoption of next-generation power supplies.

Clearly, reducing power consumption is a key method to reach the goals of ubiquitous wireless. Among all the node functions such as computation, sensing, and actuation, the wireless communication energy is still a dominant component [5]. Therefore, the high-level goal of this research is to reduce the energy dedicated to communication in wireless sensor nodes. In order to see where we can attack the power consumption problem, it is important to understand the unique network characteristics of WSN. First, packet traffic rates in WSN are generally low with small chunks of data being exchanged. Packets themselves are short; data packets with 200 bits or less is typical, with even fewer for control packets. The amount of data to be transferred and resulting packet traffic is highly dependent on the specific network application, but most sensing and monitoring applications fit this general form with sparse communications and long periods of idle time.

The natural way to take advantage of the low activity rate is heavy duty-cycling in each node. Duty-cycling is a very powerful means to reduce energy usage and increase battery life. By turning on the node's electronics for short periods of time to perform functions and then entering a low power sleep mode, average power consumption can easily be reduced by orders of magnitude. There is only one problem with spending most of the time in sleep mode: how will nodes know when to wake up? There must be some method of duty-cycle control, arranging for two neighbors to be active simultaneously to allow communication. In the WSN literature, this is called *rendezvous* [6].

1.2 Duty-cycle Control in Sensor Networks

There are several ways of solving the problem of duty-cycle control. Most methods can be described as protocol-based. In synchronous networks, a global reference clock is maintained on each node throughout the network. With a global clock, the protocol can assign communication timeslots to each node. The drawback to synchronous networks is that it may be difficult to maintain and distribute the clock in an ad-hoc network where nodes may be joining and leaving the network. In addition, the energy used to distribute and maintain synchronization can be significant.

Another type of protocol-based duty-cycle control, which avoids a global time reference, is pseudo-asynchronous rendezvous. Depending on the protocol, communication may be initiated by either the transmitting node or the receiving one [7]. Figure 1.1 shows an example of a transmitter-initiated protocol. A timer is used to activate the receiver periodically in order to monitor the channel for communication. If no signal is received, the node returns to sleep mode. When the transmitting node wants to initiate communication it repeatedly sends requests, or beacons, until the receiver wakes up and hears the request, at which time data can be exchanged. Although this method avoids the need for time synchronization between the two nodes, significant energy may be expended both by the receiver (monitoring) and the transmitter (beaconing). More importantly, there is an inherent trade-off between average power consumption and network latency. In order to reduce latency,

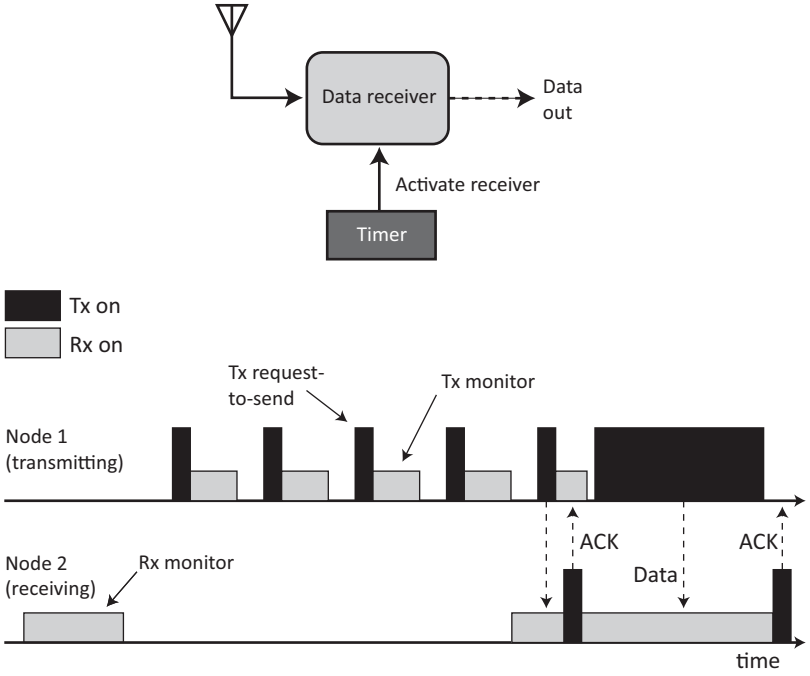

Figure 1.1: Protocol-based duty-cycle control: transmitter initiated

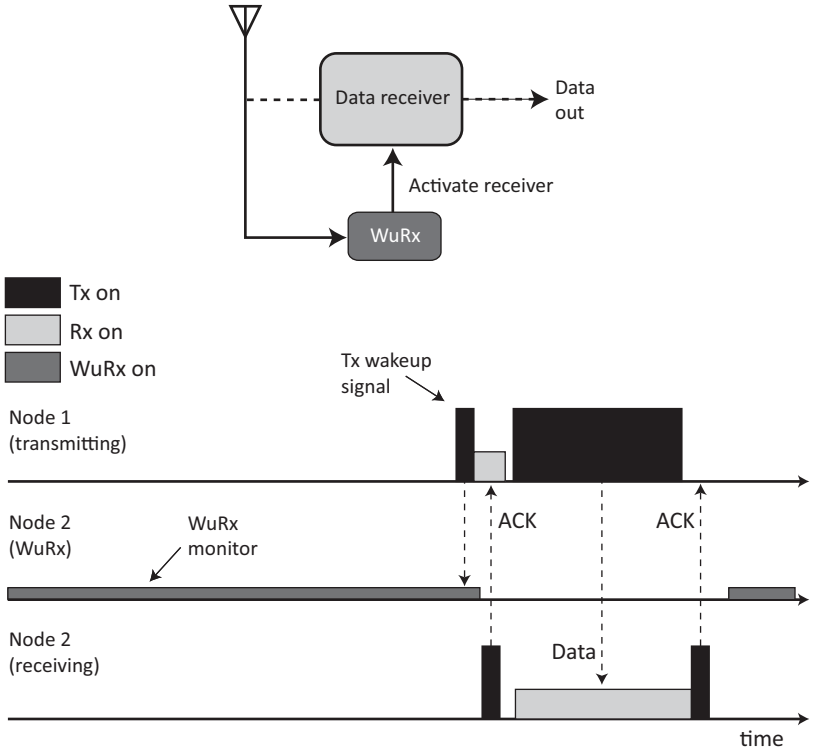

Figure 1.2: Duty-cycle control with wake-up receiver

the protocol must be adjusted for the receiving node to monitor the channel more often, increasing duty-cycle and average power.

An alternative to protocol-based duty-cycle control is based on asynchronous wake-up. This method adds an auxiliary receiver called a wake-up receiver (WuRx) to each node. Its only job is to continuously monitor the channel for communication requests, or wake-up signals. As shown in Figure 1.2, the WuRx now effectively controls the duty-cycle based on actual communication requests, taking the place of the timer used in protocol-based methods. The use of a wake-up receiver breaks the trade-off between latency and average

power consumption described earlier. The WuRx can respond immediately to requests and latency is effectively eliminated. The energy that was previously dedicated to repeated beaconing on the transmit side and periodic monitoring on the receive side is replaced by the power consumption of the WuRx. Because the WuRx is continuously monitoring the channel, its active power consumption must be very low.

Duty-cycle control based on asynchronous wake-up is an attractive alternative to protocol-based methods for many network scenarios, particularly those with low latency requirements. However, very few published wake-up receiver implementations exist in the literature. In [8], the authors extend the battery life of a personal digital assistant (PDA) by activating it only when an incoming request is received. An IEEE 802.11b wireless LAN transceiver is used for data communications in this prototype, while the wake-up receiver is implemented with a commercial off-the-shelf receiver module consuming about 7 mW in receive mode. A much simpler detection circuit with a discrete diode is proposed in [9], but no measurement results are reported to quantify the sensitivity or effectiveness of the radio trigger ciruit. Therefore, this research focuses on the implementation of a practical receiver designed specifically for the wake-up application in WSN. The first step is a high level overview of the design considerations for the WuRx and an outline of the functional specifications.

Table 1.1: Basic characteristics of PicoRadio network

Network architecture	Peer-to-peer
Routing scheme	Multi-hop
Communication range	≈10 meters
Data-rate	10s to 100s kbps
Data receiver power	400 µW
Transmitter output power	-3 to 0 dBm
Carrier frequency/modulation	1.9 GHz / OOK

1.3 Wake-up Receiver Design Considerations

The specifications and implementation of the WuRx depend heavily on the intended application. For networks like wireless LAN, data rates are high and the acceptable power consumption for the WuRx is on the order of several milliwatts, as shown in [8]. For this research, the goal is to implement a wake-up receiver specifically for sensor networks Therefore, the focus will be on the PicoRadio wireless transceiver and its particular specifications and hardware requirements, detailed in [10] and [11]. In summary, the PicoRadio transceiver uses a 1.9 GHz carrier frequency and on-off keyed (OOK) modulation, and the link is designed to operate over a 10 meter distance. The receiver consumes 400 µW when active, while the transmitter output power is about 1 mW (0 dBm) at data rates up to 300 kbps. These operating characteristics are summarized in Table 1.1.

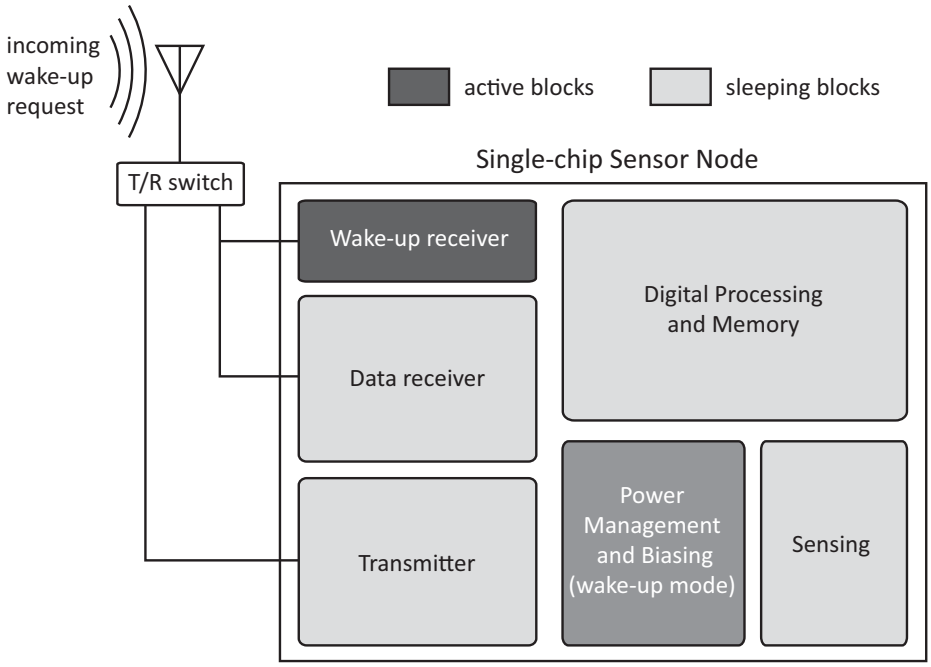

Figure 1.3: Block diagram of sensor node electronics in sleep mode

1.3.1 System Integration

At the system level, the wake-up receiver must integrate conveniently with the rest of the node's electronics. A conceptual diagram of a sensor node in sleep mode is shown in Figure 1.3. During sleep mode, most of the electronics may be powered off, with the exception of the WuRx and any required power management circuitry. From an integration perpective, it is desirable for the WuRx to share the same antenna with the other wireless blocks. To reduce hardware requirements, the WuRx should be able to receive signals from the same transmitter used for data communications, without requiring a separate wake-up transmitter. Therefore, any practical WuRx implementation

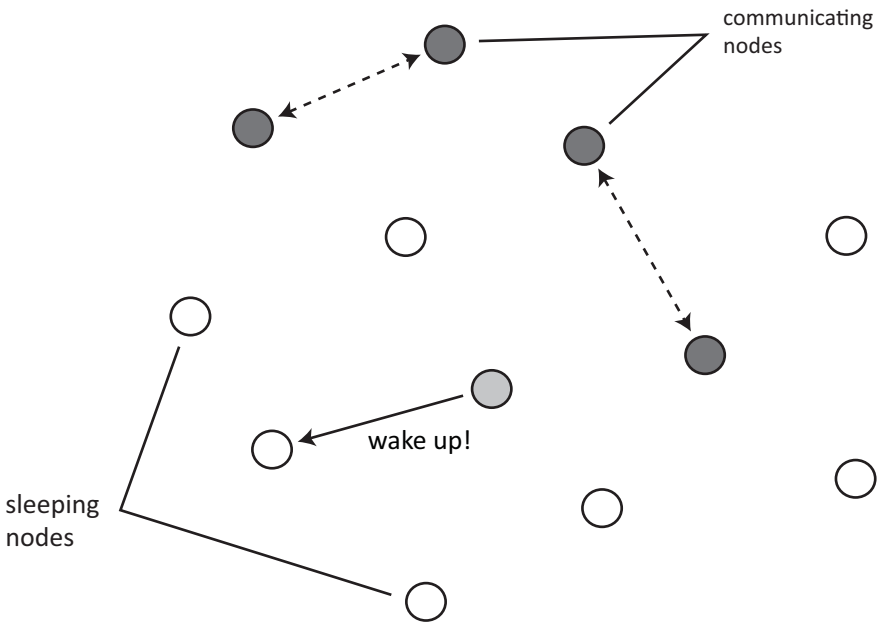

Figure 1.4: Dense wireless sensor network with wake-up rendezvous

will use a similar carrier frequency and modulation scheme as the main data transceiver.

1.3.2 Network Environment

The wake-up receiver is expected to operate in the dense network environment shown in Figure 1.4. At any given moment a few nodes will be communicating, but many will be in deep sleep mode, only monitoring the channel for wake-up requests from other nodes. In this environment, the wake-up receiver must be robust to ambient traffic in the network and avoid waking up on signals intended for neighboring nodes. From a functional perspective, the WuRx

design is not concerned with bit error rate performance as in standard receiver. Instead, the performance metrics of interest are probability of detection and conversely, probability of false alarms (FA). A missed detection means that the transmitter must re-transmit the wake-up request, increasing power and latency. A false alarm is also costly from a power perspective because the main data receiver is activated needlessly.

1.3.3 Optimizing for Active Power

The most important difference between the WuRx and a general purpose receiver is that only active power consumption, as opposed to energy efficiency of communication, is important. For general purpose communication in low-duty cycle applications like WSN, energy per bit is often the metric to be optimized when designing the wireless link [5]. With duty-cycling, high active power consumption can be tolerated as long as the data rate is high enough to result in low overall energy per bit. For general purpose communication, an energy efficient transciever can turn on, exchange a large amount of information quickly, and then go back to sleep.

The WuRx, on the other hand, is always listening for requests and cannot take advantage of duty-cycling. From a design perspective, this observation means that transceiver architectures such as ultra-wideband (UWB) are poor choices for the wake-up application because they rely on synchronization and heavy duty-cycling to achieve low energy per bit. This efficiency comes at the cost of high active power because the receiver must provide wideband gain

with low noise. Therefore, the design goal should be to optimize for active power, not energy efficiency.

1.3.4 Functional Specifications

A wake-up receiver for the PicoRadio network should be able to communicate over the same range as the data transceiver. Otherwise, it may be impossible to wake up a node that could otherwise receive data. As summarized earlier, for the PicoRadio the link range should be 10 meters with a transmitter output power of about 0 dBm. The sensitivity specification is derived using the following simple equation for path loss L_s:

$$L_s = \left(\frac{\lambda}{4\pi}\right)^2 \left(\frac{1}{d}\right)^n \tag{1.1}$$

where λ is the wavelength of the carrier frequency, d is the link distance, and n is the empirical path loss exponent. Assuming transmitter output power of 0 dBm, $\lambda = 15$ cm, and $n = 3$, Equation 1.1 gives a link distance of about 11 meters for a receiver with -70 dBm sensitivity. Therefore, for indoor wireless channel conditions, the receiver sensitivity should be at least -70 dBm.

Functionally, a wake-up receiver is essentially a single bit receiver that detects an event and asserts a signal to activate the data receiver. At the most basic level, the wake-up event could simply be a detection of RF energy. For reliability purposes, however, a practical implementation should ideally be more than just a simple energy detector. Instead, the wake-up signal will most

likely be a particular bit sequence, which allows selective wake-up among multiple nodes and avoids false alarms triggered by regular data communication between neighboring nodes.

The power consumption specification is heavily dependent not only on the main data transceiver power, but also on the network traffic conditions and desired latency. In [7], different types of rendezvous strategies are compared on the basis of average power and network latency. For a traffic rate of 0.1 packets per second with 40 kbps data rate, the WuRx must consume less than about 100 µW to have comparable average power to the synchronous and pseudo-asynchronous rendezvous methods. The analysis assumed data transceiver power of 2.5 mW and 4.5 mW for the receiver and transmitter, respectively. Therefore, the goal for this research is to implement the WuRx with less than 100 µW active power. The overall specifications are summarized in Table 1.2.

Figure 1.5 shows the landscape of previously published receivers for wireless sensor networks [12, 13, 14, 15, 16]. Although several implementations achieve an impressive level of sensitivity with very low power consumption, all of these receivers have power consumption at least 5 to 10 times higher than the budget for the WuRx. Clearly, the feasibility of implementing a functional receiver at gigahertz frequencies with less than 100 µW of power dissipation represents the most significant challenge for the WuRx design.

Table 1.2: Specifications for WuRx prototype

Parameter	Value
Architecture	Narrowband
Carrier frequency	≈2 GHz
Modulation scheme	OOK
Data rate	Unspecified
Sensitivity	-70 dBm
Functionality	Multi-bit sequence recognition
Power consumption	Minimize (≤100 µW)

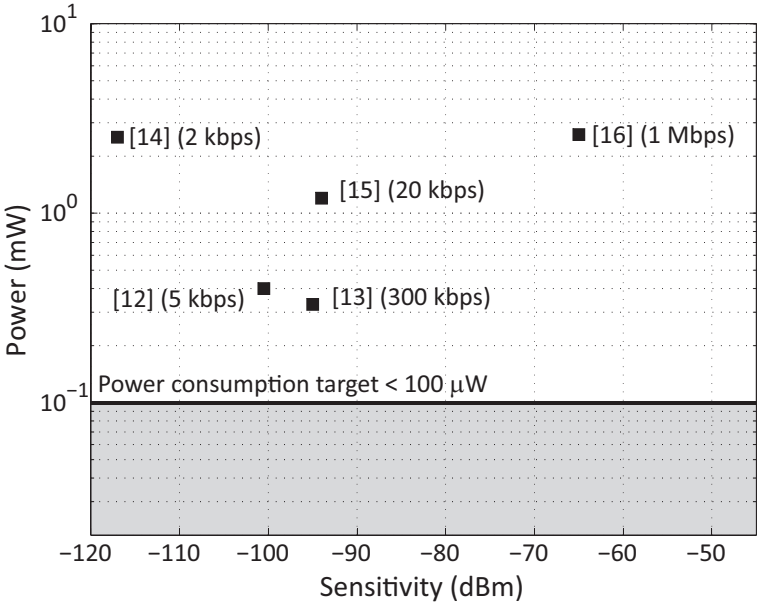

Figure 1.5: Previously published performance of receivers for WSN

1.4 Thesis Organization

This chapter has provided background on methods of controlling duty-cycle in wireless sensor networks and introduced the wake-up receiver concept. Specifications for a prototype wake-up receiver were developed and compared to state-of-the-art implementations of general purpose receivers. The goal of this research is the design and implementation of a dedicated ultra-low power wake-up receiver for the PicoRadio network.

Chapter 2 presents a survey of possible receiver architectures for the WuRx and highlights the factors limiting power consumption for each architecture. At the circuit implementation level, the limits of integrated inductors are discussed and MEMS resonators are presented as a high quality alternative to on-chip passives. Chapter 3 describes the design and implementation of a first prototype, including all the necessary blocks to perform the wake-up function. As a follow up to this first prototype, Chapter 4 details an improved receiver front-end using a novel architecture to boost sensitivity without increasing power dissipation. Finally, Chapter 5 concludes with a brief summary of results and discussion of future research directions.

Chapter 2

Exploring the Design Space

In the last chapter, we developed the specifications on functionality and power consumption for a dedicated wake-up receiver. Given that the power specification is at least a factor of 10 below state-of-the-art low power receivers, it is unlikely that a simple modification or scaling of existing designs will be able to satisfy the requirements. Accordingly, it makes sense to step back from the problem and consider fundamental limitations while exploring the available design space. In this chapter, we outline the architecture choices available for implementation and highlight the factors limiting power consumption in each case. However, because electronic device and fabrication technology is rapidly advancing, we also describe how recent progress in the area of microelectromechanical systems (MEMS) may offer new opportunities to reduce power, bypassing the limitations of integrated inductors in RF circuits.

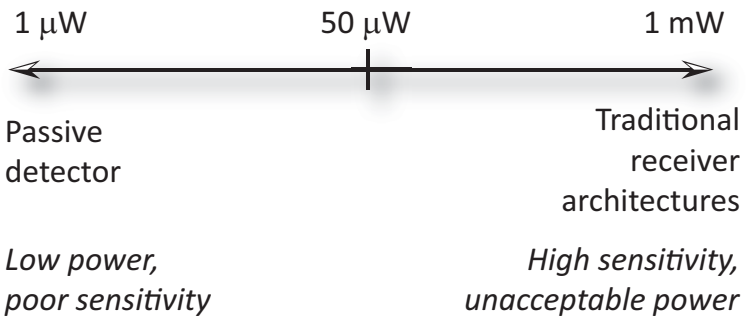

Figure 2.1: Receiver design space in terms of power consumption

2.1 Architecture Considerations

There are a wide variety of ways to build a wireless receiver and detect an RF signal. On one hand are complex receivers that can detect signals with very high sensitivity. On the other hand are simple radio frequency identification (RFID) systems, which do not even have a power supply. We can view the wide variety of receiver architectures on a continuum of power consumption and complexity versus performance, which tend to move together on the scale (Figure 2.1). The target of 50 μW for the WuRx design lies squarely in the middle between domains of low power passive detectors and high performance traditional wireless receivers.

2.1.1 Passive Detectors

Looking first at the low end of the contiuum, an RFID tag is one of the simplest, and therefore lowest power, wireless receivers. Passive tags are able to derive power from the incoming RF waveform and, after storing sufficient en-

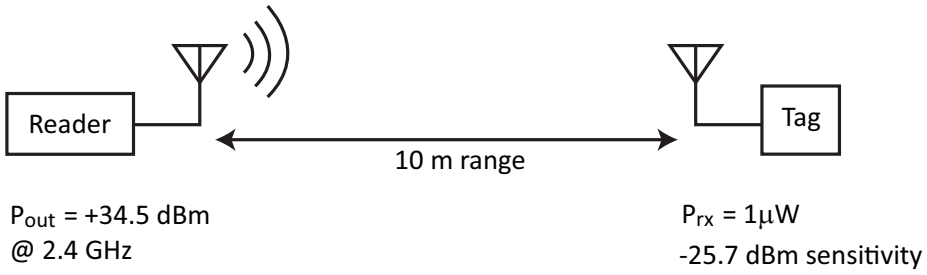

Figure 2.2: RFID link operating parameters

ergy, power up their own electronics to decode an incoming signal and transmit back to the reader. The RFID tag is inactive until it is remotely interrogated by RF energy from the reader. In this way, the operation of the tag is very similar to the desired functionality of the WuRx. One difference is that the reader in an RFID system is typically not power-constrained and is free to transmit with high output power, subject only to regulatory constraints on effective isotropic radiated power (EIRP). In the 2.4 GHz industrial-scientific-medical (ISM) band, for example, the reader may freely transmit up to 4 W EIRP for RFID applications [17].

In sensor network applications, on the other hand, wireless links are peer-to-peer and the power of the transmitter cannot be ignored. In order to quantify the effect of the transmitter in an RFID system, consider the following example of a recently published RFID tag design in the 2.4 GHz band [17]. A simple diagram of the system is shown in Figure 2.2, with the reported operating specifications. The active power consumption of the tag is only about 1 μW and well below the WuRx power budget. However, the RF sensitivity

is poor, reported at -25.7 dBm on a 300 Ω antenna. In an RFID system, the problem can be overcome by simply transmitting higher power from the reader. In order to communicate with the tag over a distance of 10 meters, the reader must transmit with +34.5 dBm output power (P_{out}) at 2.4 GHz. In this example, assume the reader is used to awaken the tag at a regular interval T_{wu} by sending a wake-up signal consisting of particular bit sequence with length N. Under these operating conditions, the average power P_{Tx} consumed by the transmitter during one interval is:

$$P_{Tx} = \frac{P_{out}}{\eta}\frac{N}{R}\frac{1}{T_{wu}} \tag{2.1}$$

where η is the transmitter efficiency[1] and R is the data rate in bits per second. Even if the transmitter efficiency is 100%, Equation 2.1 indicates that the *average* power on the transmit side to send a 15 bit sequence once per second is 425 µW (R=100 kbps). In a peer-to-peer network scenario where the transmitter is power-constrained, this power level is clearly much too high. The poor sensitivity of the tag receiver is the root cause of the high transmit power requirement. Therefore, despite the attractive low power consumption of the RFID tag receiver, a practical WuRx design will require much improved sensitivity in order to avoid shifting the burden of power consumption to the transmitter.

[1]Usually dominated by the power amplifier efficiency

2.1.2 Traditional Architectures

Traditional wireless receivers lie on the other end of the scale from RFID systems in Figure 2.1. These more complex receivers utilize active devices to achieve high sensitivity and data throughput, far beyond what is possible with passive detectors. The high-level architectures used in these receivers can generally be grouped into a few major categories. These architectures are referred to as "traditional" due to the fact that the basic architectures used have not changed substantially in recent years, although the implementation details have become immensely more complex than in the early days of radio. This overview concentrates on narrowband receivers. As mentioned in Chapter 1, UWB architectures are a poor fit for the wake-up application due to high active power consumption and long synchronization times.

The most common type of receiver architecture utilizes frequency conversion, where the input signal is shifted to (usually) lower frequency to ease implementation of signal processing blocks such as gain and filtering. Selectivity is achieved through careful frequency planning, combining narrowband low frequency responses with high purity oscillators and mixers to perform frequency conversion. For example, the super-heterodyne architecture (Figure 2.3(a)) utilizes two separate downconversion operations. First, the input RF signal is amplified by a low noise amplifier (LNA) in order to ease the noise requirements of the rest of the receiver chain. Then, the RF signal is converted to intermediate frequency (IF) with a high-accuracy, tunable local

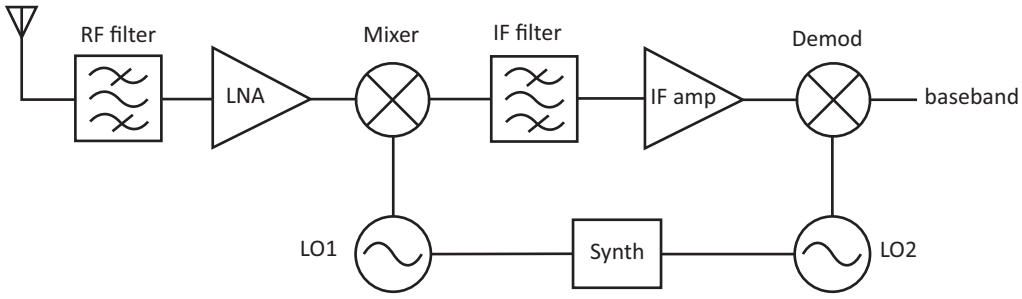

(a) Super-heterodyne architecture

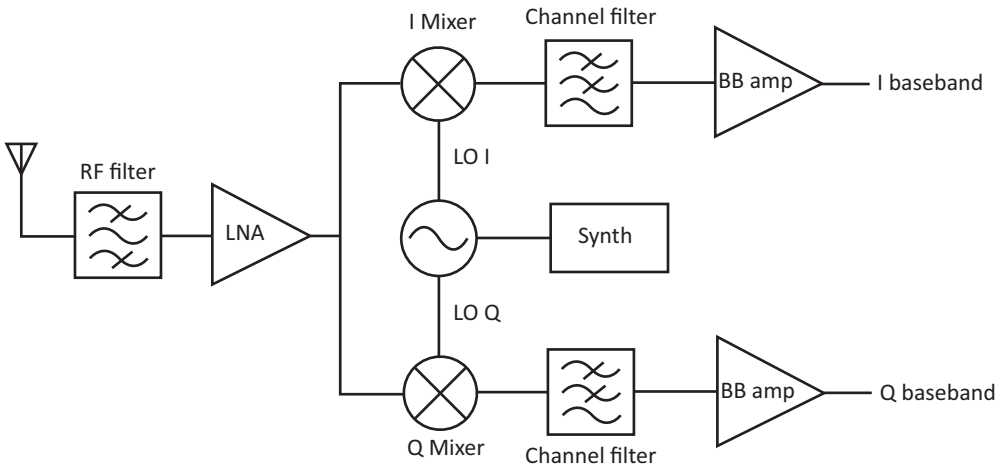

(b) Direct conversion or low-IF architecture

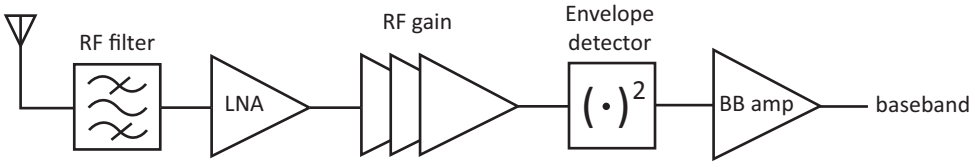

(c) Envelope detection (tuned-RF) architecture

Figure 2.3: Block level comparison of popular receiver architectures

oscillator (LO). This IF signal is amplified and filtered with a fixed frequency filter to remove the image and interferers. A second mixer converts the signal to DC using a fixed frequency oscillator at the IF frequency.

Zero-IF and low-IF receivers (Figure 2.3(b)) avoid the image problem by mixing the RF signal directly to baseband using quadrature downconversion. As in the super-heterodyne architecture, an RF LO with high spectral purity and stability is required to drive the mixer. The power consumption of these architectures, along with super-heterodyne, is fundamentally limited by the RF oscillator and synthesizer. The stringent frequency accuracy and phase noise performance typically requires a resonant LC oscillator, usually embedded in a phase-locked loop (PLL). The limited quality factor (Q) of integrated passives leads to a power floor of a few hundred microwatts.

As an example, consider the recent low-IF receiver implementation described in [13]. In order to save power, the design eliminates the typical LNA and feeds the RF input directly to the quadrature downconversion mixers. The mixers are implemented as passive switching networks using MOSFET switches, so the mixing circuits consume zero DC current. Following the mixers, the receiver circuits process the baseband signal at the low IF frequency (less than 1 MHz), so these amplifiers consume little power. The only remaining element is the oscillator to drive the LO port of the mixers. The oscillator must operate near the RF channel frequency with high accuracy and stability, while simultaneously driving the gates of the mixer switches with a large

amplitude[2] signal. For quadrature operation, the voltage-controlled oscillator (VCO) must also provide both in-phase and quadrature outputs. It is not too surprising, therefore, that the LO generation is responsible for more than 80% of the overall power consumption in the receiver. Despite the use of a large modulation index to eliminate the need for a complete PLL, the VCO itself still consumes more than 300 µW in single-phase, non-quadrature mode. This figure is several times higher than the power budget for the entire WuRx. Clearly, the power devoted to the RF oscillator must be drastically reduced.

As an alternative to frequency conversion architectures, the simplest receiver can be implemented with just RF amplification and an energy detector, similar to the first AM receivers. This architecture, also called "tuned-RF" (TRF), eliminates the power-hungry LO altogether (Figure 2.3(c)). There are two main drawbacks with the TRF architecture. First, since the self-mixing operation is insensitive to phase and frequency, selectivity must be provided through narrowband filtering directly at RF. Second, high RF gain is required to overcome the sensitivity limitations of the energy detector, usually implemented with a nonlinear element like a diode. The TRF receiver is basically an enhanced version of the simple diode rectifiers used in RFID tags, which were shown earlier to have poor sensitivity. The addition of high frequency gain is expensive from a power perspective, so TRF receivers usually exhibit inferior sensitivity compared to mixing architectures for the equal power consumption. In [18], the authors take advantage of the simplicity of the TRF

[2]Ideally square wave

architecture to implement a two-channel receiver at 2 GHz for wireless sensor networks, consuming about 3.5 mW. However, more than 80% of the total receiver power is dedicated to the RF gain stages, divided between the LNA at the antenna and the channel-select amplifiers. The power breakdown illustrates the critical problem with TRF architectures: providing adequate gain at RF usually requires large amounts of power.

One option to enhance gain and improve sensitivity is the use of positive feedback, or regeneration, in the amplifier. This technique was used in the early days of wireless communication [19] to increase the gain available from the vacuum tubes available at the time. A drawback of the technique is that the amount of feedback must be tuned and carefully controlled to enhance the gain without triggering oscillation. The super-regenerative architecture circumvents the need for feedback tuning by allowing the amplifier to oscillate at RF, achieving a large amount of gain from a single stage. The resulting high gain preceding the detector improves sensitivity substantially, to better than -100 dBm [12]. The super-regenerative receiver is fundamentally an envelope detection architecture using a super-regenerative amplifier as an RF gain stage, achieving impressive performance. The drawback is that a high accuracy LO is now required, with performance requirements similar to those of the frequency conversion architectures described above.

In summary, simple RFID receivers are not sensitive enough for peer-to-peer links, while traditional frequency conversion architectures are inherently limited by LO power consumption. In order to significantly reduce the power

of the wake-up receiver, the power contribution of the LO must be reduced.

2.2 Technology Considerations

The power and performance of any receiver will obviously be strongly influenced by the underlying technology used for implementation. For sensor network applications, the only reasonable choice for the active circuitry is the standard digital CMOS integrated circuit (IC). Single-chip integration of digital, analog, and communication circuitry is mandatory to reduce the hardware cost and scaled CMOS is proven to be a good platform for RF circuits as well as digital. For analog and RF design, however, the performance of the active devices is not the whole story. Passive devices also play a key role in determining the ultimate limits of gain and power consumption.

2.2.1 Limitations of Integrated Inductors

Figure 2.4 shows an example of a basic building block of receivers, the generalized gain stage. As shown in the figure, a basic gain stage can be modeled as a simple transconductance stage driving a load impedance. In severely power-constrained designs, the available bias current and device transconductance are limited to small fixed quantities. Therefore, in order to maximize gain, the load should be optimized for high impedance. For RF circuits, the load itself is typically implemented with a resonant LC network, where the impedance

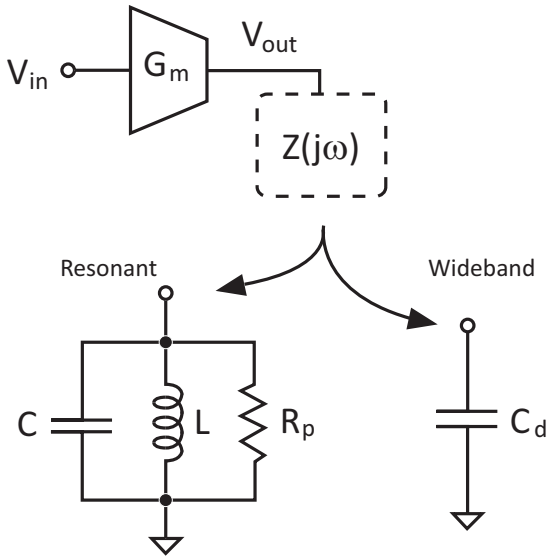

Figure 2.4: Simplified gain stage model

at resonance is given by:

$$R_p = \omega_0 L Q_L \tag{2.2}$$

where ω_0 is the resonant frequency and it is assumed that the network Q is limited by the inductor Q_L. For on-chip inductors in the low GHz regime, R_p is practically limited to a few kilohms by the size and quality of integrated passives. Figure 2.5 shows the calculated R_p at 2 GHz using Equation 2.2 for inductor quality factors of 10 and 20. Large inductors (10 nH) with quality factors of 10 on chip are considered outstanding, with Q of 15 or 20 possible for smaller inductors. Achieving an impedance greater than 1 kΩ is difficult with integrated inductors, which limits amplifier gain. As an example data point, consider a single stage amplifier using 100 µA of bias current. The maximum

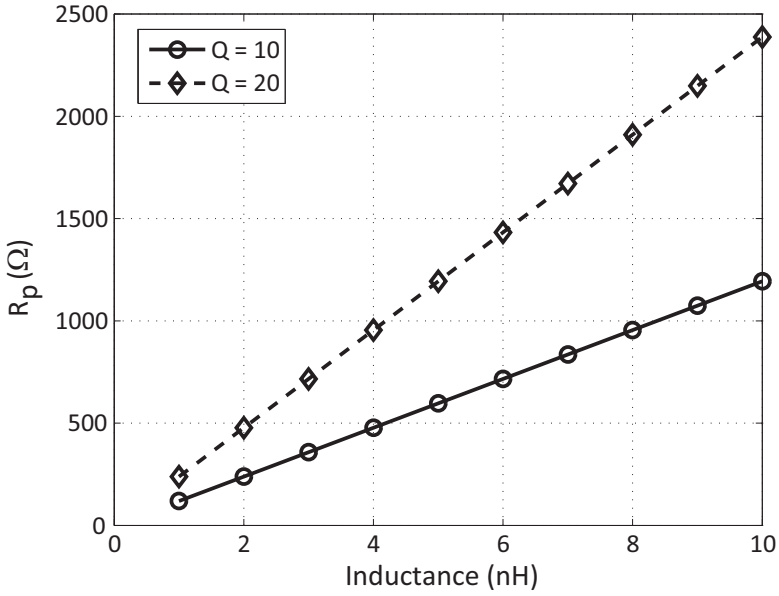

Figure 2.5: Inductor impedance at resonance

transconductance is then about 2 mS, which yields a gain of 2 with 1 kΩ load. Much higher gain will be needed to implement an RF receiver, highlighting the role of passives in low power design.

Unfortunately, technology scaling has little impact on the limitations of passive components because CMOS processes are optimized for digital performance and low cost, so the metallization used for the inductors must use relatively thin layers, increasing the loss and lowering Q. One of the few benefits of scaling is the continuing trend to add more interconnect layers, which helps move inductors further from the substrate and reduce loss. The logical extension of this concept is to post-process additional thick metal layers, specifically optimized for high quality inductors, on top of completed CMOS

wafers. Because the additional layers do not require precision lithography, minimal cost is added to the fabrication process. An example of such an "above-IC" inductor was presented in [20], using a 5 µm thick layer of copper above the CMOS. The combination of thick copper interconnect and larger distance between the coil and substrate results in a measured quality factor of 25 for a 2.5 nH inductor. Although this is almost a factor of two improvement over standard on-chip inductors, the oscillator using this coil still consumed 400 µW of power. This figure is still several times higher than the power consumption target for the entire WuRx, even with the extra processing steps and thick metals. It is therefore unlikely that the R_p available from LC networks will be improved significantly in future IC technologies.

As an alternative to resonant networks, the load impedance can also be implemented as a wideband resistive load (Figure 2.4). In this case the bandwidth is determined by the load capacitance, which is usually the input device capacitance C_d of the subsequent stage. In contrast to resonant networks, scaled CMOS technologies excel at reducing device size and capacitance. The result is that, for fixed frequencies, the impedance magnitude attainable from a wideband network is increasing rapidly with technology scaling, far surpassing resonant networks in 90 nm CMOS. Figure 2.6 illustrates this trend, comparing the impedance magnitude of an LC tank with that of transistor input capacitance at 2 GHz. For the LC tank, a very high quality inductor (L=20 nH, Q=15) is assumed to represent a best-case scenario, and as mentioned above, the impedance stays roughly constant as technology scales. In

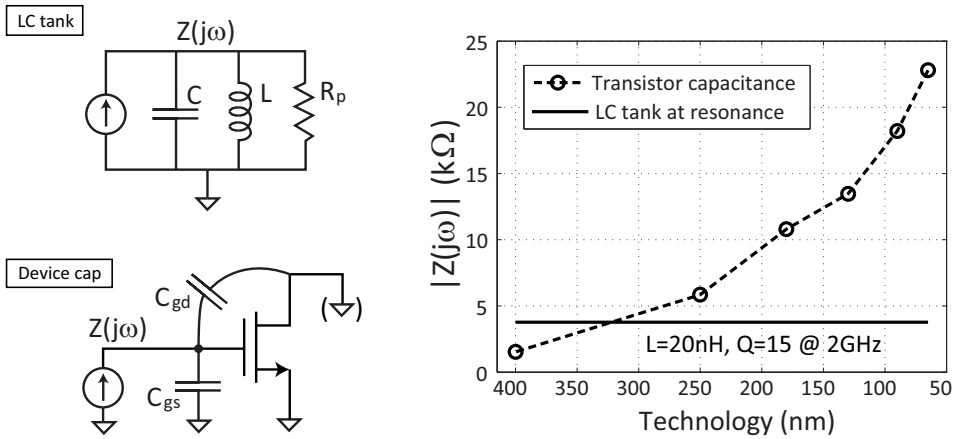

(a) *LC* tank versus device input impedance

(b) Simulated impedance magnitude

Figure 2.6: Effect of CMOS scaling on LC tank and device input impedance

the wideband case, devices in each technology are sized and biased around moderate inversion to provide a transconductance equal to 1 mS, intended to mimic the loading due to a subsequent circuit stage (Figure 2.6(a)). Clearly, the impedance magnitude due to device capacitance in modern technologies has greatly exceeded that of even a very high quality resonant tank. To maximize gain, then, wideband amplifiers and active loads are a promising choice in modern CMOS technology.

2.2.2 Micromechanical Resonators

As an alternative to on-chip passives and traditional off-chip passive components mounted on the printed circuit board (PCB), radio-frequency microelectromechanical systems (RF-MEMS) are emerging as a viable option to break

the trade-off between integration and passive quality. RF-MEMS take advantage of thin-film IC processing techniques to implement high quality resonant structures on the micro scale. Researchers have demonstrated structures with Q factors higher than 10,000 and resonant frequencies up to the low GHz, fabricated using a variety of materials from bulk silicon to diamond and others [21, 22].

Unfortunately, the reliability and stability of these research structures are not ideal for use in circuit prototypes. For that purpose, this research focuses on a type of MEMS resonator that is already in commercial production, the bulk acoustic wave (BAW) resonator. The circuit design techniques developed here to incorporate MEMS resonators will also be applicable to future MEMS devices.

2.2.3 BAW Structure

One common off-chip high quality resonator is the surface acoustic wave (SAW) resonator, where an input piezoelectric transducer uses electric signals to generate a longitudinal acoustic wave traveling on the surface of the piezoelectric substrate. Alternatively, the bulk acoustic wave (BAW) resonator employs a vertical electrode structure to generate acoustic waves that propagate through the bulk of the piezoelectric material. The basic BAW structure is a thin layer of piezoelectric aluminum nitride (AlN) sandwiched between two metal electrodes and fabricated on a silicon substrate. The whole structure must also be acoustically isolated from the substrate to allow free movement.

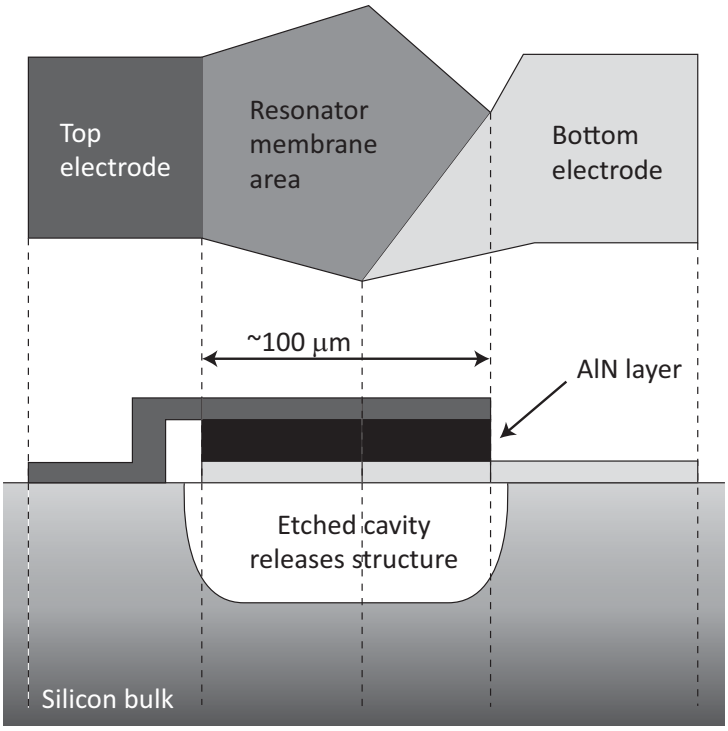

Figure 2.7: Cross-section and top view of FBAR resonator (not to scale)

There are several flavors of BAW resonators, depending mainly on the acoustic isolation method. For this research, we utilize the Film Bulk Acoustic Resonator (FBAR) manufactured by Avago Technologies [23], which uses an etch pit under the resonator. Figure 2.7 shows the structure of an FBAR, where the resonator is fabricated on the a silicon wafer using standard IC processing techniques. The bulk silicon under the resonator is etched away, allowing the structure to vibrate.

In contrast to SAW resonators, whose resonant frequency depends on the lateral spacing of the transducer electrodes, the resonance of a BAW device

depends on the thickness of the AlN layer rather than the surface feature size. This allows the BAW resonator to be made physically much smaller than a SAW device. Quality factors on the order of several hundred to a few thousand are typical, with resonance frequencies in the low GHz range. The standard IC batch fabrication method also results in low manufacturing cost. The combination of small size and low cost of FBAR technology makes it a good fit for wireless microsystems where, as described in Chapter 1, a high level of integration is imperative.

The main drawbacks of BAW resonators are manufacturing tolerance and temperature stability. Typical manufacturing tolerance is about 300 parts per million (ppm), arising partly because tighter tolerances are unnecessary for the most common application in ladder filters. Tolerance can be improved with better manufacturing methods or addressed through trimming. The frequency temperature coefficient for a single resonator is about -25 ppm/℃. Fortunately, the temperature variation is quite linear, which simplifies compensation by external circuitry [10]. The resonator itself may also be compensated by introducing extra layers in the resonator structure. A recently published 600 MHz oscillator using a temperature-compensated resonator achieves frequency variation of less than 80 ppm over a temperature range of -35 to +85℃ [24].

2.2.4 BAW Circuit Models

A simplified circuit model for the BAW resonator is shown in Figure 2.8(a), along with a photo showing a top view of the structure. Table 2.1 gives some

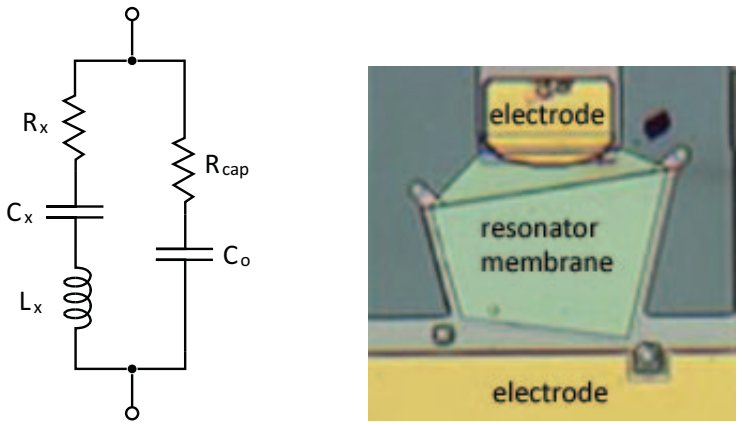

(a) Simplified circuit model

(b) Photo of FBAR

Figure 2.8: Circuit model and photo of BAW resonator

Table 2.1: Typical parameter values for 2 GHz FBAR

Model parameter	Value
L_x	82.5 nH
C_x	78.5 fF
R_x	1.5 Ω
C_o	1.4 pF
R_{cap}	1 Ω

typical values for the model parameters. The example model parameters are for an FBAR resonator at 2 GHz, which is used to build duplexers for the PCS handset band in a 50 Ω environment. However, FBARs are currently produced for various frequencies, particularly 900 MHz and within the range from 1.7 to 2.2 GHz.

As shown in Figure 2.9(a), a BAW structure is characterized by two different resonances. As frequency increases, the series resonance occurs first at a frequency f_s, determined by the motional inductance L_x and capacitance C_x. As expected for a series resonant circuit, the impedance reaches a minimum equal to R_x at f_s. The low motional resistance value of R_x is a major advantage of BAW resonators compared to other types of MEMS devices like bulk silicon resonators. Although the polysilicon resonator published in [21] possesses high Q (greater than 14,000), the motional impedance is 282 kΩ, making it difficult to couple energy into the structure and interface with circuits.

Past the series resonance, the impedance of the structure rises and peaks at the parallel resonance f_p. The resonator appears inductive between the series and parallel resonance frequencies. Outside this range, the response is dominated by the physical parallel plate capacitor C_o. Varying the shunt capacitance in parallel with C_o changes f_p and the impedance at the parallel resonance (R_p), but leaves the series resonance unchanged [25]. Figure 2.9(b) illustrates this effect by shunting the resonator with an additional capacitance C_p in parallel with C_o. The ratio of R_p to R_x falls as loading from C_p increases, although the quality factor of the resonance remains the same if C_p is loss-

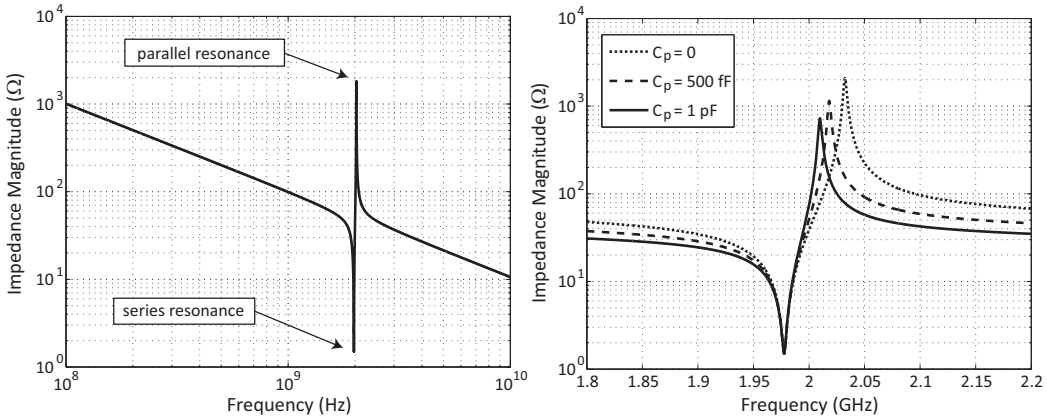

(a) Wideband response (b) Zoomed response for different C_o

Figure 2.9: Simulated BAW resonator impedance response

less. If the resonator is used in the parallel resonant mode where high R_p is desirable, it is critical to minimize the loading from C_p, which can come from external circuitry or wiring parasitics. Although BAW resonators possess high Q factor, they are nevertheless subject to similar limits in R_p [26] and therefore power consumption. The resonator impedance plotted in Figure 2.9 reaches a maximum between 1 and 2 kΩ with a realistic load capacitance. However, the Q factor and frequency stability of these resonators is still much better than what is achievable with integrated passives.

2.2.5 Circuit and BAW Integration

Of course, the most obvious drawback of MEMS components is the reduced level of integration and increased cost. From this perspective, BAW resonators are an attractive choice because they are fabricated on silicon substrates with-

out the use of exotic materials. In fact, several research groups have succeeded in post-processing resonators on top of finished CMOS wafers [27, 28, 29]. The extra processing steps increase cost but result in highly integrated solutions. If the cost of post-processing is too high, the small size of the resonators means that they are good candidates for flip-chip packaging [10, 30]. This well-known "system-in-package" technique can yield very compact implementations with volumes of just a few cubic millimeters. Although MEMS resonators have a reputation as research components that are impractical for real-world products, these recent advances in packaging mean that the use of BAW resonators is well within reach for modules targeting low cost and small size. Accordingly, the use of BAW resonators for low power RF circuits has been popularized recently, using the resonator both in high quality oscillators and as a filtering element [10, 26].

This chapter has summarized the architecture options and limitations for the design of an ultra-low power receiver. At the circuit level, the limited current consumption means that amplifiers utilizing on-chip passives will suffer from low gain. Incorporation of RF-MEMS technology, in particular the BAW resonator, was identified as a possible means to achieve the required receiver functionality without excessively increasing cost or size. With this background in hand, the next chapter describes a first effort to implement the complete WuRx.

Chapter 3

Tuned-RF Receiver

The previous chapter showed that the most difficult specification to meet in the WuRx design is the extremely low power consumption. This chapter develops a first WuRx prototype using a simple receiver architecture that meets the power specification.

3.1 Tuned-RF Receiver Background

The earliest wireless receivers were very simple, consisting of just an antenna to couple energy from the atmosphere and a nonlinear circuit element to demodulate the signal. The most common example is the crystal set, consisting of just an antenna, tuning circuit, and nonlinear envelope detector [19]. As implied by the name, the envelope detection process discards all frequency and phase content of the input signal and simply detects the amplitude of

the RF carrier. Therefore, this type of receiver can only be used to detect amplitude-modulated signals, most commonly on-off keying (OOK). In this case, a "one" is encoded by transmitting the RF carrier, while a "zero" is simply the absence of the carrier. Although OOK is inferior to other modulation methods like frequency or phase modulation from the perspective of link efficiency [31], it offers the advantage of substantial simplification of the circuit implementation and results in large power savings compared to more complex methods.

The use of envelope detection in the TRF receiver makes the operation fundamentally different from more standard architectures like super-heterodyne, entailing different analysis techniques and design trade-offs. Therefore, as the first step in designing a complete TRF receiver, the next section develops a method of analyzing the receiver sensitivity. The results of this analysis can then be applied to the unique design constraints presented by the WuRx application.

3.2 Tuned-RF Sensitivity Analysis

Due to the nonlinear nature of the envelope detector, it is not meaningful to analyze the linear noise figure (NF). In this section we analyze the sensitivity of a hypothetical envelope detection receiver shown in Figure 3.1, consisting of a front-end amplifier with a specified voltage gain (A_v) and noise factor (F_{amp}) followed by a simple envelope detector. The RF filter is assumed to limit the

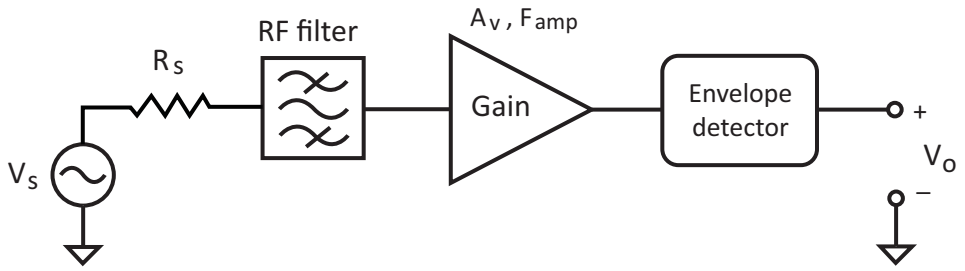

Figure 3.1: Generic envelope detection receiver

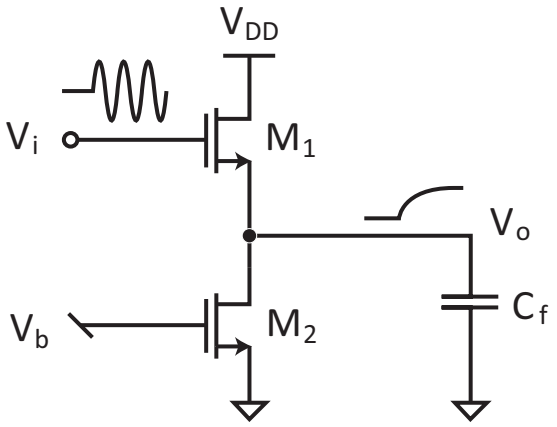

Figure 3.2: Schematic of basic envelope detector circuit in CMOS

noise bandwidth to approximately the same bandwidth as the signal.

3.2.1 Envelope Detector Conversion Gain

The first step is to determine the nonlinear response of the envelope detector. The detection circuit can be implemented using any nonlinear circuit element, such as a diode. However, in a CMOS process it is convenient to realize the detector with the circuit shown in Figure 3.2. This circuit is a CMOS version

of the standard bipolar topology described in [32], and is basically a band-limited source follower. The operation of the circuit in CMOS is similar to the bipolar version if device M_1 is biased in weak inversion, where its drain current is an exponential function of gate-source voltage instead of the weaker nonlinearity of square-law behavior in strong inversion. Device M_2 acts as a simple current source to bias M_1 with a constant current. A large filter capacitor C_f is connected to node V_o. The bandwidth at the output is set by the pole at $f_{p,det}$ formed by C_f and the output impedance of the detector, which is approximately $1/g_{m1}$ neglecting body effect:

$$f_{p,det} = \frac{g_{m1}}{2\pi C_f} \tag{3.1}$$

This pole is designed to be low enough to filter out any signal at the fundamental and higher harmonics, while still affording enough bandwidth to avoid attenuating the baseband signal. For a typical OOK signal, the detected baseband waveform is a square wave with a given baseband data rate, so the detector bandwidth must be high enough to avoid filtering this desired signal.

An AC input signal is applied to the input at V_i in Figure 3.2. Since the output bandwidth is much smaller than the input signal frequency, the full signal appears across the gate-source terminal V_{GS} of M_1. Device M_1 generates an output current that is an exponential function of the input voltage. The nonlinear transfer function contributes a DC term at the output in response to the AC input signal. In order to calculate a simple expression for the ef-

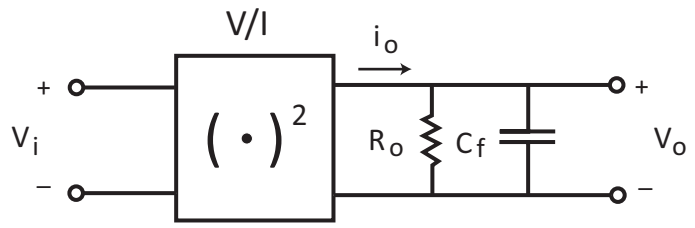

Figure 3.3: Simple model of envelope detector to calculate conversion gain

fective conversion gain from input AC to output DC, the exponential can be approximated by using its Taylor series expansion and dropping terms above the second order. This yields the simple model shown in Figure 3.3, where the detector circuit is modeled as a squaring function that converts an input voltage V_i to an output current i_o. The linear term at the fundamental frequency, along with higher order harmonics, will be filtered out by C_f. Although higher order terms will also generate DC components, these contributions are small compared to the squaring term. The output impedance R_o is simply $1/g_{m1}$.

Using the model in Figure 3.3, the conversion gain k from the AC input voltage to the DC output response can be calculated. First, the large signal drain current of M_1 in weak inversion is modeled as [33]:

$$\begin{aligned} I_D &= I_0' \exp\left(\frac{V_{GS}-V_{th}}{nV_t}\right)\left(1-\exp\left(\frac{-V_{DS}}{V_t}\right)\right) \\ &\approx I_0' \exp\left(\frac{V_{GS}-V_{th}}{nV_t}\right) \end{aligned} \tag{3.2}$$

where I_0' is a constant depending on process and device size, V_{th} is the threshold voltage, V_t is the thermal voltage (kT/q), and n is the subthreshold slope

factor. In the 90 nm CMOS process used for this research, n is approximately 1.5, leading to an nV_t product of 40 mV at room temperature. The variables V_{GS} and V_{DS} are the gate-source and drain-source terminal voltages, respectively. The approximation of I_D holds when the transistor is in saturation [1], which is valid for this source follower circuit.

Next, we find DC output signal current i_o in Figure 3.3 due to an input signal $V_i = V_s \sin(\omega_s t)$. Expanding Equation 3.2 in a Taylor series and focusing on the second order term:

$$\begin{aligned} i_o &= \frac{V_i^2}{2}\frac{\partial^2 I_D}{\partial V_i^2} \\ &= \frac{V_i^2}{2}\frac{\partial}{\partial V_i}\left(\frac{I_D}{nV_t}\right) \\ &= \frac{V_i^2}{2}\frac{I_D}{(nV_t)^2} \end{aligned} \tag{3.3}$$

Substituting for V_i and recognizing that $\frac{I_D}{nV_t} = g_m$:

$$\begin{aligned} i_o &= \frac{g_m}{2nV_t}V_s^2 \sin^2(\omega_s t) \\ &= \frac{g_m}{2nV_t}V_s^2\left(\frac{1-\cos 2\omega_s t}{2}\right) \end{aligned} \tag{3.4}$$

The second harmonic term will be filtered by the detector output pole, giving a DC output current:

$$i_o = \frac{g_m}{4nV_t}V_s^2 \tag{3.5}$$

Finally, we arrive at the DC output voltage by multiplying the output signal

[1] V_{DS} greater than about 150 mV

current by the detector output impedance:

$$V_o = i_o R_o = \frac{i_o}{g_m} = \frac{V_s^2}{4nV_t} \tag{3.6}$$

Therefore the voltage conversion gain k from peak AC input amplitude V_s to output DC voltage V_o is given by:

$$k = \frac{V_o}{V_s} = \frac{V_s}{4nV_t} \tag{3.7}$$

The derivation above holds for small input signals where the response is dominated by the second order term and higher order effects are not significant. For the purposes of analyzing the detector sensitivity, the signals of interest are small and the simple form of Equation 3.7 is a convenient way to represent the detector response. Using the full Bessel function representation in [32], a more accurate expression for gain can be derived [34]. Figure 3.4 compares the Bessel function model with the simple gain expression of Equation 3.7, along with full circuit simulation. The simulation results are for the circuit in Figure 3.2 in a 90 nm CMOS process with $\left(\frac{W}{L}\right)_1$=(10/0.2) µm/µm and bias current of 2 µA. Equation 3.7 is calculated without any parameter fitting and using $n = 1.5$. The simple model is within 20% of the simulated gain for input amplitudes up to 30 mV. If needed, even better accuracy can be obtained by using n as a fitting parameter to match simulations.

Interestingly, Equation 3.7 predicts that the gain is independent of the de-

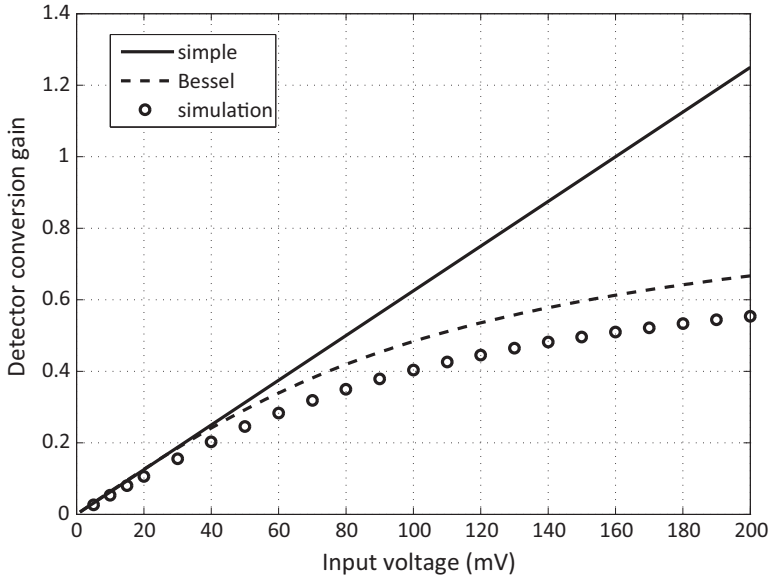

Figure 3.4: Comparison of envelope detector calculations and simulation

vice sizing and transconductance. The derivation above assumes that the device drain current follows an exponential characteristic, so the transistor must be biased in weak inversion. In order to minimize loading on the preceding amplifier, the detector device sizing should be optimized for low input capacitance while still maintaining the device in weak inversion. In deep submicron technologies like 90 nm, minimum channel length should also be avoided due to the high drain-source conductance g_{ds} observed for devices with short channel length. An additional consideration is the output bandwidth, which is determined by the output pole (Equation 3.1) and may affect the bias design. Finally, the transistor may need to be sized larger to lower flicker noise, if it becomes dominant in the overall receiver noise calculation. Noise considerations

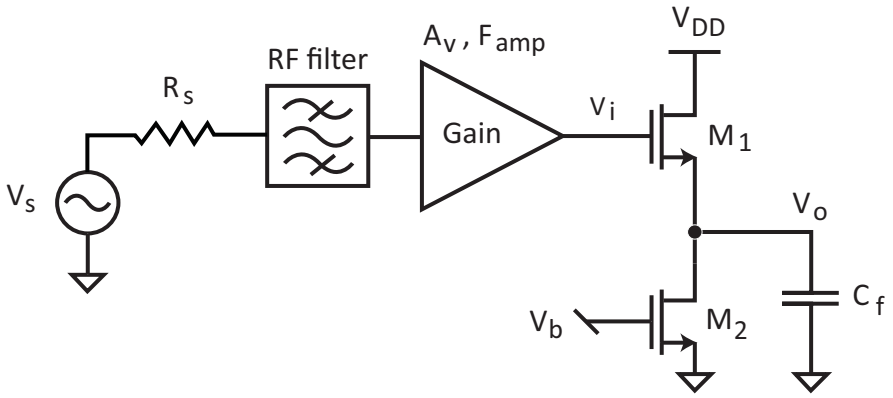

Figure 3.5: Generic TRF receiver with envelope detector

are addressed in Section 3.2.2.

It should be emphasized that this k factor is the *conversion* gain for high frequency signals at the detector input. Any input signals, including noise, at frequencies below the detector output bandwidth will experience the linear transfer function instead, with approximately unity gain ($k_{DC} \approx 1$).

3.2.2 Sensitivity Calculation

Having established a simple expression for the conversion gain of the detector as a function of input voltage, we are ready to re-visit the complete receiver, reproduced here with the detector in Figure 3.5. The ultimate sensitivity can be determined by analyzing the various noise contributions and gain factors to the detector output and calculating an effective *NF* that depends on input signal power. For the simple receiver of Figure 3.5, there are three main noise sources:

1. Noise is added by the amplifier in front of the detector, which is captured by its linear noise factor F_{amp}.

2. The noise of the envelope detector itself, due to M_1 and M_2, appears directly at the output. This noise, $N_{o,ED}$ (V^2/Hz), can be written as:

$$N_{o,ED} = 4kT\gamma\frac{1}{g_{m1}}\left(1+\frac{g_{m2}}{g_{m1}}\right) \tag{3.8}$$

3. Any practical amplifier implementation will exhibit low frequency noise (within the detection bandwidth) at its output. This noise, N_{LF} (V^2/Hz), passes through the detector with gain k_{DC} as described above, and depends on the particular design of the amplifier.

Each noise source is normalized to bandwidth to facilitate the calculation of an overall receiver noise factor, which is defined for a 1 Hz bandwidth. In order to take into account flicker noise and confirm calculations, circuit simulation is used. Figure 3.6 plots the simulated output noise for an example envelope detector design. The relevant design parameters are summarized in Table 3.1. The simulated noise is integrated over the entire band and normalized to the detector bandwidth in order to approximate an equivalent "brickwall" noise density in a 1 Hz bandwidth, as shown in the figure. This equivalent noise density can then be used in the noise factor calculation. The same approach is used with simulations for the low frequency amplifier noise N_{LF}.

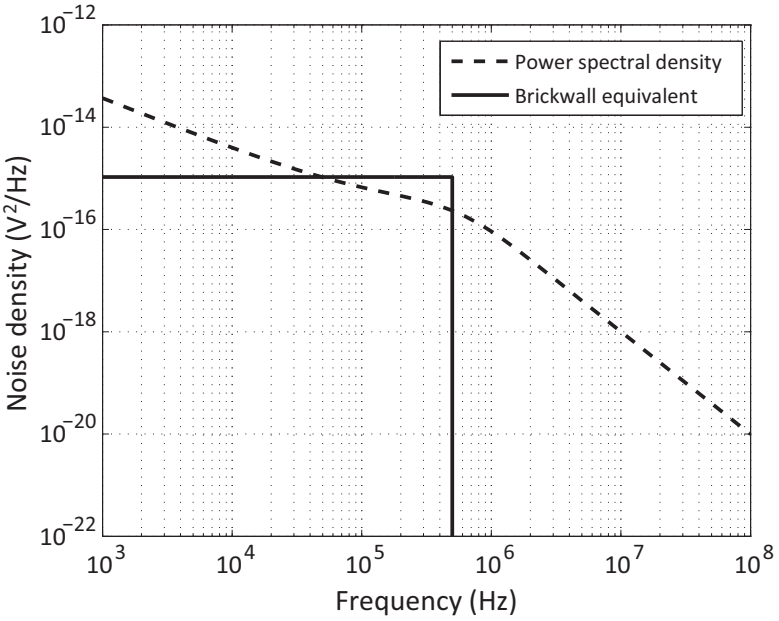

Figure 3.6: Simulated envelope detector noise density

Table 3.1: Example envelope detector design parameters

Parameter	Value
I_{d1}	2.5 µA
g_{m1}	70 µS
g_{m2}	55 µS
C_f	20 pF
$f_{p,det}$	≈ 500 kHz

The total noise factor F_{tot} of the entire receiver can now be written as [34]:

$$F_{tot} = 2F_{amp} + \frac{N_{LF}k_{DC}^2}{N_{src}A_v^2k^2} + \frac{N_{o,ED}}{N_{src}A_v^2k^2} \tag{3.9}$$

where N_{src} is the noise from the source resistance ($4kTR_s$) and A_v is the gain of the front-end amplifier. Because of the dependence of k on signal level, F_{tot} increases with decreasing input power. Using $NF_{tot} = 10\log F_{tot}$ and the detector bandwidth BW_{det}, we can calculate an input-referred noise for the receiver in dBm:

$$P_{n,in} = -174 + 10\log(BW_{det}) + NF_{tot} \tag{3.10}$$

If the minimum signal-to-noise ratio (SNR) for reliable detection is SNR_{min}, the minimum detectable signal (P_{mds}) is the input power for which:

$$P_{mds} = P_{n,in} + SNR_{min} \tag{3.11}$$

where the quantities in Equations 3.10 and 3.11 are expressed in dB. This relationship can be visualized by plotting the noise power $P_{n,in}$ and ($P_{in} - SNR_{min}$) versus P_{in} and finding the intersection. For a typical value of 12 dB for SNR_{min}, the curves are compared for two different front-end amplifiers in Figure 3.7, one with $A_v = 20$ dB and $NF = 10$ dB and the other with 40 dB gain and 20 dB NF. For this example, the low frequency amplifier noise, N_{LF} in Equation 3.9, is ignored.

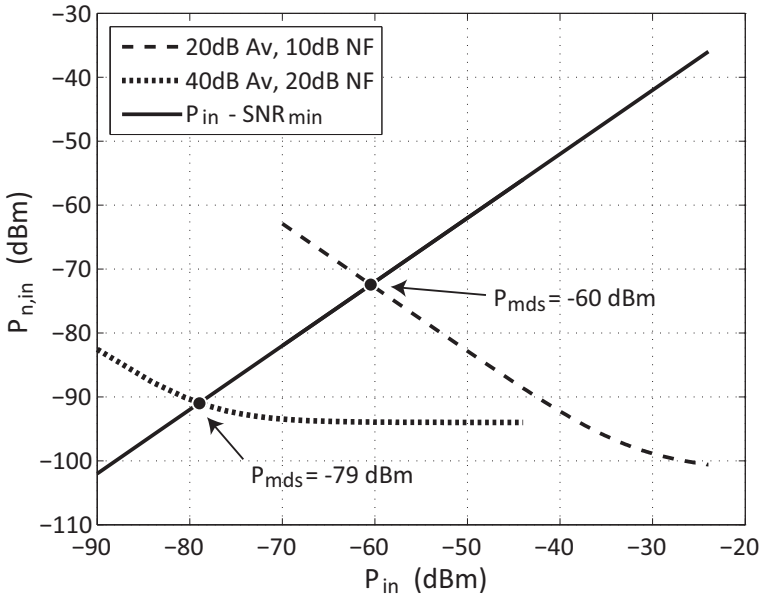

Figure 3.7: Effect of amplifier gain and *NF* on envelope detection receiver sensitivity

The receiver with higher gain has almost 20 dBm better sensitivity, despite 10 dB extra *NF* in the front-end. The example clearly illustrates the benefit of increasing gain in the front-end of an envelope detection receiver, even if the increase in gain results in degraded front-end noise performance. With these general principles established, the next section describes the receiver circuit design.

3.3 Receiver Circuit Design

For this first prototype, the TRF architecture is chosen for its simplicity and because no local oscillator is required. A block diagram of the proposed TRF receiver is shown in Figure 3.8. The architecture is similar to the one presented in [18], although the circuits must be re-designed to lower the power consumption. The input RF signal first passes through a matching network that embeds a BAW resonator to simultaneously filter the input with a sharp bandpass response. The front-end amplifier (FEA) then provides RF gain before the envelope detector, which yields the analog baseband signal. For a complete receiver, the baseband chain is also included, consisting of a low power analog/digital converter (ADC) driven by a programmable gain amplifier (PGA). A reference voltage generator is also implemented to produce the reference levels required by the ADC. With the exception of the ADC sampling clock, the blocks shown in Figure 3.8 constitute a complete receiver capable of listening for an RF wake-up signal.

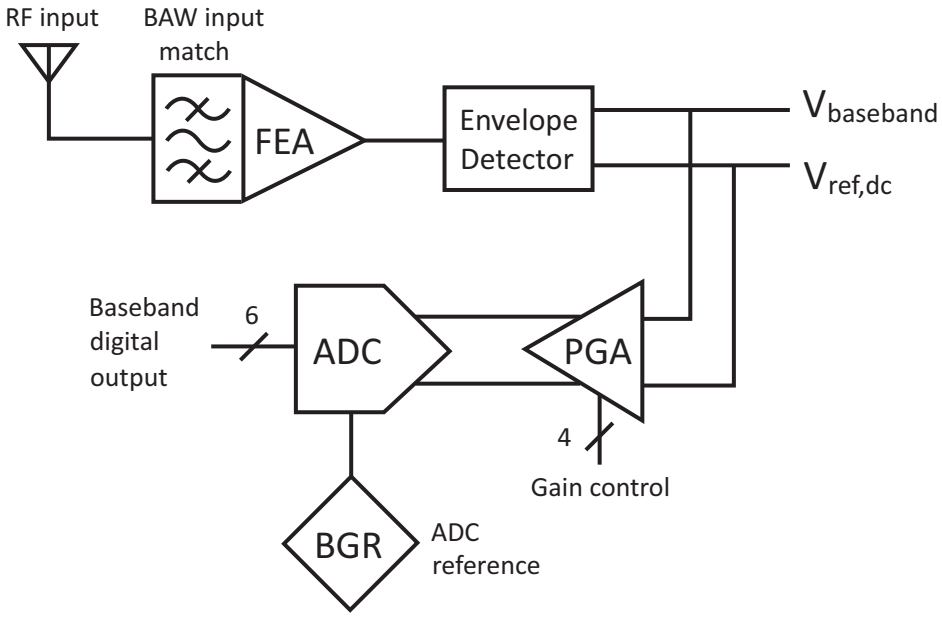

Figure 3.8: Block diagram of complete TRF receiver

3.3.1 Input Matching Network

The matching network serves two purposes. First, it must supply a stable impedance match to the 50 Ω input source. Second, the network should provide a narrow RF filter to remove out-of-band noise and interfering signals. From a filtering perspective, the high quality factor of the BAW resonator is an attractive choice. As shown in Section 2.2.4, the BAW circuit model contains a series resonant branch and large shunt capacitance C_o that dominates the response outside the narrow resonant frequency range. Between the series and parallel resonance frequencies, the resonator acts as a high quality inductive element. If used in series mode as a short circuit, C_o still allows signal

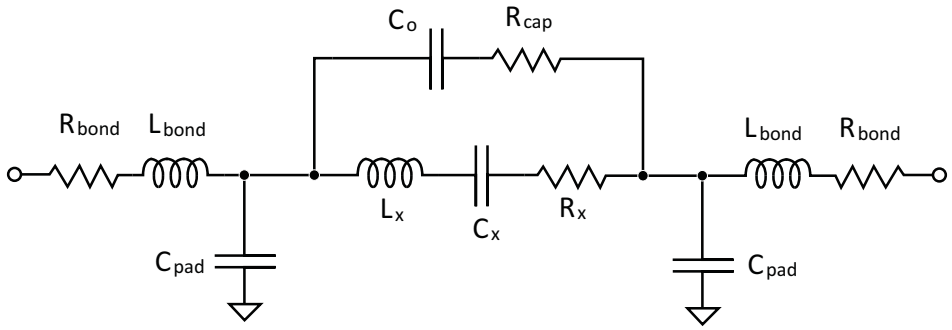

Figure 3.9: Complete resonator model including parasitics

feedthrough away from resonance. For this reason, if only a single resonator is to be used, it is better to use the parallel resonant mode to build a filter.

For actual design of the matching network, a more complex resonator model including parasitic effects is appropriate (Figure 3.9). There are a variety of possible packaging techniques and methods for connecting the resonator to the electronics, each with different effects on parasitic elements external to the resonator. For prototyping purposes, the FBAR chip is simply placed adjacent to the CMOS and wirebonded directly to pads on the CMOS die. In this configuration, the pad capacitance C_{pad} is about 100 fF and the short bonds (L_{bond}) can be modeled with about 500 pH of inductance. The quality factor of these bonds is quite high due to the short length and low loss, so a Q of 30 is assumed for design.

Compared with other common matching networks, a capacitive transformer is appealing because it contains no inductors, which are typically large and lossy when integrated on the CMOS die. Instead, the resonator itself can pro-

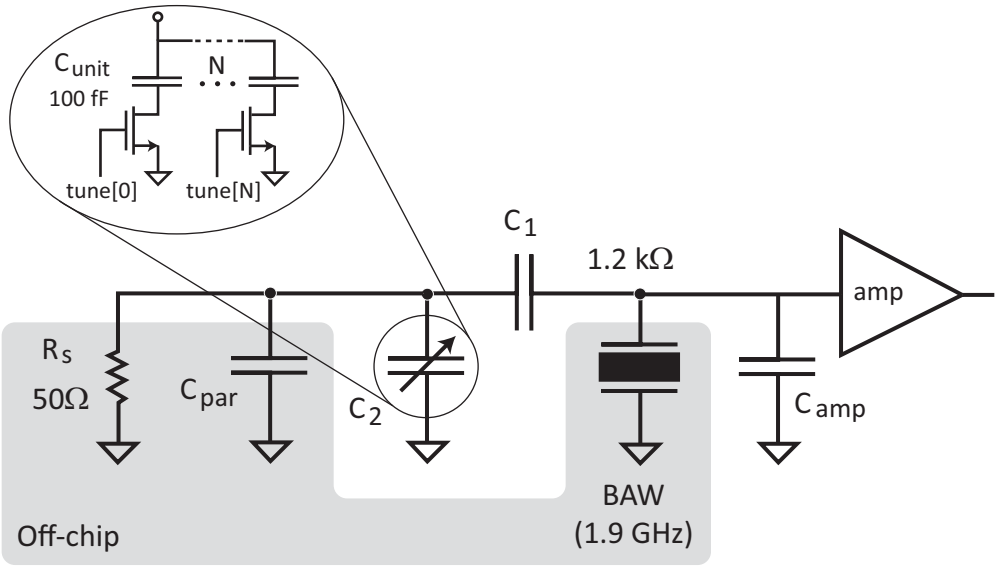

Figure 3.10: Schematic of BAW resonator input matching network

vide an inductance to resonate with the capacitive network. Furthermore, the capacitive parasitics of the resonator (C_o and C_{pad} in Figure 3.9) are conveniently lumped with capacitors in the transformer.

A schematic of the input matching network is shown in Figure 3.10, utilizing a 1.9 GHz resonator. Metal-insulator-metal (MIM) capacitors C_1 and C_2 transform the low antenna impedance up to match the resonator impedance. The input capacitance of the following amplifier stage C_{amp} can then be absorbed with the resonator capacitance, without requiring a real impedance at the amplifier input. The resonator C_o is about 1.5 pF, so the relatively small C_{amp} has little influence on network response.

Determining the optimal transformer ratio C_1/C_2 for the resonator is not as straightforward as the case with a real inductor. This is because, unlike a real

inductor, the equivalent inductance and parallel impedance of the resonator changes with shunt capacitance. From [25], the impedance of the resonator at its parallel resonance is:

$$R_p = \frac{1}{\omega_0^2 C_T^2 (R_x + R_{cap})} \tag{3.12}$$

where C_T is the total capacitance in shunt with the resonator, R_x is the resonator motional impedance, and R_{cap} represents the loss from the capacitive network including R_s.

Practical considerations dictate the value of C_2, which includes input parasitics from chip pads and the printed circuit board (PCB). For this design, the target value of C_2 is fixed at 1.5 pF and a switched capacitor network is included on-chip, digitally tunable between 0 and 1 pF. Any value of C_{par} between 500 fF and 1.5 pF can then be accommodated with the tunable C_2 (Figure 3.10). The final value of 700 fF for C_1 is verified with simulations to optimize the input match, using typical values for the BAW resonator model. The simulated $|S_{11}|$ is shown in Figure 3.11, including bondwires and pad parasitics. The matching network voltage gain is also shown in Figure 3.11. The amplifier input transistor is sensitive to voltage, so an additional benefit of the impedance transformation is approximately 12 dB of passive voltage gain [13]. The drawback of matching to the resonator impedance is that it presents a real resistance, thus degrading the noise figure by 3 dB compared to methods like inductive degeneration. As shown in Figure 3.7, the goal of

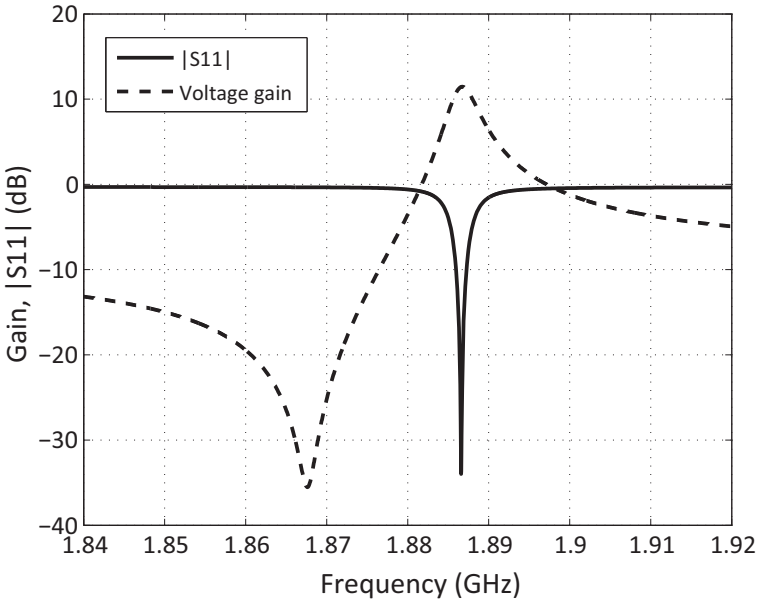

Figure 3.11: Simulated input match and voltage gain

maximum gain takes precedence over noise considerations for this amplifier, so the design choice is justified.

3.3.2 Front-end Amplifier Design

The FEA is a critical block in the receiver, since the gain and power consumption of this element will largely determine the overall performance of the receiver. The main objective is to provide the maximum possible gain while staying within the 50 µW power budget. As shown in the analysis in Section 3.2, noise performance is a secondary concern.

A schematic of the amplifier is shown in Figure 3.12. In order to conserve voltage headroom, the bulk terminals of devices M_1, M_2, M_3, and M_4

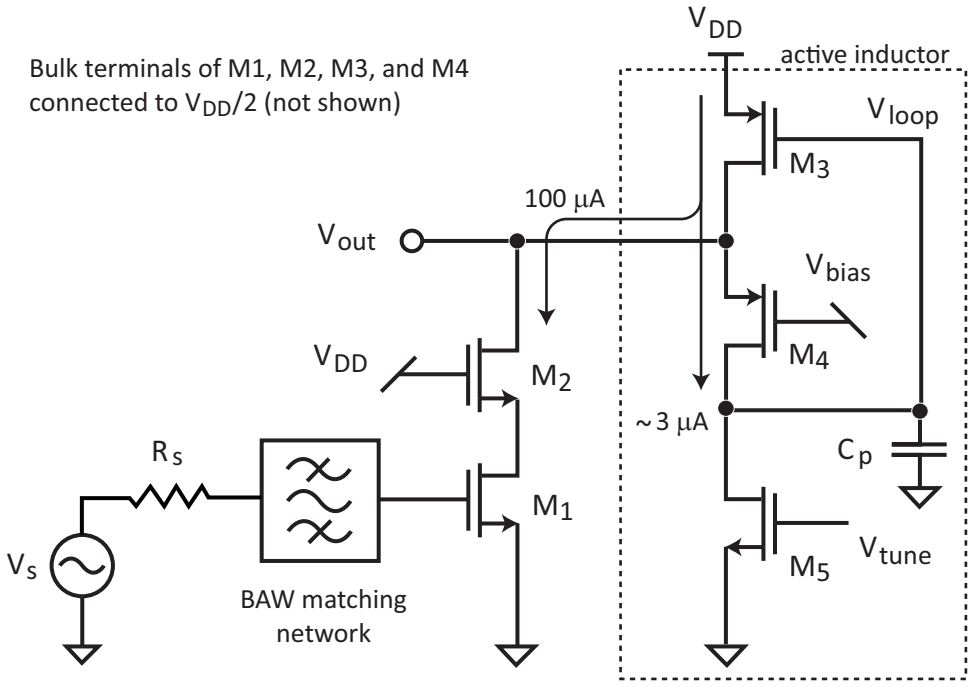

Figure 3.12: Schematic of front-end amplifier

are connected to $V_{DD}/2$ to lower the threshold voltage. With a 0.5 V supply, the forward bias voltage is small and there is no danger of turning on the source-bulk junction. After the matching network, devices M_1 and M_2 form a standard cascode transconductor, with the input capacitance of M_1 absorbed into the resonator C_o in the matching network. Device M_1 is sized (16/0.1) µm/µm for a current density of about 6 µA/µm, which biases the device in moderate inversion. The maximum available device f_t is not needed for low GHz frequencies, and this region of operation offers higher transconductance efficiency ($\frac{g_m}{I_d} \approx 19$). The cascode device M_2 is sized only half as wide in order to reduce capacitive loading at the output.

Typically, RF amplifiers utilize an inductor to form a resonant load network. The inductor resonates with the load capacitance, providing a high equivalent impedance without consuming any voltage headroom. However, as described in Section 2.2.1, the equivalent impedance at resonance is limited. For this reason an active inductor structure [35], comprised of M_3, M_4, and M_5 in Figure 3.12, is chosen for the load network of the FEA. The bulk of the cascode bias current flows through M_3, but a small fraction is drawn through M_4 and M_5 as determined by the tuning voltage V_{tune}. The key point is that this network can synthesize a higher impedance at the RF frequency than an on-chip inductor. This particular active inductor topology is chosen because the three stacked devices interface conveniently with the cascode and M_3 requires headroom of only a saturation voltage V_{dsat}, instead of a full V_{GS}.

The behavior of the active inductor circuit can be understood by investigating the impedance looking into the source of M_4. At low frequencies, the impedance is low due to the loop gain provided by M_3. As the frequency increases, the loop gain is attenuated by any capacitance C_p at the loop node V_{loop}, causing the impedance to increase. The small-signal model of the structure is shown in Figure 3.13. It can be shown that the input impedance is given by:

$$Z_i(s) = \left(\frac{r_{o3}}{1 + sr_{o3}C_o}\right) \| \left(\frac{g_{o4} + g_{o5}}{g_{m3}g_{m4}} + s\frac{C_p}{g_{m3}g_{m4}}\right) \tag{3.13}$$

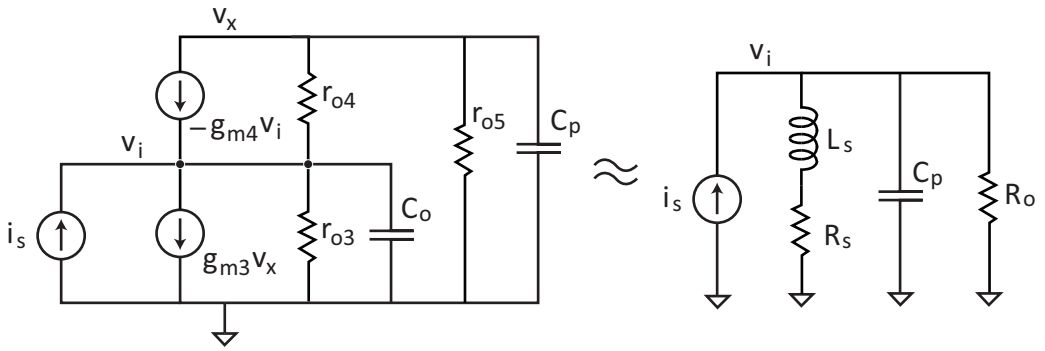

Figure 3.13: Active inductor small-signal model

Therefore, the tank model parameters are:

$$\begin{aligned} L_s &= \frac{C_p}{g_{m3}g_{m4}} \\ R_s &= \frac{g_{o4}+g_{o5}}{g_{m3}g_{m4}} \\ C_o &= \text{ output shunt capacitance} \\ R_o &= r_{o3} \| \frac{1}{g_{m4}} \end{aligned} \tag{3.14}$$

Equations 3.13 and 3.14 provide intuition for design, with the goal of maximizing the impedance at 1.9 GHz. The final design values are obtained taking into account the output loading of the envelope detector and verified with simulations. The device sizes and bias conditions are shown in Table 3.2. Using the parameters in the table with estimated output loading C_o of 20 fF and C_p of 10 fF, the calculated impedance is shown in Figure 3.14. A simulation of the final structure while embedded in the amplifier and driving the detector is also plotted. The discrepancy is about 3 dB at the peak, which is most likely

Table 3.2: Final design values for active inductor

Device	W/L (µm/µm)	Current (µA)	g_m (µS)	r_o (kΩ)
M_3	4/0.1	100	803	3.5
M_4	6/0.1	3	54	207
M_5	0.12/0.1	3	32	107

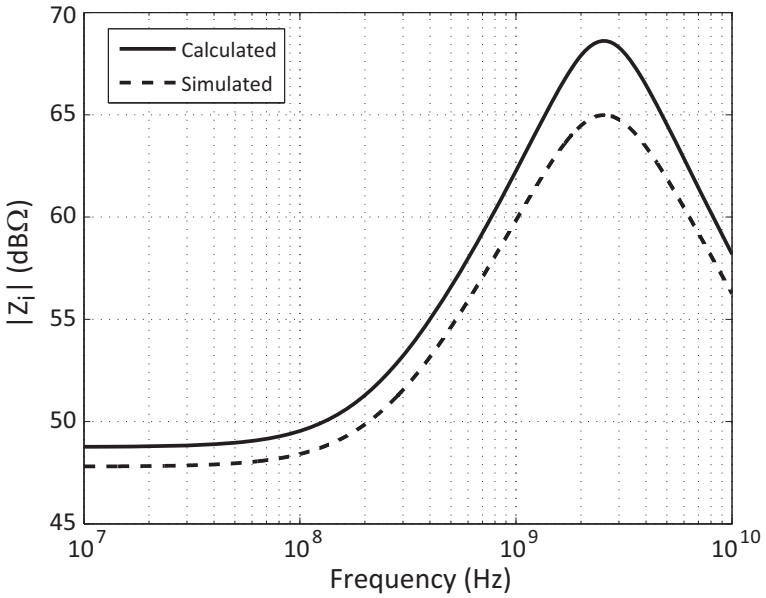

Figure 3.14: Active inductor input impedance response

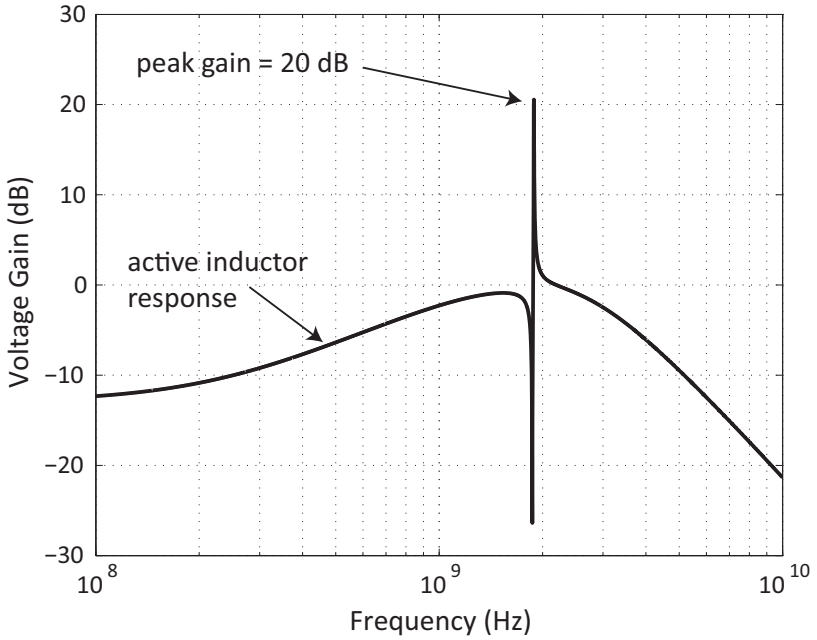

Figure 3.15: Simulated FEA voltage gain response

caused by departure from ideal transconductor behavior in the small-signal device models used for calculations. The low supply voltage means that the devices operate close to the triode region with low V_{DS}. Nevertheless, the active inductor realizes about 1.8 kΩ of impedance at 2 GHz. A real 10 nH inductor with Q of 15 would be required to provide the same load impedance, which would be difficult or impossible to achieve on-chip.

A simulation of the overall AC gain response of the amplifier, including matching network, is shown in Figure 3.15. The sharp resonance of the BAW filter is clearly visible, with the peak gain occuring at the parallel resonance. In the background, the low Q response of the active inductor rolls off the gain

at low frequencies. The overall gain is about 20 dB for a current consumption of 100 µA.

Although the use of an active inductor increases the gain of the FEA, the penalty is increased noise for the active structure over a real inductor. This penalty can be quantified by defining an excess noise factor β:

$$\beta = \frac{\overline{v_{o,n}^2}}{4kTR_{eff}} \tag{3.15}$$

where $\overline{v_{o,n}^2}$ is the noise current density at the active inductor output and R_{eff} is the effective resistance synthesized by the structure at resonance. The output noise voltage of the active inductor is dominated by M_3 and M_5 at resonance:

$$\overline{v_{o,n}^2} = 4kT\gamma \left(g_{m3} R_{eff}^2 + g_{m5} R_x^2 \right) \tag{3.16}$$

where R_x is the transresistance gain from the thermal drain noise of M_5 to the output voltage. The transfer function R_x can be calculated from the small-signal model shown in Figure 3.16. The calculated R_x for the design values in Table 3.2 is also plotted in the figure. Table 3.3 compares the calculated output noise with simulations. The simulations match calculations only if the noise parameter γ is unusually small. This discrepancy was checked with noise simulations of single devices and several bias conditions, confirming that the noise models produce noise corresponding to a γ value of about 0.25 or less. Although this is unrealistically low, it confirms the validity of the noise

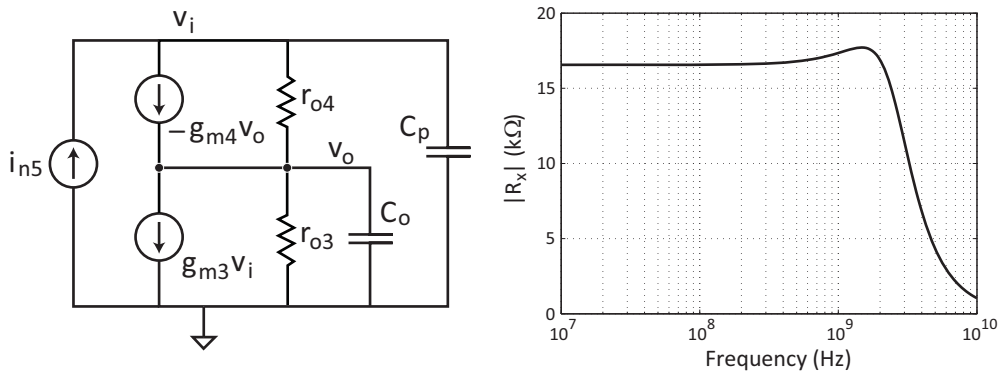

Figure 3.16: Small-signal model and calculation for M_5 noise contribution

Table 3.3: Active inductor output noise breakdown

Device	Simulation (V^2/Hz)	Calculation, $\gamma = 1$ (V^2/Hz)	Calculation, $\gamma = 0.25$ (V^2/Hz)
M_3	8.4e-18	4.3e-17	10.8e-18
M_5	4.1e-17	1.5e-16	3.8e-17

analysis presented here. Using Equation 3.15 and the values in Table 3.3, the excess noise factor β for the active inductor is 1.6 if $\gamma = 0.25$ and 6.4 if $\gamma = 1$.

The noise figure of the complete amplifier can now be calculated using Equation 3.17:

$$F = \frac{1}{A_v^2} \frac{\overline{V_{o,n}^2}}{4kTR_s} \tag{3.17}$$

where A_v is the overall voltage gain of the amplifier, including passive voltage gain in the matching network. The total output-referred noised density $\overline{V_{o,n}^2}$ is given by:

$$\overline{V_{o,n}^2} = 2(4kTR_s)A_v^2 + (4kT\gamma g_m)R_{load}^2 + \beta(4kTR_{load}) \tag{3.18}$$

where R_{load} is the effective resistance of the active inductor load at the RF frequency. Thus, using Equation 3.17, the noise factor is:

$$\begin{aligned} F &= \frac{1}{A_v^2} \frac{2(4kTR_s)A_v^2 + (4kT\gamma g_m)R_{load}^2 + \beta(4kTR_{load})}{4kTR_s} \\ &= 2 + \frac{\gamma g_m R_{load}^2}{A_v^2 R_s} + \frac{\beta R_{load}}{A_v^2 R_s} \end{aligned} \tag{3.19}$$

The first term in Equation 3.19 is a factor of two noise penalty, due to the input match to the real resonator impedance. The second two terms represent the noise stemming from the main transconductor device M_1 and active inductor load, respectively. Using $\gamma = 1$, which is expected to give a more realistic noise estimate for comparison with measurement, Equation 3.19 yields a noise figure of 12 dB.

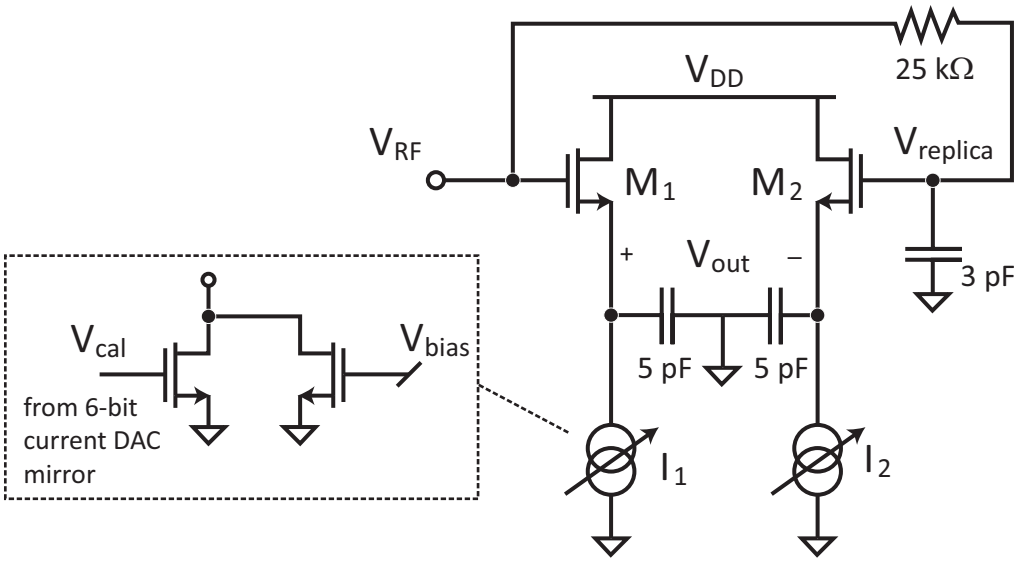

Figure 3.17: Schematic of envelope detector with offset calibration

3.3.3 Envelope Detector

The envelope detector design uses the same topology as described in Section 3.2.1. The detailed circuit schematic is shown in Figure 3.17. The detector device M_1 is sized with W/L of (5/0.25) µm/µm and M_2 is an identical device set up as a DC replica path. In order to derive a reference level for the ADC, the replica path filters the input with an RC to match the DC levels at V_{RF} and $V_{replica}$. Offset between the signal path and reference path can be removed via digital calibration of the tail currents I_1 and I_2 of the two detector paths. The tail currents share a bias voltage to set the primary bias current in M_1 and M_2, but have independent fine-tuning via current mirrors from a 6-bit current DAC. The bias DACs can then be used to make slight

changes to the DC level at the output and calibrate offset due to mismatch in the detector. This scheme also provides a simple way to remove input-referred offset of the continuous-time PGA. The nominal bias current in each branch is 800 nA and adjustable over an additional 800 nA in 12.5 nA increments, which translates to an LSB offset step of about 250 µV. For this prototype, offset is canceled manually before testing. In the final design, an offset calibration algorithm could easily be implemented in the digital domain using ADC samples to adaptively adjust the bias DACs.

3.3.4 Programmable Gain Amplifier

The PGA drives the ADC and provides some level of gain control in order to utiltize the full dynamic range of the ADC. A simplified schematic of the amplifier is shown in Figure 3.18. Due to the low supply voltage, a two stage architecture is chosen to minimize the number of stacked transistors. Miller compensation with a zero cancellation resistor is used to ensure amplifier stability. Miller capacitor C_c is a 50 fF MIM device and the 12.5 kΩ R_z is implemented with a $p+$ polysilicon resistor. In order to stabilize the output common mode voltage, a common mode feedback (CMFB) network senses the output common mode with two PMOS devices and adjusts the bias of the load devices in the first stage. Although not as robust or accurate as a typical CMFB using an auxiliary OTA, the simplified CMFB scheme is efficient to implement with the low supply voltage. To increase the device output resistance r_o, all the PGA transistors use channel lengths of 0.35 µm.

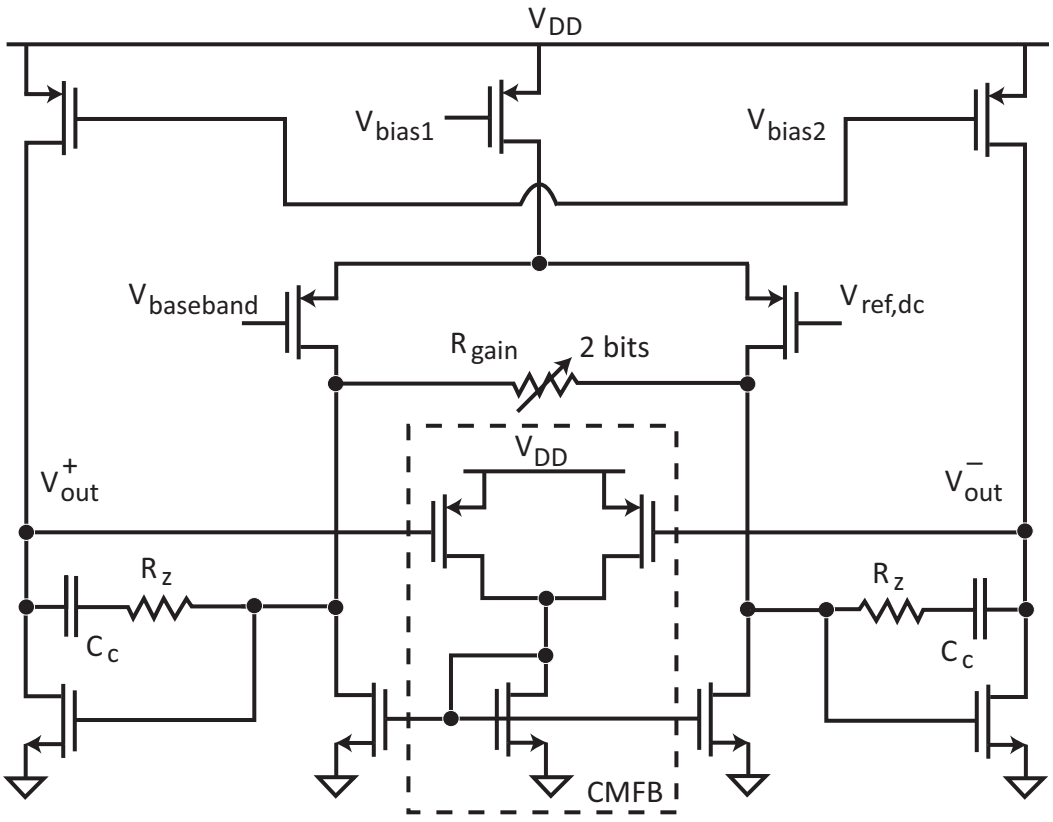

Figure 3.18: Schematic of programmable gain amplifier

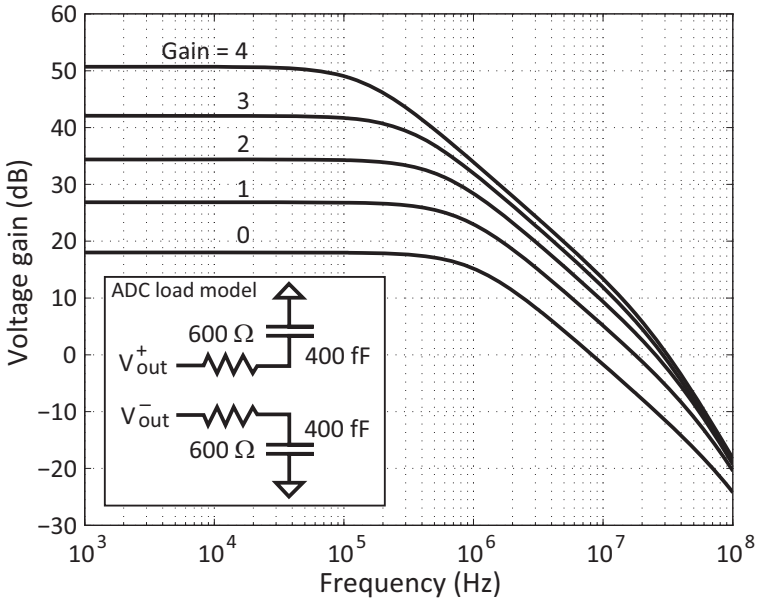

Figure 3.19: Simulated PGA gain for different gain settings

Gain control is accomplished by varying the load resistance of the first amplifier stage. The resistors are implemented with triode devices to realize large resistances without consuming excessive area. An AC gain simulation of the PGA driving the ADC is shown in Figure 3.19 for each of the five programmable gain settings. The ADC input sampling circuit is modeled with the network shown in the figure inset. The PGA gain is programmable from 18 to 50 dB in approximately 8 dB steps. The -3 dB bandwidth is at least 100 kHz across all gain settings, which is more than enough for the nominal 40 kbps data rate, and should be adequate up to 100 kbps. The power consumption of the PGA when driving the ADC is 2.5 µW.

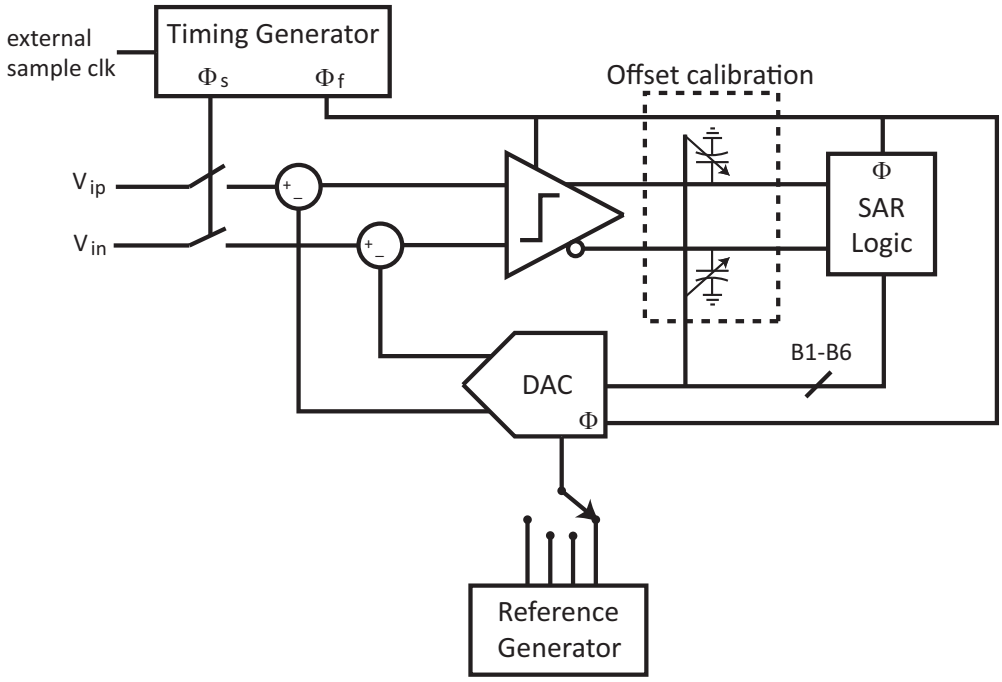

Figure 3.20: Block diagram of baseband ADC

3.3.5 ADC Design

The ADC was designed by Simone Gambini[2] and is capable of sampling at 1 MSample/s with 6 bit resolution. The complete details of the design are available in [36], but some of the relevant characteristics are summarized here. The block diagram of the converter is shown in Figure 3.20. A successive approximation register (SAR) architecture was employed because it is a good fit for the relatively low resolution and sampling rate requirements of the receiver baseband. Furthermore, the SAR architecture uses a comparator rather than high gain linear amplifiers, making it more easily compatible with the 0.5 V

[2]Department of EECS, UC-Berkeley

supply. For testing purposes, the ADC sampling clock is provided externally at 16 times the desired sampling rate.

In order to reduce power in the PGA driving the ADC, the converter input is optimized for low input capacitance in two ways. First, the switched-capacitor feedback DAC uses tri-level unit elements instead of the usual binary elements, which has the effect of halving the required DAC capacitance [36]. In addition, the capacitors in the DAC are implemented using vertical capacitors between Metal5 and Metal6, which have lower capacitance density per unit area than alternative devices such as MIM capacitors. This allows the DAC elements to be sized large enough to meet matching requirements while minimizing the capacitance of each element. The result is a differential input capacitance of just 155 fF for the ADC, which is easily driven by the PGA with mininmal power requirements.

3.3.6 Reference Generator

As mentioned in Chapter 1, one of the requirements for a practical WuRx design is that all the necessary components are included for complete functionality while the rest of the node's electronics are asleep. Therefore, the ADC reference generation and its power requirements cannot be ignored. As a companion to the ADC, the reference was also designed and implemented by Simone Gambini[3].

The schematic of the reference generator is shown in Figure 3.21 and the

[3]Department of EECS, UC-Berkeley

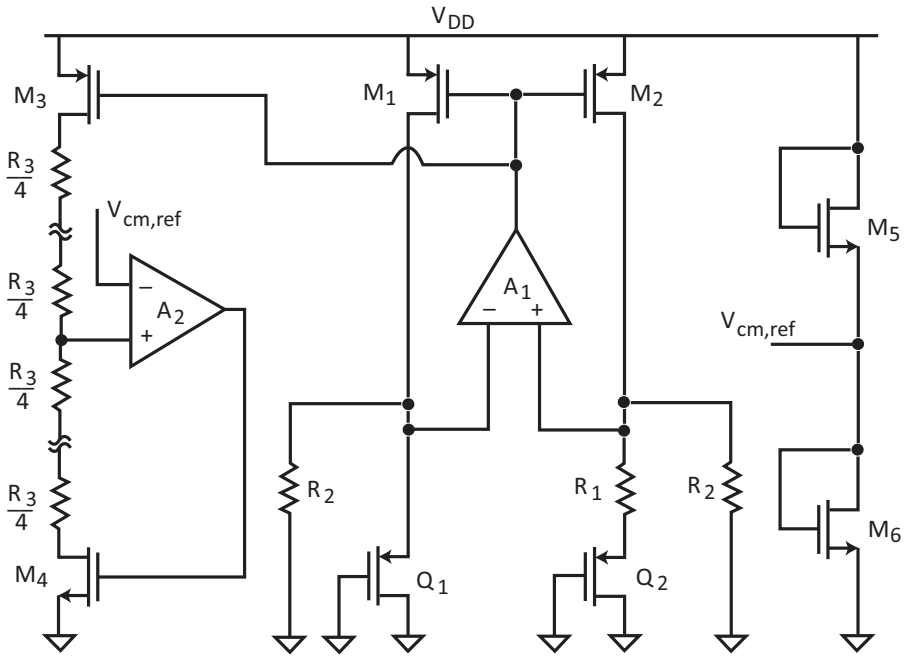

Figure 3.21: Schematic of ADC reference generator

design is described in detail in [36]. The main challenge in the design of the reference is that the 0.5 V supply does not accomodate the V_{be} drop normally used in bandgap references. This design makes use of subthreshold PMOS devices Q_1 and Q_2 [37] and a resistive division technique [38] to provide an output lower than the silicon bandgap voltage. In addition, the output is made programmable by dividing R_3 into a tapped resistor string and digitally selecting the desired differential output for the ADC. Varying the reference voltage changes size of the least-significant bit (LSB) and effectively changes the DC gain of the ADC. The effect is illustrated in Figure 3.22, where the measured transient output samples from the receiver are plotted for two different ref-

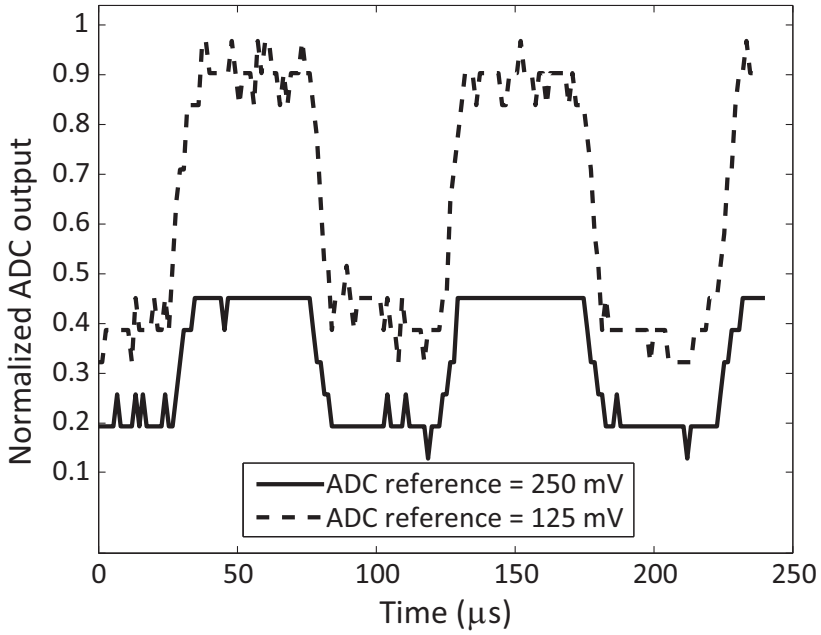

Figure 3.22: Measured ADC output for different reference settings

erence voltage settings. With the reference voltage configured for 250 mV, the LSB size is maximized and the DC gain is minimized. The LSB size is reduced when the reference is set to 125 mV, increasing the effective gain. With four possible output settings, the programmable reference yields about 12 dB of additional programmable gain in the receiver chain. The total power consumption of the reference is 11 µW, which is largely determined by settling time constraints as the reference output charges the feedback DAC of the successive approximation ADC. The average simulated temperature coefficient is 136 ppm/℃.

3.3.7 Design for Testability

The complete receiver, including RF front-end through the baseband and ADC, is a complex system with several functional settings and biasing required for multiple blocks. One way to simplify the testing and reduce the amount of instrumentation required is to make these settings digitally tunable. Some adjustments, such as reference voltage output and input match tuning, are controlled by switch networks and are inherently digital. The bias currents and voltages, however, require a digital-to-analog converter (DAC). For this purpose, a single general-purpose DAC is designed and then used throughout the chip. A schematic of the design is shown in Figure 3.23. A simple current mirror-based topology is used, where the gate-source voltage from a reference mirror is distributed through a switch network to a device array, made up of unit devices grouped into binary-weighted sets. The current range is from 0 to 63 µA with LSB steps of 1 µA, controlled via a digital word applied to the bitlines. The unit devices are sized with a length of 0.35 µm to improve r_o and the overall output resistance of the current DAC. Cascoding the output devices is not possible due to the low supply voltage, but linearity is not a primary concern for the intended bias application.

With the basic DAC building block, bias voltages and currents for all receiver blocks can be implemented using current mirrors to expand or compress the DAC current range as needed. A digital serial peripheral interface (SPI) and register set is included on the chip to receive configuration words from a

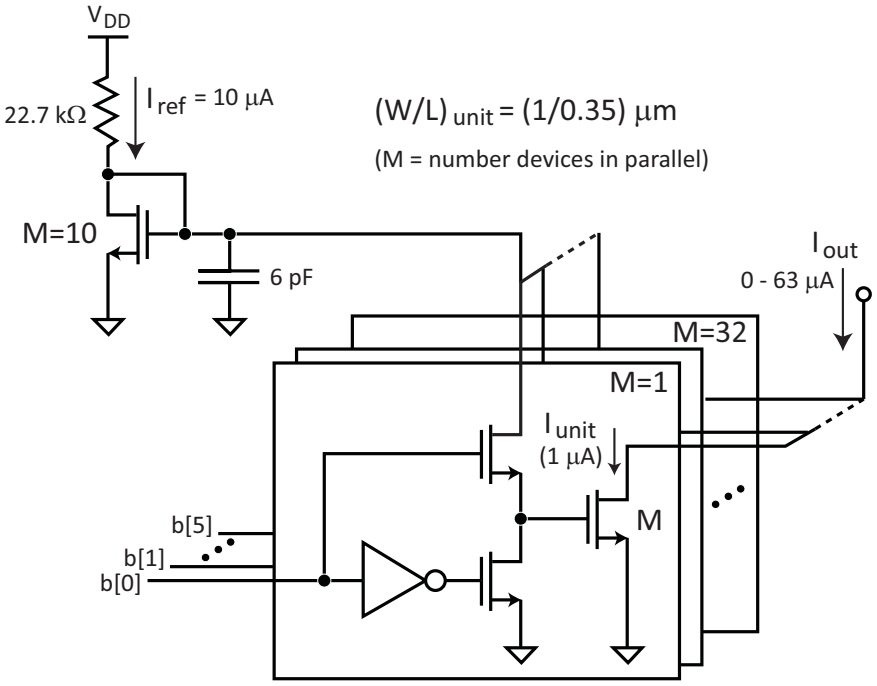

Figure 3.23: Schematic of current DAC for bias generation

laptop computer and store the digital settings on-chip. This digital approach to control the receiver saves I/O pads and simplifies testing by reducing the number of supplies needed for biasing.

The chief drawback of this flexibility in testing is the power penalty incurred by using a single general purpose bias DAC across the entire design, which is not included in the receiver power measurements. After the first prototype design and characterization, much of the testing flexibility is not needed for the final implementation. For example, the measured FEA performance was found to be relatively insensitive to the precise bias voltages used in the active inductor. For a real implementation, the receiver bias circuits could be optimized and re-designed for much lower power.

3.4 Measurement Results

The WuRx prototype is fabricated in 90 nm standard digital CMOS with the MIM capacitor option. A micrograph is shown in Figure 3.24, using a standard chip-on-board (COB) technique. The CMOS die is glued onto the circuit board and wirebonds are made directly from the chip pads to landing sites on the PCB. For prototyping purposes, the packaged resonator is simply connected to the CMOS die using wirebonds. The COB packaging allows convenient connections between the CMOS and MEMS chips.

The BAW package actually contains two separate resonators, only one of which is required for this design. The inset shows the unpackaged resonator

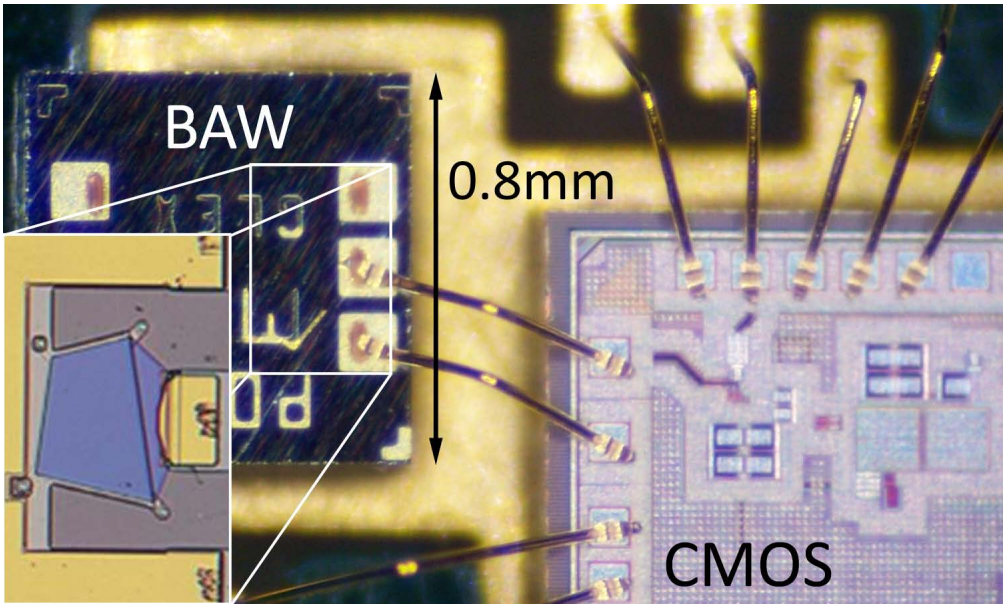

Figure 3.24: Die photo of CMOS prototype bonded to packaged BAW

with the corresponding size scale. On the CMOS side, Figure 3.25 shows a magnified veiw of the active die area. The active area is approximately 76,000 μm^2, of which about 20% is taken by the bias DACs and associated decoupling capacitance. The capacitor feedback DAC used by the ADC is largest single block. Due to the small number of transistors and lack of inductors in the FEA, the RF front-end consumes negligible area compared to the baseband circuits.

3.4.1 Standalone Front-end Amplifier

A standalone front-end amplifier was included on the test chip in order to characterize the amplifier performance. To facilitate S-parameter measurements

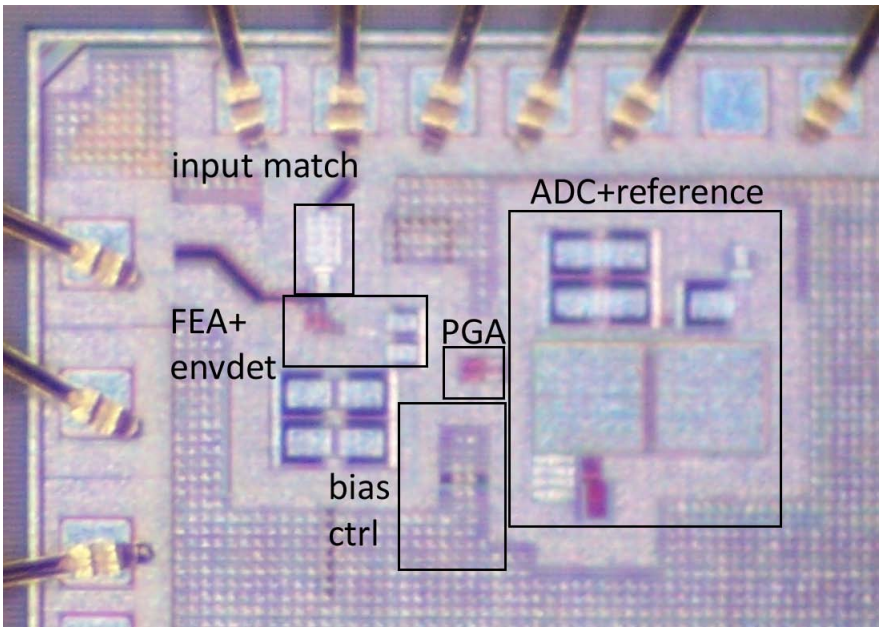

Figure 3.25: Annotated die photo

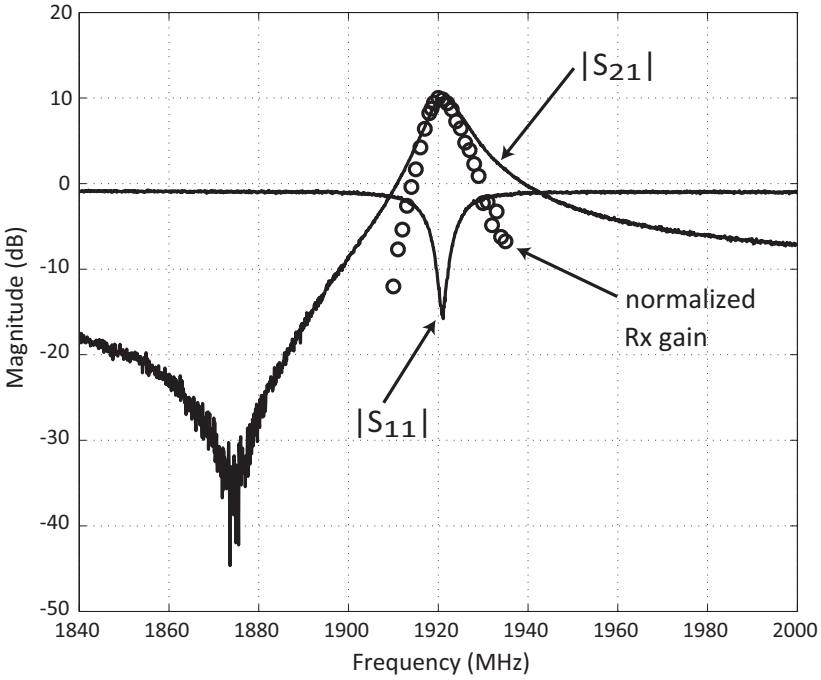

Figure 3.26: Measured FEA S-parameters and normalized gain to baseband

with a network analyzer, a 50 Ω output buffer is included on the chip to drive the amplifier output. In the complete receiver, the FEA output is loaded only by the capacitance of the envelope detector. In order to match the loading of the test amplifier, the buffer design uses a two-stage topology to provide a 50 Ω output without presenting excessive capacitance to the amplifier. The amplifier S-parameters were measured using a Hewlett Packard (HP) 8717C network analyzer and the results are shown in Figure 3.26. For this measurement the amplifier was biased at its nominal operating point with 100 µA. The input match is about -15 dB and the peak $|S_{21}|$ is 10 dB, with the resonator response clearly visible in the measurement. Using simulations to estimate the loss due

Table 3.4: Comparison of FEA measurements to simulation

Parameter	Simulated	Measured
$\|S_{21}\|$ (dB)	15	10
Voltage gain A_v (dB)	21	≈17
Noise Figure (dB)	8	10.3

to the output buffer, the *in-situ* voltage gain for the amplifier when driving the envelope detector is approximately 17 dB. The noise figure of the standalone amplifier was measured at 10 dB with an Agilent N8974A *NF* tester and HP 346C noise source. A comparison of measurements and simulations is shown in Table 3.4. The gain is a few dB less than expected, which could be due to excessive capacitance in the active inductor load network or deviation from the simulation model used in the matching network. The calculated *NF* is between simulations and measurements, and is most likely due to uncertainty in γ and lower than expected FEA gain.

Figure 3.26 also plots the normalized gain of the complete receiver all the way to baseband. The gain is determined by applying an RF carrier with square wave amplitude modulation to the receiver input and calculating the amplitude of the baseband square wave from the digital output. The measurement shows that RF bandwidth of the complete receiver is less than that of the amplifier itself. This effect is evidence of the nonlinear gain of the envelope detector as explained in Section 3.2.1. As the input frequency

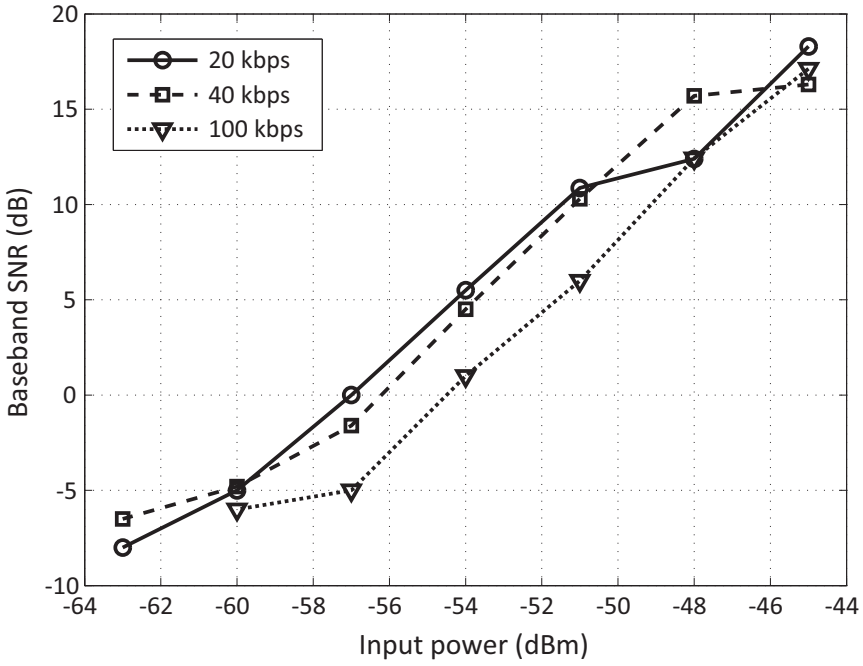

Figure 3.27: Calculated baseband SNR for different data rates

moves off the peak, the amplifier gain falls. However, the detector gain falls even more with decreasing input signal. The overall bandwidth is narrowed to about 7 MHz.

3.4.2 Receiver Sensitivity

The raw sensitivity of the receiver is quantified by calculating the signal-to-noise ratio (SNR) of the baseband digital samples in the presence of an on-off keyed (OOK) RF input. An alternating series of ones and zeros is used for the modulation signal input and the SNR is calculated offline in Matlab. Figure 3.27 shows the resulting SNR as the input power varies. Measurements

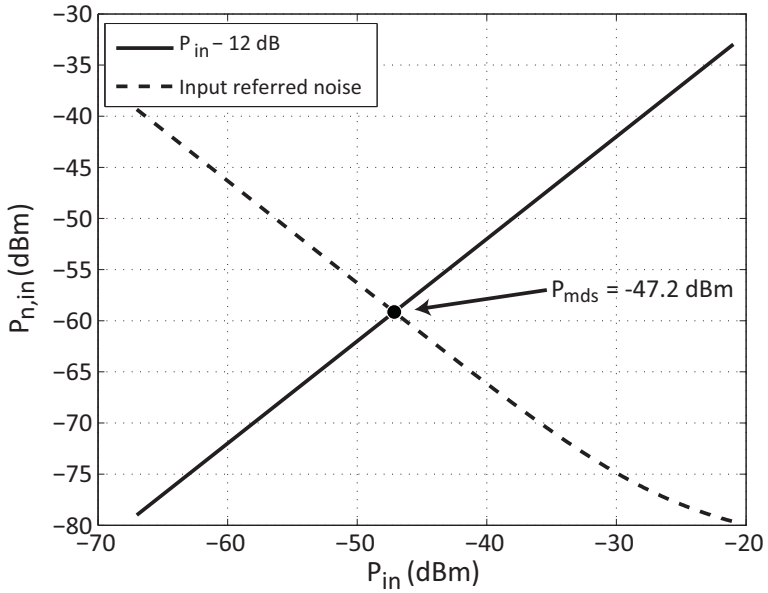

Figure 3.28: Calculated sensitivity

are shown for 20, 40, and 100 kbps modulation rates. A baseband SNR of about 12 dB is typical for reliable detection of OOK data. Therefore, Figure 3.27 shows that the sensitivity for this performance is about -49 dBm. As expected, the SNR degrades rapidly as input power decreases, again due to the nonlinear gain of the envelope detector.

Figure 3.28 shows the predicted sensitivity using Equation 3.11 and the analysis method outlined in Section 3.2.2. The envelope detector noise parameters are obtained from simulations and calculations, while the FEA gain and NF are measured results. The predicted sensitivity for 12 dB baseband SNR is -47.2 dBm, which is less than 2 dB from the measured value.

3.4.3 Digital Baseband and Wake-up Sensitivity

The raw sensitivity measurement is a valuable metric for comparing the receiver performance with other general-purpose wireless receivers. For a wake-up receiver, however, a better metric of interest is the rate of false alarms and missed detections of the wake-up sequence. As mentioned in Chapter 1, the wake-up receiver will be more useful if it is more than a simple energy detector and instead is able to recognize a particular sequence of bits.

In order to trigger on a bit pattern, the receiver requires additional baseband processing on the digital samples. For this test, long transient captures of the ADC output are saved and processed off-line with a digital baseband implemented in Matlab. Figure 3.29 shows a diagram of the complete measurement setup. First, the pattern generator is programmed to output a particular pseudo-noise (PN) code sequence of length N, wait at least N cycles, and then repeat the wake-up sequence. For example, the wake-up sequence is programmed to be 111010 for $N = 7$. The data bits are modulated onto the RF carrier using OOK modulation and this signal is fed to the WuRx prototype. The receiver ADC is configured to sample at 4 times the bit-rate R_b, but there is no synchronization between the transmitted signal and the WuRx sampling clock. The samples from the ADC are captured and saved into a file by a logic analyzer. The files can then be loaded in Matlab and run through the baseband off-line.

The digital baseband processing must be able to recognize the target wake-

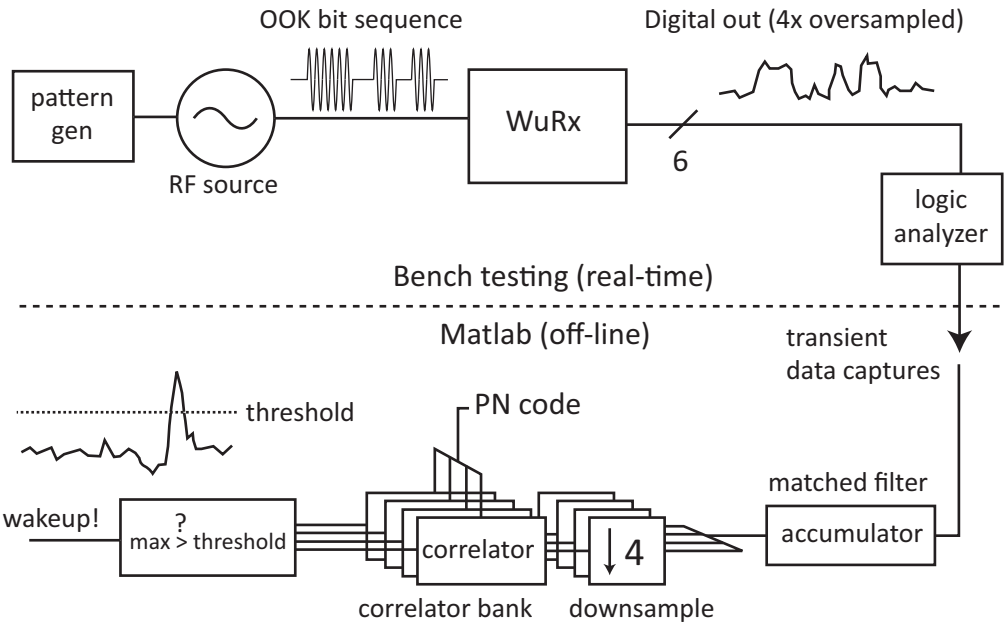

Figure 3.29: Measurement setup for wake-up sensitivity

up sequence with minimum complexity. The architecture is based on the timing estimation algorithm described in [39], where a baseband synchronization system is described for OOK receivers. Detection of the wake-up sequence is similar to timing estimation using a packet preamble. First, the signal passes through a matched filter at full rate ($4R_b$), which is simply an accumulator for an OOK signal. Next the signal is downsampled by a factor of four and sorted into four parallel paths at the original rate (R_b). This results in four copies of the signal, each sampled with a phase-shifted clock. Each correlator compares its own shifted input signal to the target sequence and generates an output peak for a match. In a full synchronization scheme, the correlation path with the best phase match would be used to sample the remaining bits in

the packet. For the wake-up baseband, however, there is no further data and a decision can be made if any one of the correlator outputs exceeds a decision threshold.

PN code sequences are chosen for the target wake-up sequences because they have low correlation with other sequences, which helps avoid false alarms (FA) due to other communicating nodes or wake-up signals intended for neighboring nodes. The threshold level setting for the correlator decision depends on the desired probabilities of detection (P_{det}) and FA. For these measurements, the input RF power level is swept and the correlator threshold is adjusted to maintain a constant P_{det} of 90%. At each input power level and threshold setting, FA occur at some average interval due to random noise. The simulation results are shown in Figure 3.30, where the average time between FA (T_{FA}) is plotted against input power for wake-up sequence lengths of 7, 15, 31, and 63 bits. As expected, FA occur less often as input power increases, but the curves are very steep due to sharp SNR roll-off of the receiver (Figure 3.27). The most relevant region is where T_{FA} is on the same order as the rate of packet traffic in the network. For example, if each node is receiving one packet per second on average and T_{FA} is 10 seconds, then FA will be very rare for all practical purposes. Using Figure 3.30, the sensitivity to a 31-bit wake-up sequence is -56 dBm for 90% P_{det} and T_{FA} of about 10 seconds.

From a practical standpoint, the measurements and Matlab simulations show that there is little benefit to be gained from coding using this envelope detection receiver. Even for a 31 bit code, sensitivity improves only by a

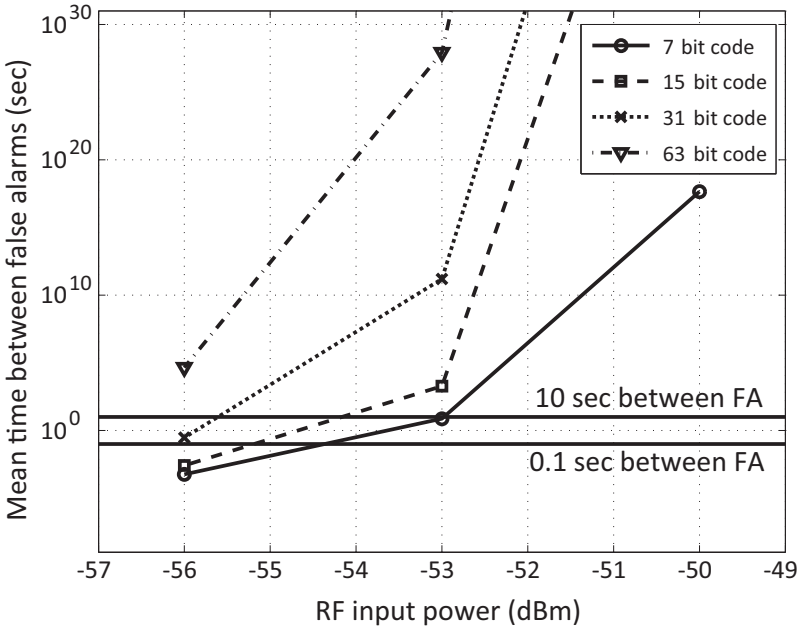

Figure 3.30: Mean time between false alarms for different sequence lengths

few dB. Nevertheless, this prototype receiver and the Matlab digital baseband provide a useful framework for evaluating the performance of the wake-up receiver and investigating different baseband implementations.

In an actual wake-up receiver implementation, the Matlab digital baseband would be synthesized to run on-chip using the standard CMOS libraries. In 90 nm CMOS, the power consumption of the digital circuitry is expected to be quite small, especially given the low frequency of operation. The correlator blocks are responsible for the bulk of the computation and run at the bit rate, which is only 40 kHz. In order to rougly estimate the power, a 20-bit adder simulation can be used because the average adder width in the correlator pipelines is 20 bits for a 31-bit sequence (6-bit ADC samples). The adder power is then simply scaled by the number of adders in the correlator and multiplied by four for the complete correlator bank. This estimation method assumes that the full digital precision is maintained throughout the correlation. In reality, the later stages could be truncated to reduce adder bit width.

The total simulated digital power using the standard CMOS library cells is about 14 µW when running at 40 kHz with a 0.5 V supply. This estimate is worst-case, since no architecture or precision optimization is performed. Of the total estimated power, about 98% is consumed by leakage because of the very low clock rate. Due to the low speed requirements, the baseband digital circuits should be optimized and deeply pipelined to reduce the leakage contribution [40]. Alternatively, the baseband implementation is an excellent candidate for a logic style optimized for low leakage and low supply voltage [41].

Table 3.5: Tuned-RF receiver performance summary

Parameter	Measurement
Global supply voltage (V)	0.5
Carrier frequency/modulation	1.9 GHz / OOK
Total power dissipation (µW)	65
RF bandwidth (MHz)	7
ADC performance	6 bit, 1 MSample/s
Date rate (kbps)	40/100 (nom/max)
Raw sensitivity for 12 dB SNR (dBm)	-49 (100 kbps)
Sensitivity to 31b sequence for 90% P_{det} and 10 FA/s (dBm)	-56

3.4.4 Measurement Summary

The overall performance of the prototype is summarized in Table 3.5 and Figure 3.31 shows the breakdown of power consumption among the different receiver blocks. The raw sensitivity is -49 dBm at 100 kbps while dissipating 65 µW from the 0.5 V supply. In order to overcome the nonlinearity of the envelope detector, the bulk of the power consumption is spent in the FEA to supply RF gain.

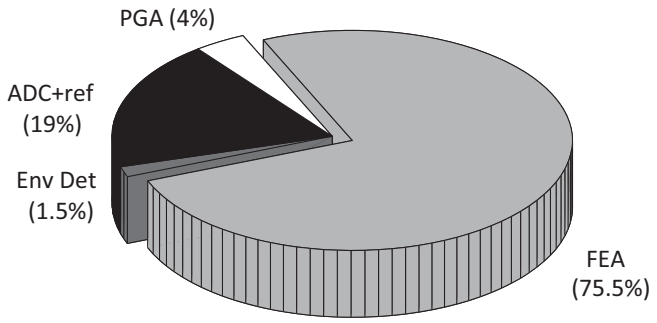

Figure 3.31: Receiver power consumption breakdown

3.5 Conclusion

This chapter described a first prototype of the wake-up receiver. All required circuitry is demonstrated, including the baseband section and ADC, with power dissipation well below 100 µW. Nevertheless, the sensitivity for this TRF prototype receiver is inadequate to receive RF signals over more than a meter or two. As shown in Section 3.2, the only way to improve sensitivity with this architecture is to add more RF gain. This is not an attractive prospect, given that the power consumption of the prototype is already dominated by the FEA. Some other method of increasing the gain without increasing power excessively must be found. In the next chapter, a novel receiver architecture is proposed to do just that.

Chapter 4

Uncertain-IF Receiver

The last chapter presented a receiver using the tuned-RF architecture to eliminate the need for a local oscillator. Although the power consumption meets the specification, the receiver sensitivity is inadequate due to the difficulty of efficiently realizing gain at high frequencies. In this chapter, an improved receiver prototype is described which uses a novel architecture to circumvent the gain limitations of the tuned-RF receiver.

4.1 Architecture Development

As described in Chapter 2, the generation of a local oscillator signal often sets the lower limit on power consumption for a wireless receiver. In order to address the problem of oscillator power consumption, it is useful to review the fundamental power limitations for oscillation.

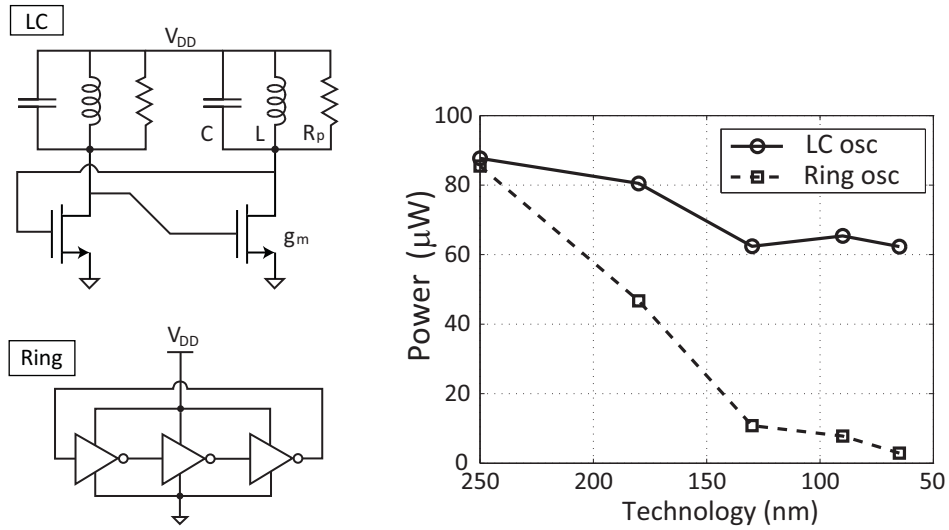

(a) LC oscillator and ring oscillator schematics

(b) Simulated power consumption

Figure 4.1: Effect of technology scaling on oscillator power

4.1.1 Oscillator Power Limitations

In order to overcome the losses in any real circuit implementation, an oscillator requires active gain sufficient to sustain oscillation. Section 2.2.1 showed that technology scaling has resulted in the input impedance of small devices surpassing the impedance available from integrated LC tanks. In the context of oscillators, this observation leads to the expectation that a simple ring oscillator, consisting of wideband inverting gain stages, can achieve lower power oscillation than its LC oscillator counterpart. Simulations confirm this expectation. Figure 4.1 compares the simulated power consumption of a simple 3-stage ring oscillator with an LC oscillator as technology scales. The ring oscillator V_{DD} is reduced as technology scales to maintain the frequency constant

at 2 GHz.

For the LC oscillator, the loop gain $A_l = g_m R_p$ must be greater than one for startup. The power consumption required for startup diminishes slightly due to the reduced device threshold voltage in scaled technologies, enabling lower supply voltages with the same bias current. However, since the power of a 3-stage CMOS ring oscillator scales with the total switched capacitance and the square of supply voltage, its power consumption drops much more rapidly. For modern 90 nm and 65 nm technologies, the 2 GHz ring oscillator results in about 20x power savings over an LC oscillator.

4.1.2 Uncertain-IF Architecture

The preceding comparison addressed only the minimum power required to achieve oscillation at RF frequencies, without considering phase noise or frequency accuracy. Of course, these are important considerations for frequency conversion architectures, and the ring oscillator is known to have inferior frequency stability compared with an LC oscillator [42]. However, the receiver presented here overcomes these problems at the architecture level, by employing an "uncertain-IF" to ease the phase noise and frequency accuracy requirements. The relaxed specifications allow the use of a free-running ring oscillator for LO generation.

The frequency plan and method of operation for the uncertain-IF architecture is shown in Figure 4.2. The desired signal is first filtered at the front-end to remove image and interferers. It is then mixed with an LO whose frequency

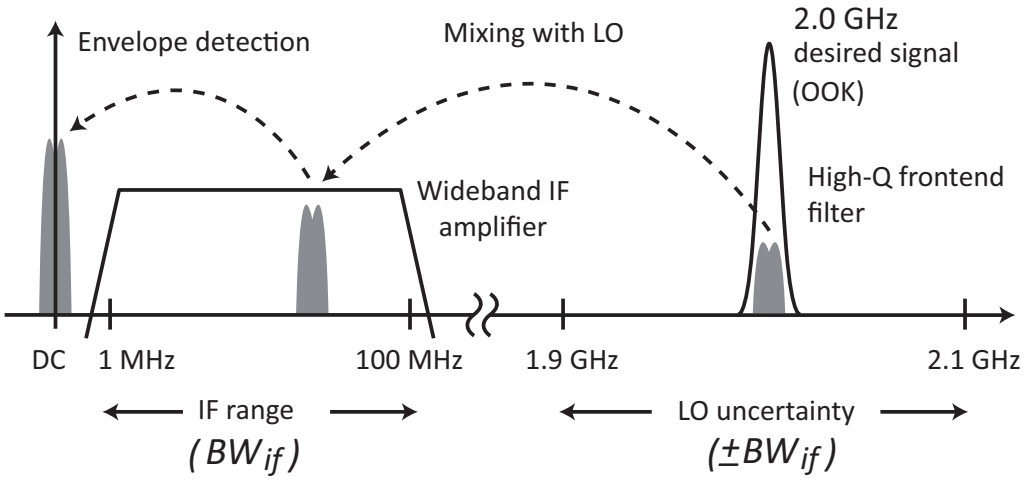

Figure 4.2: Uncertain-IF frequency plan and method of operation

is not well-defined. In fact, the LO must only be guaranteed to lie within some pre-determined frequency band $\pm BW_{if}$ ($\pm$100 MHz in this implementation) around the RF channel frequency. Due to the uncertainty of the LO frequency, the exact IF frequency will vary, but the downconverted signal will lie somewhere around DC within BW_{if}. The signal is then amplified at this IF frequency, which is much more power efficient than achieving the equivalent gain at RF. Finally, envelope detection performs the final downconversion to DC. Note that the use of envelope detection again limits the receiver to detection of amplitude-modulated signals, most commonly OOK, because the envelope detector removes all phase and frequency content in the IF signal.

As shown in Figure 4.2, AC coupling is used to limit the low end of the IF bandwidth to a frequency above the baseband bandwidth. For now, it is sufficient to mention that this bandwidth limiting is used to ensure proper

operation of the envelope detector and avoid the situation where the RF signal is directly converted to DC. However, the gain roll-off near DC means that the receiver cannot detect signals at the RF channel frequency if the LO frequency happens to fall on that channel frequency. The implications of this failure mode are discussed in more detail in Section 4.2.3.

The uncertain-IF architecture may be viewed as super-heterodyne, where the second downconversion is simply self-mixing, obviating the need for a precise LO at the IF frequency. For an ultra-low power receiver like the WuRx, the uncertain-IF architecture holds several advantages over the architectures described in Chapter 2. First, LO phase noise and frequency accuracy requirements are significantly relaxed. Frequency variation of the LO simply appears as IF frequency variation, to which the envelope detector is insensitive. An initial calibration step is only required to account for process variation and tune the LO within the desired range. As discussed earlier, it may also be necessary to adjust the LO to ensure that it does not coincide with the RF channel. Thereafter, re-calibration is required only to counteract frequency drift due to aging or temperature and supply variation. Furthermore, as in the heterodyne architecture, signal amplification can be performed at IF instead of RF, resulting in substantial power savings. The result is essentially performance similar to a TRF receiver with dramatically increased gain before envelope detection, improving performance compared to receivers using only RF gain.

Like any TRF receiver, however, a disadvantage of the uncertain-IF archi-

tecture is its susceptibility to interferers. Any undesired signal within $\pm BW_{if}$ of the LO frequency that passes through the front-end filter will be mixed down and detected by the envelope detector. Therefore, a narrow and accurate RF bandpass filter is required to improve robustness to interferers. In effect, the burden of selectivity has been shifted from the LO to the front-end filter. Here, as in the previous prototype, filtering is performed by a bulk acoustic wave (BAW) resonator.

One important architecture-level design consideration is the trade-off between LO tuning accuracy and IF bandwidth. If the LO can be tuned very close to the channel frequency, the required bandwidth of the IF amplifier can be narrowed and its power reduced proportionately. On the other hand, the LO must now be kept within a smaller frequency range, increasing vulnerability to oscillator frequency drift. If the IF bandwidth is made large enough, the receiver will be relatively immune to frequency drift and the LO will be able to run for long periods without calibration. For this implementation, a relatively wide IF bandwidth of 100 MHz is chosen to maximize tolerance of LO frequency drift, without requiring excessive power in the IF amplifier.

4.2 Circuit Design

A block diagram of the complete receiver is shown in Figure 4.3. The OOK input signal is first filtered by the matching network containing the BAW resonator, followed directly by the mixer. The resulting IF signal is amplified

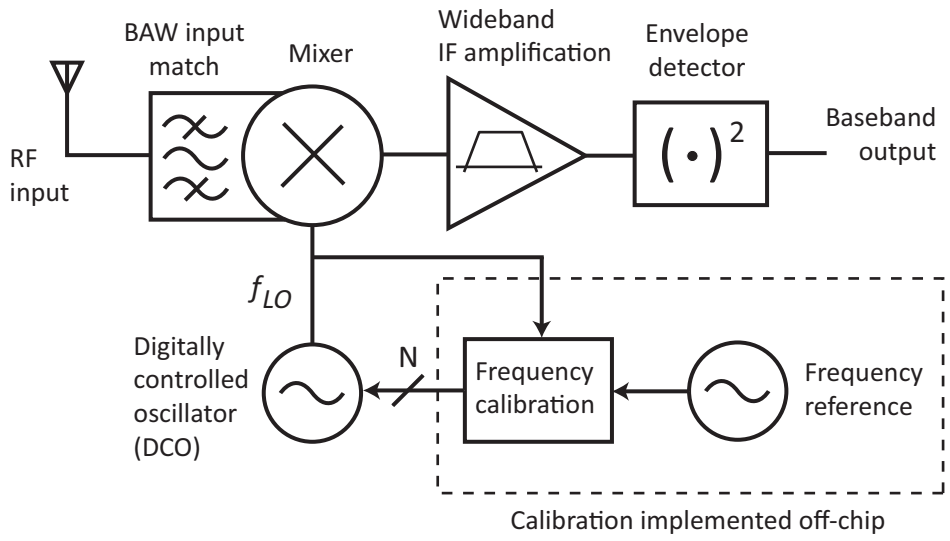

Figure 4.3: Block diagram of prototype uncertain-IF receiver

with a gain block covering the entire IF range and finally converted to DC by the envelope detector. On the LO side, a free-running digitally-controlled oscillator (DCO) drives the mixer. Digital frequency control is used to calibrate the LO within the desired frequency range only when necessary, instead of maintaining an analog control voltage during normal operation. This section describes the design of each component in detail. In implementing each receiver block, the primary goal of reducing power consumption motivates simplicity in the circuit design. To further reduce power, the entire receiver is optimized to operate from a single 0.5 V supply.

4.2.1 Input Matching Network

The uncertain-IF receiver requires a narrow RF filter similar to the TRF receiver presented in the last chapter. The BAW-based network with tapped-capacitor match is also a good fit for this receiver, especially since it was already well-characterized in Chapter 3 and shown experimentally to work effectively. The matching network topology is identical to the one from Chapter 3, except a 2 GHz resonator is used due to higher availability from the manufacturer. Capacitor C_1 is reduced to 550 fF for optimum matching with the slightly different resonator frequency. As before, the mixer input capacitance can be absorbed with the resonator capacitance, eliminating the need for a real input impedance at the mixer input. If the mixer input is designed to be sensitive to voltage, passive voltage gain from impedance transformation will also boost mixer conversion gain.

4.2.2 Dual-gate Mixer

The mixer design is driven by two goals: maximizing conversion gain and minimizing LO drive requirements. A single-ended dual-gate topology (Figure 4.4(a)) is chosen because the LO port is conveniently driven from a single-ended ring oscillator. A differential ring oscillator would require at least two times the power of the single-ended implementation. RF and LO feedthrough inherent to the single-balanced design are filtered by the load network and the IF amplifier stages before arriving at the envelope detector. The input RF

signal is coupled onto the gate of M_1 through the capacitive transformer in the matching network, while DC bias is applied to the gate by an on-chip 50 kΩ resistor (not shown). Devices M_1 and M_2 are sized with W/L of (10/0.1) µm/µm, with M_2 presenting only about 10 fF of capacitive load to the LO. Although the cascode device M_2 generally modulates the transconductance of M_1, the CMOS buffers drive the LO port with a rail-to-rail signal, effectively switching the RF transconductor M_1 on and off (Figure 4.4(b)). Therefore, the output signal current i_o at the IF frequency can be calculated by approximating the time-varying transconductance $g_m(t)$ as switching between g_{m0} and zero [43]:

$$i_o = g_m(t)v_i = g_{m0}p(t)v_i \tag{4.1}$$

where $p(t)$ is a pulse train with 50% duty-cycle (square wave). Using the Fourier series representation of $p(t)$:

$$p(t) = \frac{1}{2} + \frac{2}{\pi}\cos(\omega_{LO}t) - \frac{2}{3\pi}\cos(3\omega_{LO}t) + \cdots \tag{4.2}$$

the output current is:

$$i_o = g_{m0}v_i\left(\frac{1}{2} + \frac{2}{\pi}\cos(\omega_{LO}t) - \frac{2}{3\pi}\cos(3\omega_{LO}t) + \cdots\right) \tag{4.3}$$

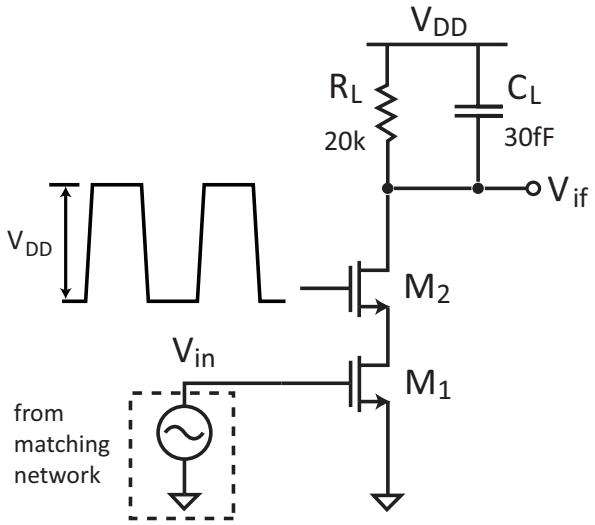

(a) Mixer schematic

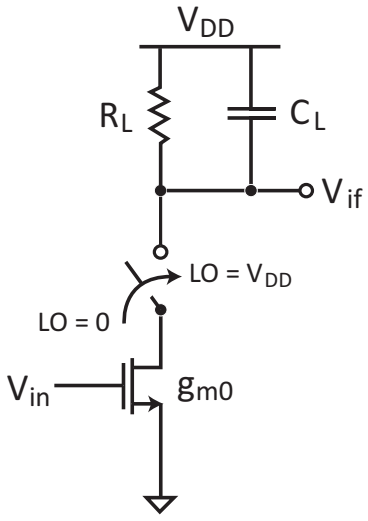

(b) Model of operation

Figure 4.4: Dual-gate mixer

The input RF signal is $v_i = v_s \cos(\omega_s t)$, leading to a final conversion transconductance g_{conv}:

$$\frac{i_o}{v_i} = \frac{g_{m0}}{2}\cos(\omega_s t) + \frac{g_{m0}}{2}\left(\frac{2}{\pi}\cos\left(\omega_{LO} \pm \omega_s\right)t\right) \cdots \rightarrow g_{conv} = \frac{1}{\pi}g_{m0} \qquad (4.4)$$

To obtain the overall voltage conversion gain G_{conv} from RF to IF, g_{conv} is multiplied by the output resistance of the mixer at the IF frequency:

$$G_{conv} = \frac{1}{\pi}g_{m0}\left(R_L \parallel R_{o,mix}\right) \qquad (4.5)$$

where R_L is the load resistance and $R_{o,mix}$ is the output resistance looking into the drain of M_2 when the LO voltage is at V_{DD}. The mixer load resistor R_L is made as large as possible to maximize the conversion gain within the available voltage headroom. The final resistor design value is 20 kΩ, implemented with a $p+$ polysilicon resistor. The quiescent transconductance g_{m0} is controlled by the DC bias voltage on the gate of M_1, which is set at 330 mV. Under these bias conditions and with the LO running, the simulated average current in the mixer is 13 µA. Including the voltage gain in the matching network and using Equation 4.5, the calculated G_{conv} is 13.9 dB, which closely matches the value of 14.5 dB obtained with SpectreRF periodic steady-state (pss) simulations.

A CAD layout of the mixer core devices is shown in Figure 4.5. The RF and LO devices are laid out as a single cascode compound device. The contacts can then be removed from the intermediate node and the poly gates moved closer

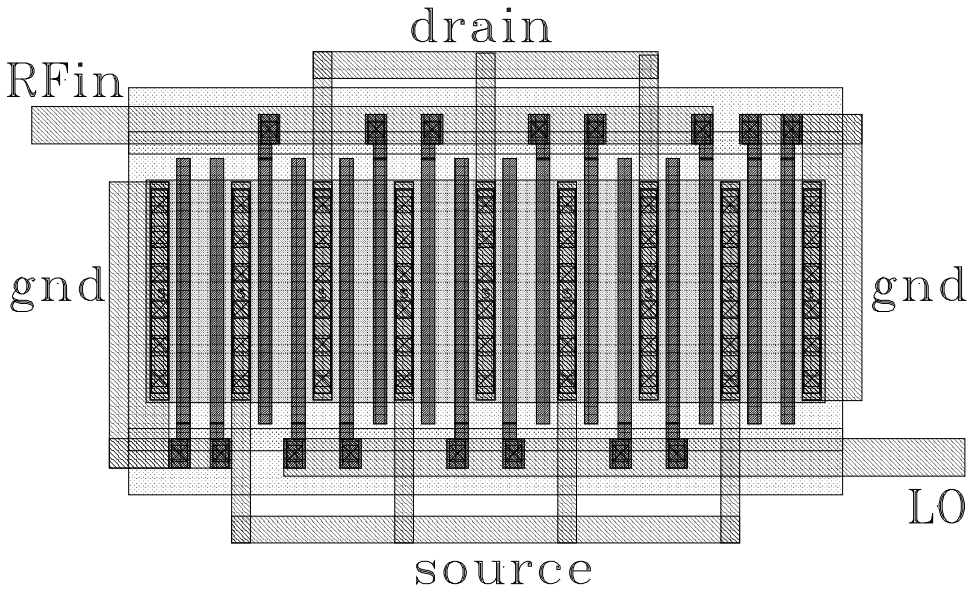

Figure 4.5: CAD layout of dualgate mixer core

together, reducing junction capacitance. This results in a compact layout with convenient connection points for the RF input port and LO drive. Due to the relatively small size of the compound device, two dummy devices are included on both ends of the structure with gates tied off to ground to reduce edge effects and improve matching with simulation models.

4.2.3 IF Amplifier

As specified in the architecture design, the IF amplifier must provide gain across the bandwidth of 100 MHz. In scaled CMOS technology, this frequency performance is easily met using a wideband differential pair with resistive loads. In order to operate under the low supply voltage a multi-stage architec-

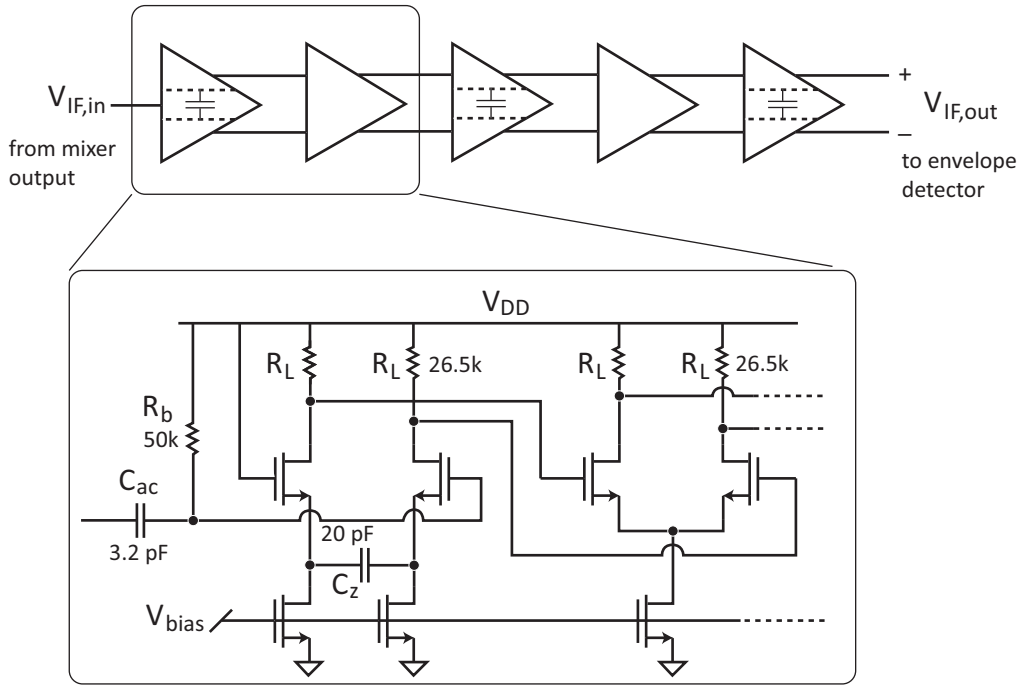

Figure 4.6: Schematic of IF amplifier

ture is chosen, using five differential pair gain stages optimized for maximum gain-bandwidth product for a given power consumption. Accordingly, each stage provides a gain of about 8 dB [19]. The input is AC coupled to the mixer output as shown in Figure 4.6. The differential pair devices are sized (6/0.2) µm/µm and biased in the subthreshold regime for high transconductance efficiency ($\frac{g_m}{I_d}$). The gain stages together produce more than 40 dB of total gain, with each stage consuming 8 µA of current. The use of identical stages and resistive loads simplifies biasing and allows simple DC coupling between stages (Figure 4.9). The bias currents of all five stages are matched and controlled simultaneously via a single voltage, which is shared among all

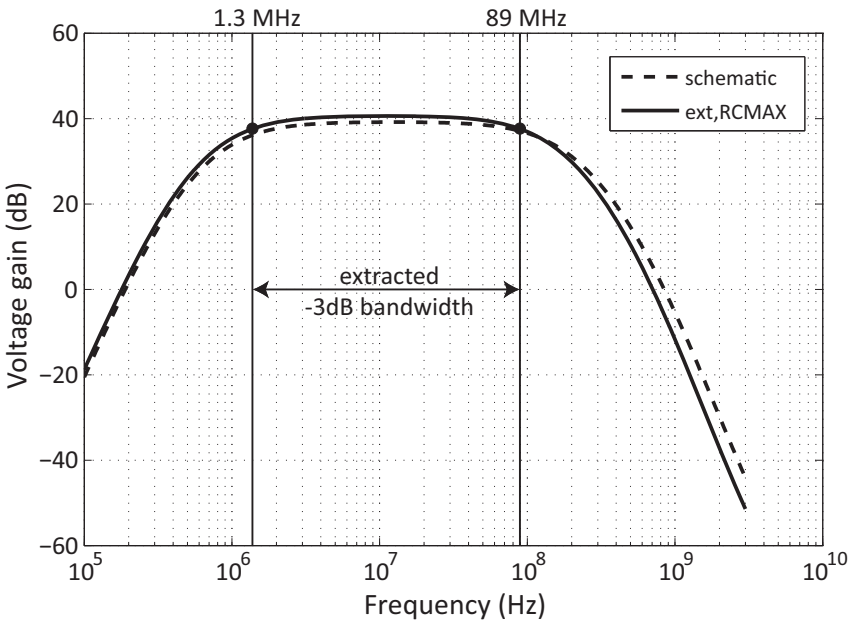

Figure 4.7: Simulated IF amplifier frequency response

stages.

In the first, third, and fifth stages, the tail current source is split into two halves with a coupling capacitor C_z [44] of 20 pF, introducing a zero at DC in the differential transfer function. Combined with AC coupling between the mixer and the first IF stage, this technique rolls off the IF gain close to DC, where the IF signal would be too close to the baseband bandwidth. The lack of gain at DC also prevents large accumulated offsets through the IF amplifier chain [16]. The simulated frequency response of the complete IF amplifier is shown in Figure 4.7 for both the plain simulation schematic and with parasitic extraction from the final layout. For the extraction, only capacitances

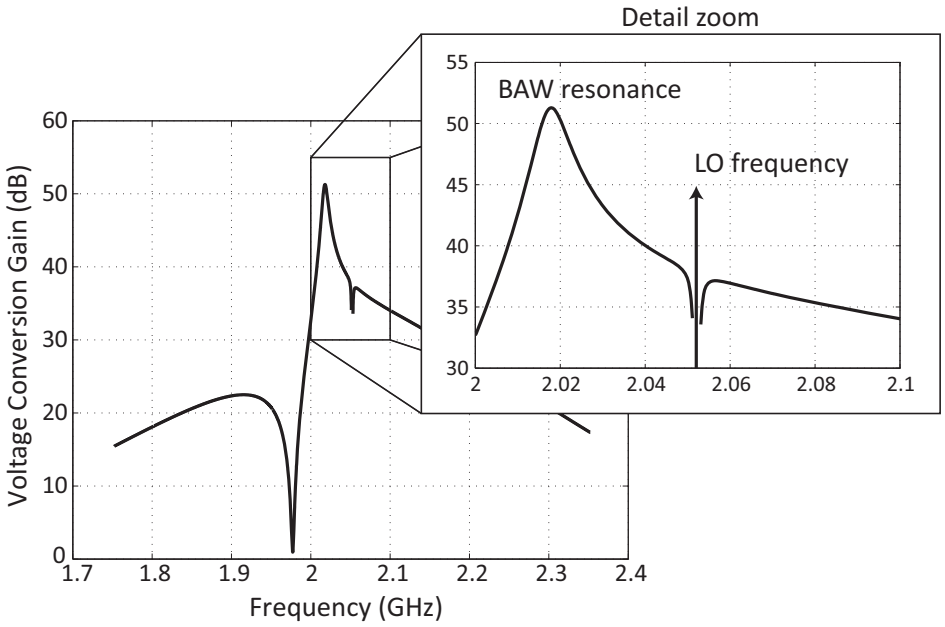

Figure 4.8: Simulated front-end conversion gain

to ground and coupling capacitances are considered[1]. The statistical models for maximum, worst-case capacitance were used to make sure that the amplifier would have adequate bandwidth under worst case conditions. The -3 dB bandwidth is marked in Figure 4.7, verifying that the amplifier has high gain across the band from 1 to 100 MHz, with a peak gain above 40 dB.

The simulated voltage conversion of gain of the combined mixer/IF amplifier front-end is about 50 dB to the IF output, with a corresponding noise figure of 23 dB. Figure 4.8 shows the simulated conversion gain from periodic transfer function (pxf) analysis. For this simulation, the LO frequency is set to about 2.05 GHz. As mentioned earlier, the roll-off of gain at DC causes a

[1]Distributed resistance was not extracted

null in the gain response at the LO frequency. The width of this dead band is determined by the high-pass cutoff frequency of the IF amplifer, due to C_z and AC coupling to the mixer. Without the null, if the LO does happen to fall directly on top of the desired channel frequency, the input signal would be converted directly to DC, bypassing the nonlinear function of the envelope detector and corrupting the baseband output. In order to avoid problems, the width of the null should be kept larger than the baseband bandwidth, but also as small as possible relative to the IF bandwidth. This minimizes the probability of the LO frequency aligning with the RF channel frequency. For this design, the allowed LO range is approximately 200 MHz and the dead band is less than 2 MHz. In the unlikely event that the LO falls in the wrong place, it can also be re-tuned. By designing the IF bandwidth about 10% larger than the LO calibration step size, there are guaranteed to be at least two calibration settings within the desired range. Therefore, the receiver could always flip to another LO frequency setting.

4.2.4 Differential Envelope Detector

The envelope detection circuit is implemented with a differential pair [16] biased in weak inversion with 1 µA of current per side for maximum nonlinearity. A simplified schematic of the complete receiver is shown in Figure 4.9, including the differential detector. When a differential IF signal drives the gates of M_3 and M_4, the nonlinear bias point shift appears at the drain of the tail current source, converting the IF energy to a DC baseband signal. In order to

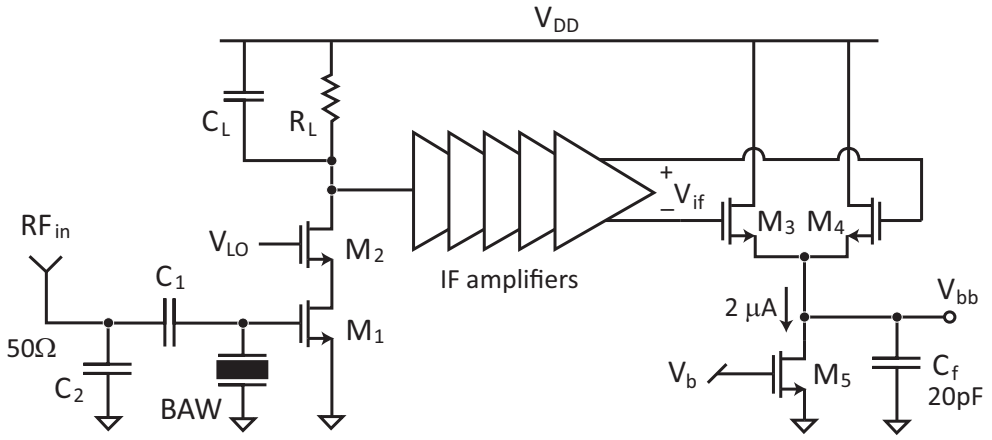

Figure 4.9: Simplified schematic of complete prototype receiver

avoid loading the IF amplifier excessively, the detector pair must not be sized too large. Devices M_3 and M_4 have an aspect ratio of (10/0.2) µm/µm, with current source device M_5 sized at (5/1) µm/µm. A 20 pF capacitor at the output filters any feedthrough from the IF signal or higher harmonics, with a baseband bandwidth of about 600 kHz.

For signals inside the detector's baseband bandwidth, the differential topology rejects the differential mode, but common mode signals pass through with gain $k_{DC} \approx 1$. The output noise is given by:

$$\overline{V_{o,n}^2} = 4kT\gamma \frac{1}{2g_{m3}} \left(1 + \frac{g_{m5}}{2g_{m3}}\right) \tag{4.6}$$

For the same total bias current, the differential detector has the same output noise as the single-ended version from Figure 3.2.

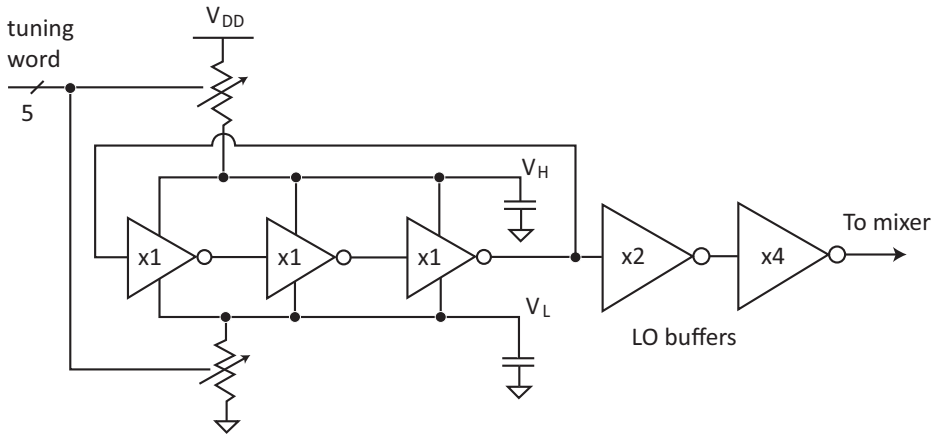

Figure 4.10: Digitally-controlled oscillator (DCO) schematic

4.2.5 Digitally-Controlled Oscillator

The DCO is implemented with the simplest type of ring oscillator, a 3-stage CMOS ring using standard library inverters. Frequency tuning is accomplished through the use of two identical resistive DACs that modify the virtual supply rails (V_H, V_L) of the ring (Figure 4.10). Two DACs are used in order to keep the voltage swing near the middle of the range, so that the output levels can be restored to full swing using an inverter chain operating with the full V_{DD}. The scaled inverter chain serves as a non-resonant buffer to drive the mixer LO port. Low threshold devices are used to ensure sufficient speed with the 0.5 V supply. The 5-bit resistive tuning DACs are simple switched resistor networks. The resistor values are designed using Monte Carlo simulations to guarantee that the LO frequency can always be tuned within the desired range across process and temperature. The frequency tuning step size, which

defines the calibration precision, is approximately 50 MHz. A variety of well-known techniques can be used to calibrate the DCO, which is similar to a coarse tuning algorithm in standard digital PLLs [45]. Re-calibration of the LO frequency is only required to adjust for process variations and changes in temperature and voltage that occur over time. Because calibration cycles will be relatively infrequent, the active power of the calibration circuitry can be amortized over the entire period between calibrations. The average power devoted to calibration is given by:

$$P_{avg} = P_{active}\frac{T_{cal}}{T_{interval}} \tag{4.7}$$

where P_{active} is the combined power of the calibration circuitry and frequency reference, T_{cal} is the time required to complete a calibration cycle, and $T_{interval}$ is the time between calibrations. As a worst case example, assume that P_{active}=1 mW and T_{cal}=1 millisecond. If calibration is performed once every 10 seconds, the average power is 1 µW. In reality, calibration will not be required so often, and the calibration time is likely to be shorter than 1 millisecond. In Section 4.3.1, measurement results show that, in the absence of large rapid temperature changes, calibration may only be required after several minutes or even hours. Nevertheless, even for these worst-case estimates, the calibration power is not a significant fraction of the total power budget.

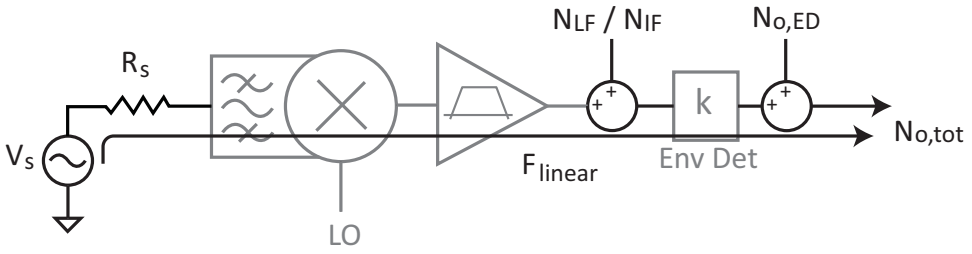

Figure 4.11: Noise sources for the uncertain-IF receiver

4.2.6 Complete Sensitivity Analysis

With the receiver design parameters established as described above, the overall sensitivity can now be predicted following the analysis method from Section 3.2.2. Figure 4.11 shows the sources of noise in the uncertain-IF receiver and the transfer functions to the output. The mixing front-end is a linear block and is modeled with its noise factor F_{linear}. The envelope detector noise is added at the output ($N_{o,ED}$), given by Equation 4.6. Noise from the IF amplifier is added at the input of the detector and originates from two different mechanisms. The first is analogous to the low frequency noise of the FEA for the receiver in Chapter 3. The differential detector used in this receiver rejects low frequency *differential* noise from the IF amplifier, but common mode low frequency noise (N_{LF}) must still be taken into account, as it will pass through to the detector output with $k_{DC} \approx 1$.

The other noise source at the IF amplifier output is unique to the architecture and arises due to the wide IF bandwidth. Since the high-Q filter occurs at the input of the receiver, the noise of the front-end entering the detector

is integrated across the entire IF bandwidth of 100 MHz. This noise source (N_{IF}) passes through the nonlinear transfer function of the detector with the desired signal. The noise density *at the detector output* due to N_{IF} can be calculated as [34]:

$$N_{o,IF} = \frac{\left(2\sigma^2\right)^2}{\left(4nV_t\right)^2}\frac{1}{BW_{if}} \tag{4.8}$$

where σ^2 is the noise variance at the IF output integrated across the entire IF bandwidth. The value of σ^2 is determined by periodic steady-state simulation with periodic noise analysis.

The output noise is added as an additional factor in Equation 3.9 to arrive at the complete noise factor for the uncertain-IF receiver:

$$F_{tot} = 2F_{linear} + \frac{N_{LF}k_{DC}^2}{N_{src}G_{conv}^2k^2} + \frac{N_{o,ED}}{N_{src}G_{conv}^2k^2} + \frac{N_{o,IF}}{N_{src}G_{conv}^2k^2} \tag{4.9}$$

where F_{linear} and G_{conv} are the linear noise figure and voltage conversion gain of the mixer/IF amplifer combined front-end. As before, final values for the noise densities in Figure 4.11 are derived from simulations and normalized over a brickwall detector bandwidth as illustrated in Figure 3.6.

The relative contributions to the noise factor for each term in Equation 4.9 are shown in Figure 4.12, using simulations of the final design to establish values for all noise and gain variables. The integrated IF noise ($N_{o,IF}$) dominates at the sensitivity limit due to the wide IF bandwidth. Reducing the IF bandwidth will proportionately reduce this noise component, at the expense of increased LO tuning accuracy and less tolerance to LO drift.

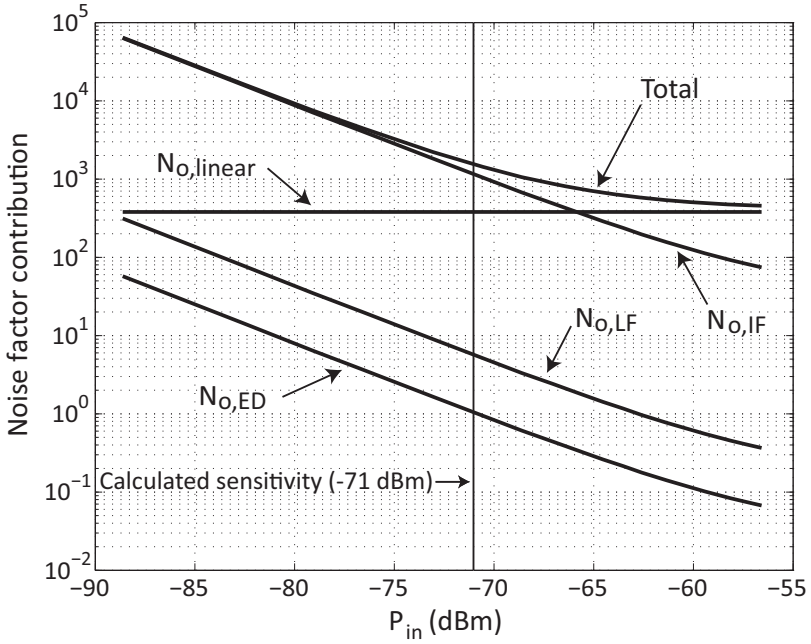

Figure 4.12: Breakdown of noise figure contributions

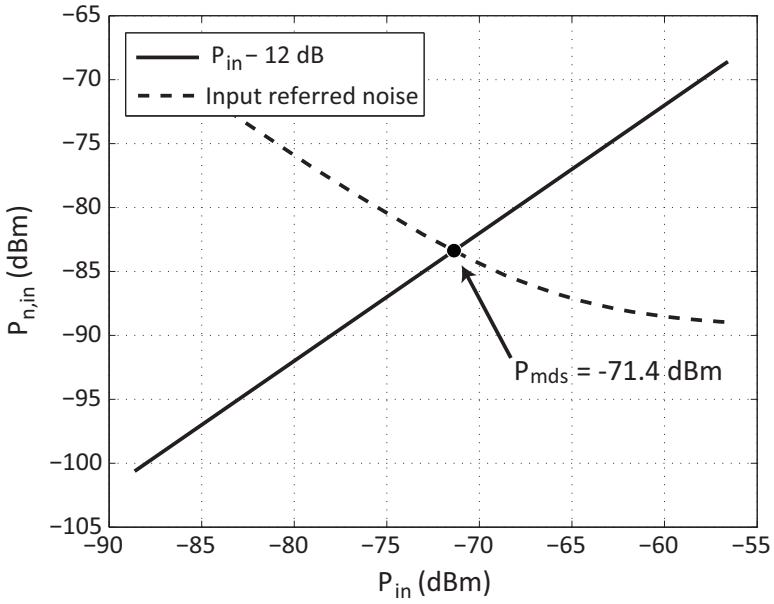

Figure 4.13: Calculated sensitivity for uncertain-IF receiver

The overall sensitivity is predicted by using Equation 4.9 to plot the input-referred noise versus the power of the RF input signal. Figure 4.13 shows that the minimum detectable signal (P_{mds}) to guarantee 12 dB baseband SNR is -71.4 dBm. The improved sensitivity of this receiver over the prototype in Chapter 3 is due to the higher gain of the frequency conversion front-end. The mixer and IF amplifier combination realizes more than 50 dB gain before the detector, compared to just 16 dB for the FEA in Chapter 3.

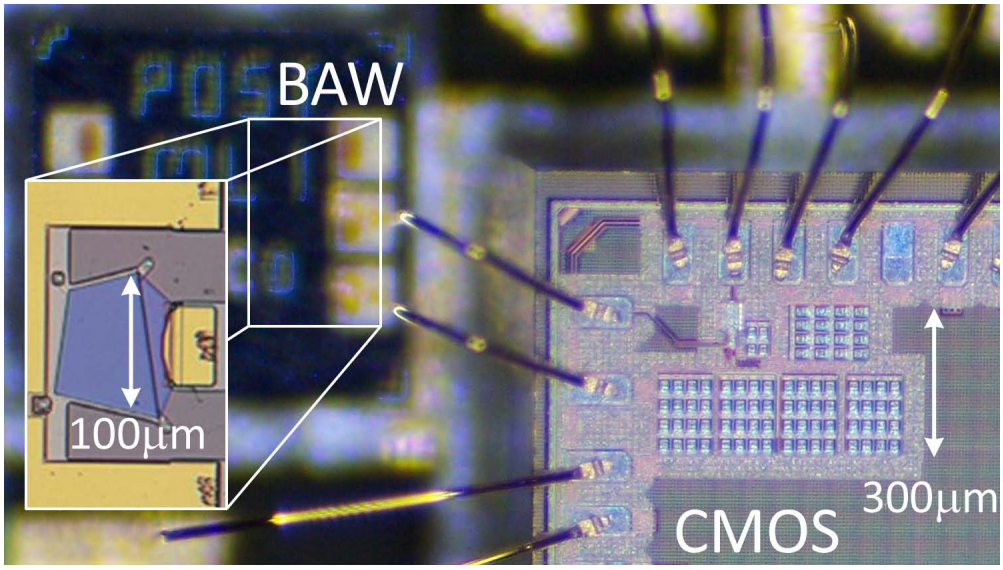

Figure 4.14: Die photo of receiver prototype bonded to packaged BAW

4.3 Measurement Results

The prototype receiver is fabricated in 90 nm standard CMOS technology with MIM capacitors (Figure 4.14). The active area is approximately 0.1 mm^2, again with no external components required except a single BAW resonator. The packaged resonator can be seen wirebonded to the die similar to the prototype in Chapter 3. On the CMOS side, Figure 4.15 shows a magnified view of the active die area. Due to the simple circuit design and lack of on-chip inductors, the silicon area devoted to active circuitry is extremely small. The majority of the area is taken by MIM capacitors for the source-coupled IF amplifier stages and power supply decoupling. For ease of layout, these capacitors were positioned adjacent to the circuits, but the area could

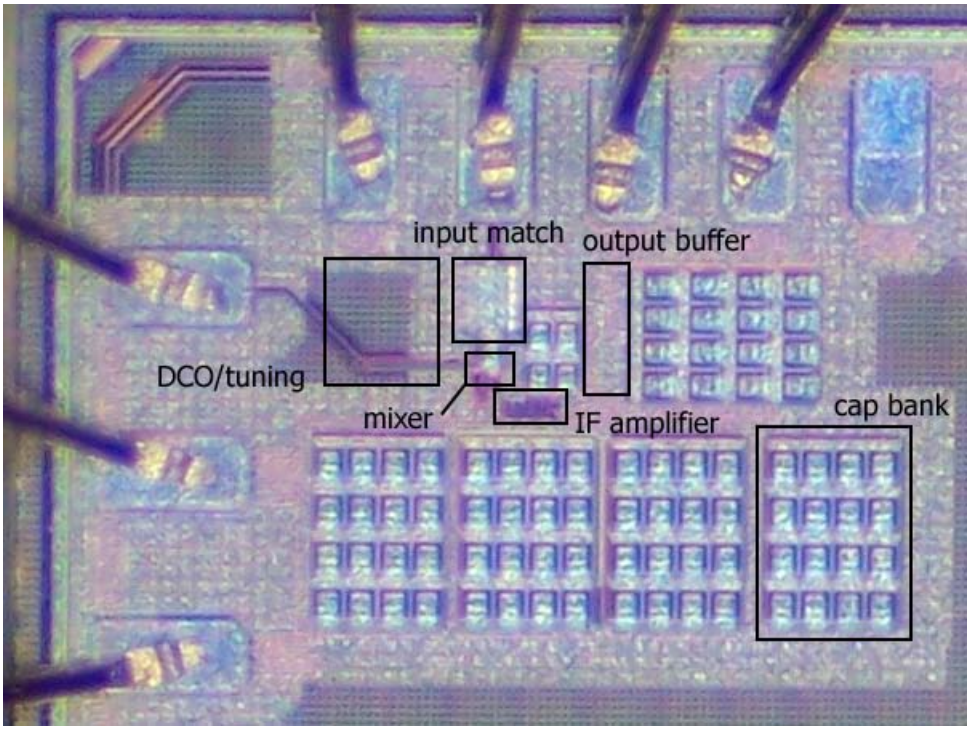

Figure 4.15: Annotated die photo

be reduced by moving the MIM capacitors above the receiver blocks. The complete WuRx fits conveniently in the corner of the chip, which is appropriate for its role as an auxiliary receiver.

4.3.1 Standalone LO Measurements

A standalone LO test block is included on the prototype chip for characterization purposes, consisting of a DCO and LO buffers identical to the circuits used in the receiver, along with an open-drain buffer to drive off-chip instrumentation. For receiver functionality, the chief metrics of interest for the LO are process compliance, temperature compliance, and transient stability. The first two factors are addressed through frequency calibration, while the latter determines how often calibration is required. To compensate for process variation, the measured tuning range of the LO is from approximately 1 to 3 GHz, with the tuning curves for five different samples plotted in Figure 4.16(a).

Three samples were also measured across a temperature range from 0 to 90℃ (Figure 4.16(b)), using an off-chip state machine to control the oscillator tuning. Chip temperature was swept using a Temptronic ThermoStream TP04100A thermal forcer setup. The frequency of the oscillator is allowed to drift with temperature until it leaves the preset limits, which triggers an automatic calibration cycle to re-center the LO. For all three samples, the frequency remains well within the desired region around 2 GHz across the entire temperature range.

To quantify long-term stability, the test oscillator frequency was measured

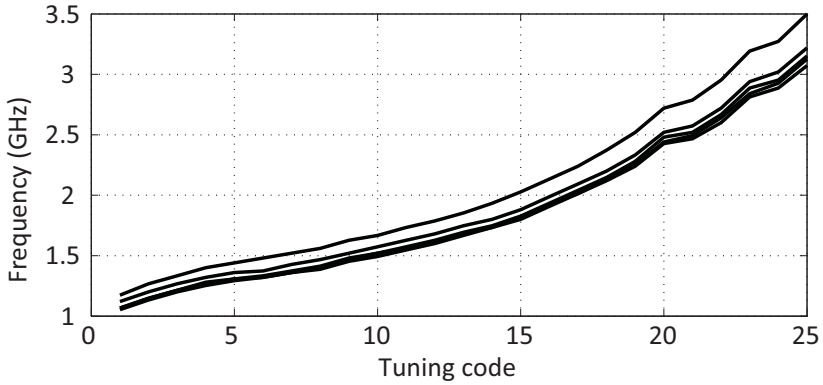

(a) Frequency tuning range

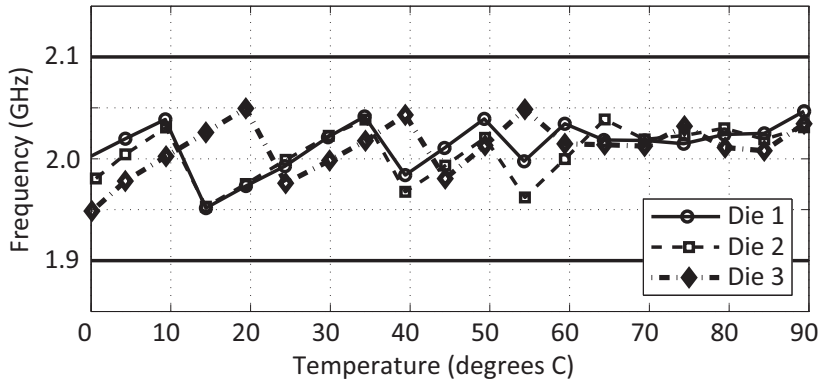

(b) Temperature compliance

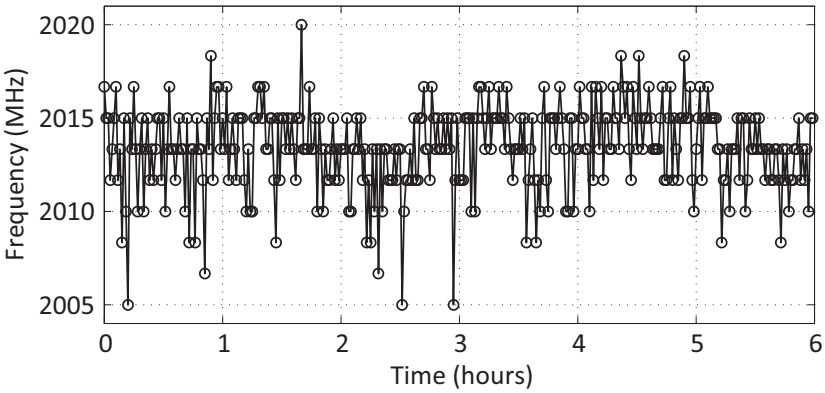

(c) Long-term stability

Figure 4.16: Measurements for standalone LO test block

open-loop over a six hour period at one minute intervals in an office environment, without changing the frequency control word. Figure 4.16(c) shows that the frequency ranges from 2005 MHz to 2020 MHz, verifying that LO calibration would not have been required during the entire six hour period. This robustness is a direct result of the wide 100 MHz IF bandwidth chosen for this implementation, illustrating that the receiver is able to remain functional despite variations in process and temperature. Similarly, long-term drift due to aging is not an issue as long as the reference used for calibration is stable over time. The tuning range available to cover process variation should be more than enough to handle any age-related frequency drift of the ring oscillator.

The receiver architecture is also robust to short-term LO frequency variation, or jitter. The measured time domain waveform of the standalone test LO is shown in Figure 4.17 from an Agilent Infiniium 54855A sampling oscilloscope. The asymmetry of the waveform shape is due to the open-drain buffer included on-chip to drive instrumentation. The period of the LO signal is 500 ps, with peak-to-peak jitter of about 70 ps. Despite the poor jitter performance, the receiver functions with no problems due to the architecture's inherent tolerance of variation in the LO frequency.

4.3.2 Receiver Gain Response

The receiver's overall RF-to-baseband gain response versus frequency is plotted in Figure 4.18 for four different samples. The response of the BAW resonator is evident in the plot, with the peak gain occurring at 2.02 GHz on the parallel

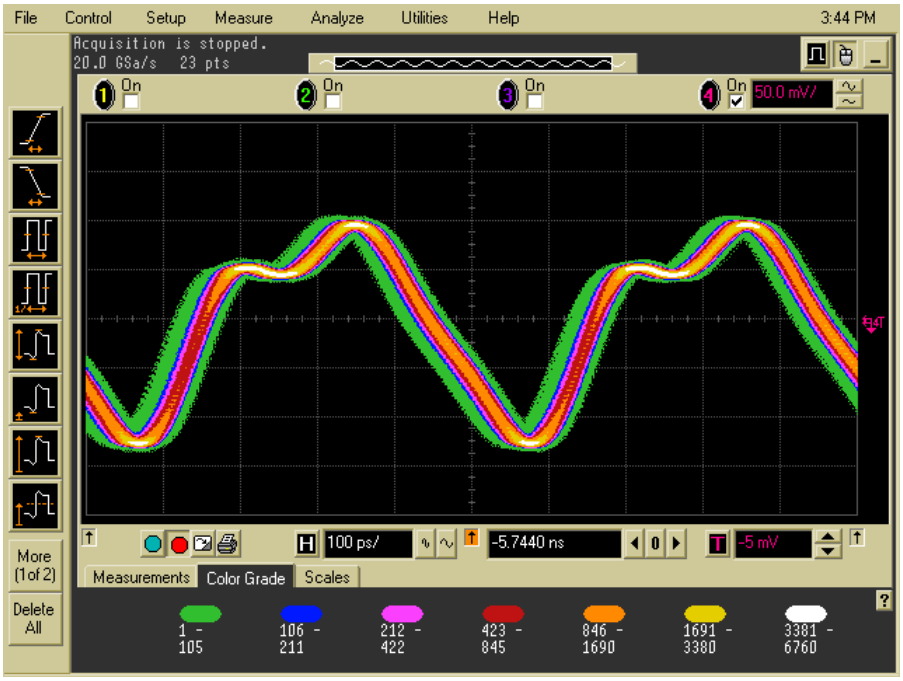

Figure 4.17: Measured transient waveform of LO test block

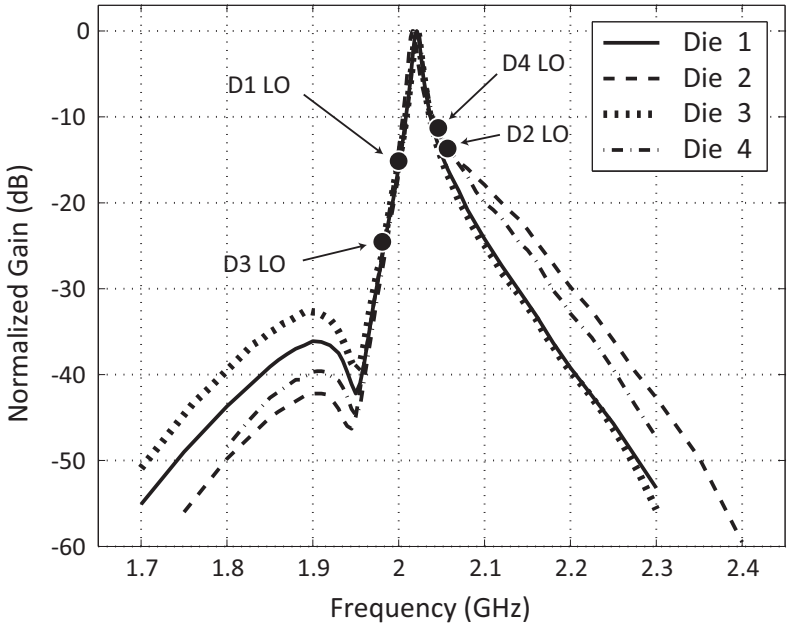

Figure 4.18: Normalized receiver gain to baseband for four different samples, with LO frequency marked for each sample

Table 4.1: LO frequencies after calibration for different samples

Sample	LO Frequency (GHz)	Offset from RF channel (MHz)
1	1.998	-22
2	2.060	+40
3	1.980	-40
4	2.044	+24

resonance of the BAW. The -3 dB bandwidth is about 9 MHz at the peak. Input $|S_{11}|$ (not shown) is about -10 dB. Although not as well-matched as expected, the quality of the match is sufficient for testing. For each sample, the LO was first calibrated to the frequency setting closest to the RF channel frequency. The LO frequencies after calibration are listed in Table 4.1 and marked on the gain plot in Figure 4.18, showing the natural variation in LO frequency for different samples. Each calibrated LO is well within the required range, and none are close enough to the channel to null the desired signal as described in Section 4.2.3. Although the resonator dominates the frequency response, samples 2 and 4 show a slight knee on the high frequency side of the peak. This is due to the IF gain, which extends 100 MHz on both sides of the LO frequency. This effect is masked by the resonator on samples 1 and 3.

4.3.3 Receiver Sensitivity and Robustness

The receiver sensitivity is measured by modulating the input RF carrier with an OOK pseudorandom bit sequence and buffering the baseband analog output signal off-chip, where the raw waveform is directly sliced by a comparator to generate digital bits for the bit error rate (BER) tester. An op-amp buffer is included on the chip, designed to provide 12 dB of gain while driving off-chip loads up to 20 pF. A simplified schematic of the buffer is shown in Figure 4.19. Because the buffer is for measurement purposes only, it operates from the 1 V pad ring supply voltage. The baseband output from the envelope detector is single-ended, so an external common-mode reference voltage is used to adjust the output DC level.

The complete BER measurement setup is shown schematically in Figure 4.20. An additional 20 dB of gain is provided by a commercial op-amp on the PCB before slicing. For a BER of 10^{-3} and a data rate of 100 kbps, the sensitivity is about -72 dBm (Figure 4.21). The measured value closely matches the sensitivity prediction from analysis in Figure 4.13. For higher data rates, the bandwidth of the envelope detector begins to limit the response, degrading sensitivity by about 2 dBm at 200 kbps.

The measured sensitivity exhibits about 1 dB variation among the four different samples. BER measurements are plotted in Figure 4.22 for the same four samples as measured in Figure 4.18. The peak gains of all four receivers are within 1.6 dB of each other. The variation could be due to a number of

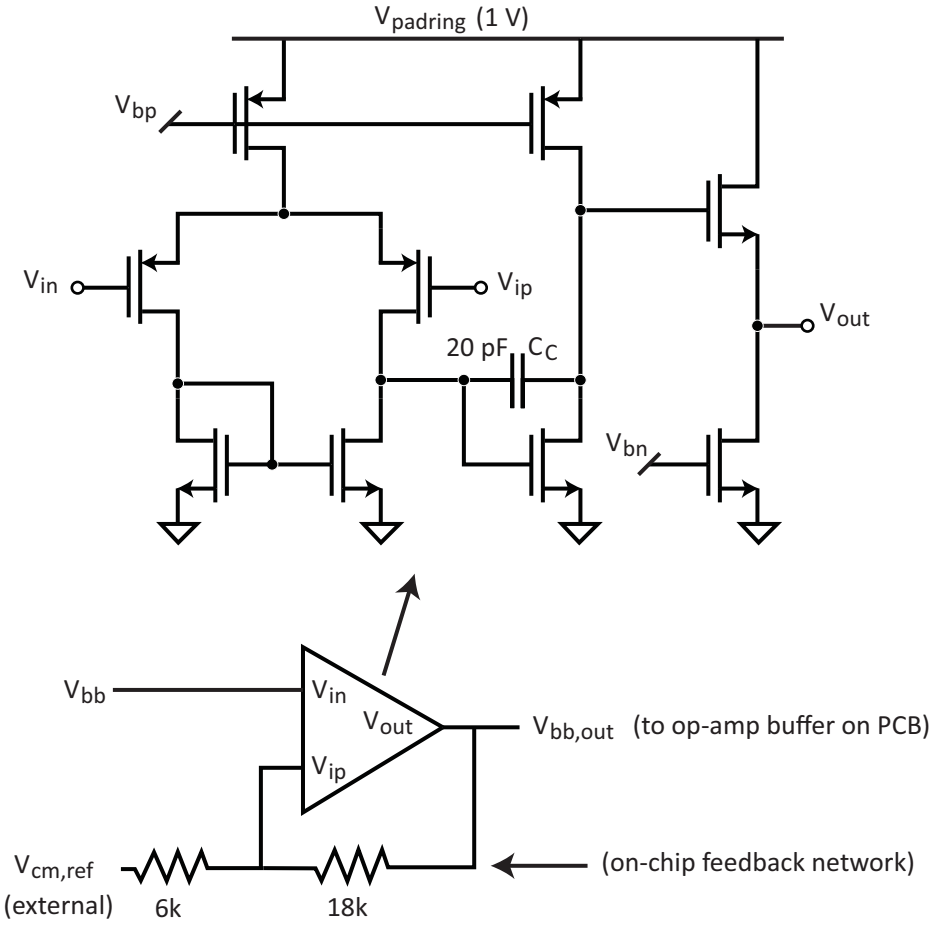

Figure 4.19: Simplified schematic of baseband buffer

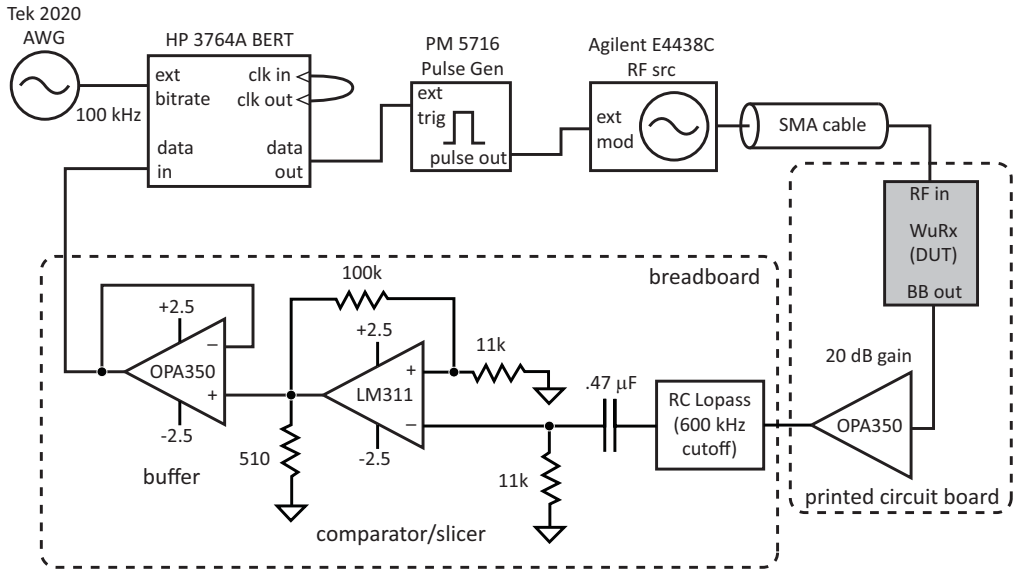

Figure 4.20: Test setup for BER measurements

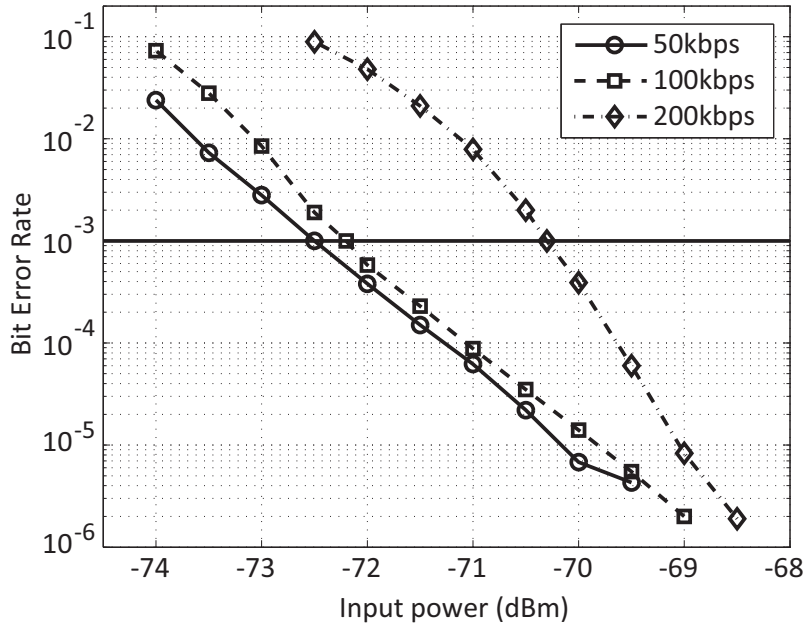

Figure 4.21: Measured BER versus input power for different data rates

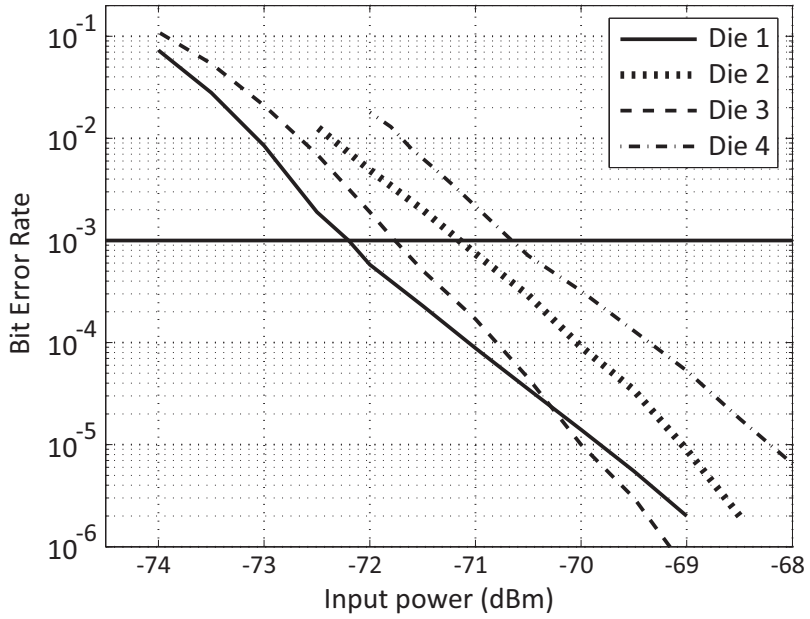

Figure 4.22: Measured BER versus input power for different samples

factors, including differences between individual BAW resonators and process variation of the CMOS transistors and polysilicon resistors. In any case, the relatively close matching of receiver sensitivity means that the highly variable LO frequency inherent to the uncertain-IF architecture does not have a large impact on sensitivity. It is worth mentioning that bit error measurements were conducted over several hours with no observed outages and no LO re-calibration required.

4.3.4 Selectivity and Interference Rejection

The receiver performance in the presence of interfering signals is a concern for the WuRx, because false alarms will needlessly activate the main receiver. The interferer performance is quantified using the following measurement setup. The desired signal is injected with a power level +3 dB above the sensitivity limit, combined with a continuous wave interferer at a given frequency offset from the desired channel. The interferer power level is increased until the BER rises above 10^{-3}, yielding a signal-to-interferer ratio (SIR) at that frequency offset, which represents the maximum interferer power level that can be tolerated without blocking the receiver. The results of the measurement are plotted in Figure 4.23, overlaid with the normalized gain response.

Clearly, the SIR points correspond closely to the gain response, indicating that the interferer performance is dominated by the front-end filter, rather than any nonlinear effects. As stated earlier, selectivity is now determined by the precision and high Q of the filter, instead of a highly precise frequency synthesizer. This is a characteristic common to all receivers based on envelope detection: any residual undesired signal after the filter is detected with the desired signal. For example, in an interference scenario where the undesired signal is within $\pm BW_{if}$ of the LO, the interfering signal itself will pass through the filter with finite attenuation before experiencing any nonlinear effects. The distortion products of these interferers will be negligible compared with the blockers themselves. A frequency domain picture of the situation is shown

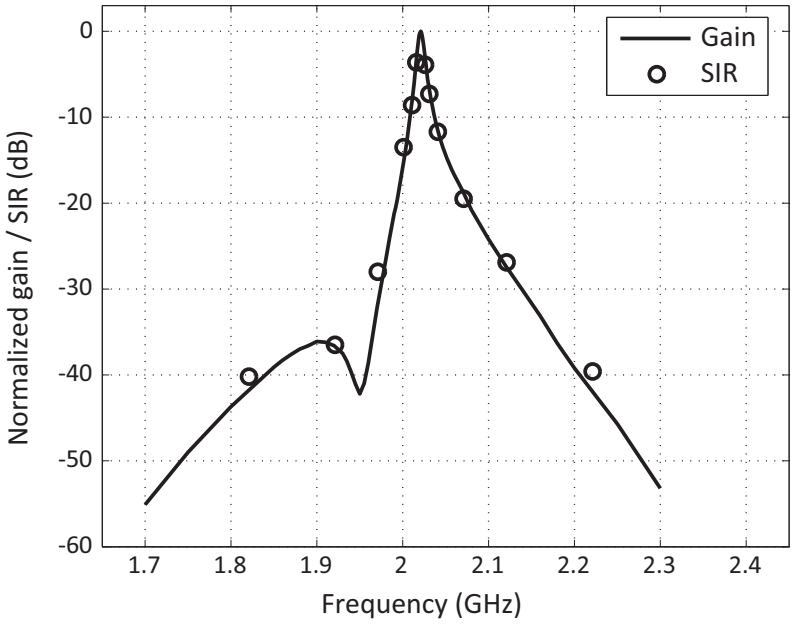

Figure 4.23: Acceptable signal-to-interferer ratio and normalized gain versus frequency

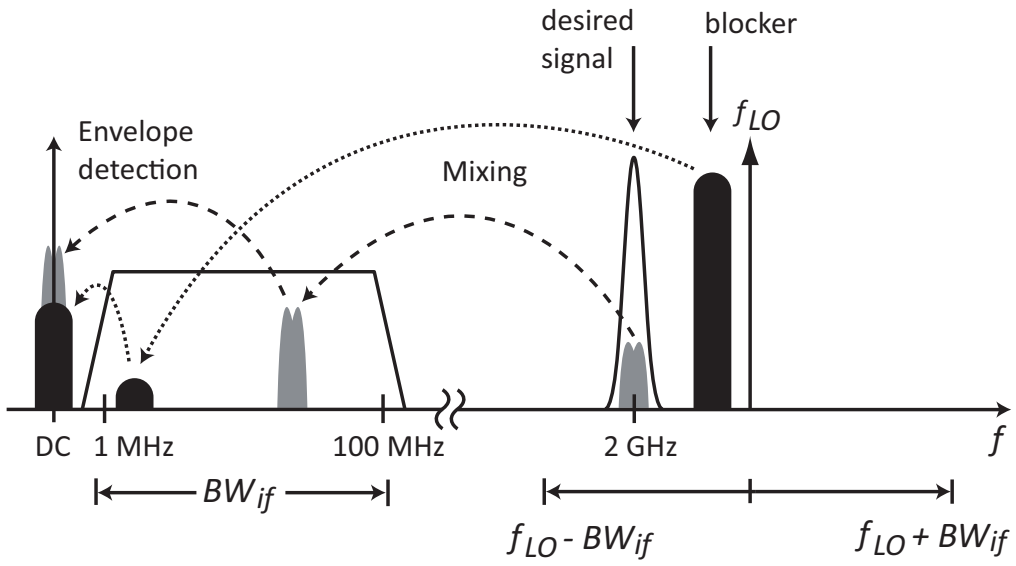

Figure 4.24: Blocker within IF bandwidth

in Figure 4.24. After filtering, a blocker will experience the linear transfer function and simply appear in the IF band along with the desired signal. This situation, where the blocker is outside the desired channel but inside the IF bandwidth, is the most likely scenario and also the worst case because the blocker will experience full IF gain. The only way to eliminate the blocker problem is improved filtering. Nevertheless, the fact that interference performance depends chiefly on the filter is a key observation. It means that very simple front-end circuitry may be used to reduce power consumption, despite poor linearity.

Although nonlinear effects are not a problem in a majority of cases, it is possible to construct a few special situations where front-end linearity is

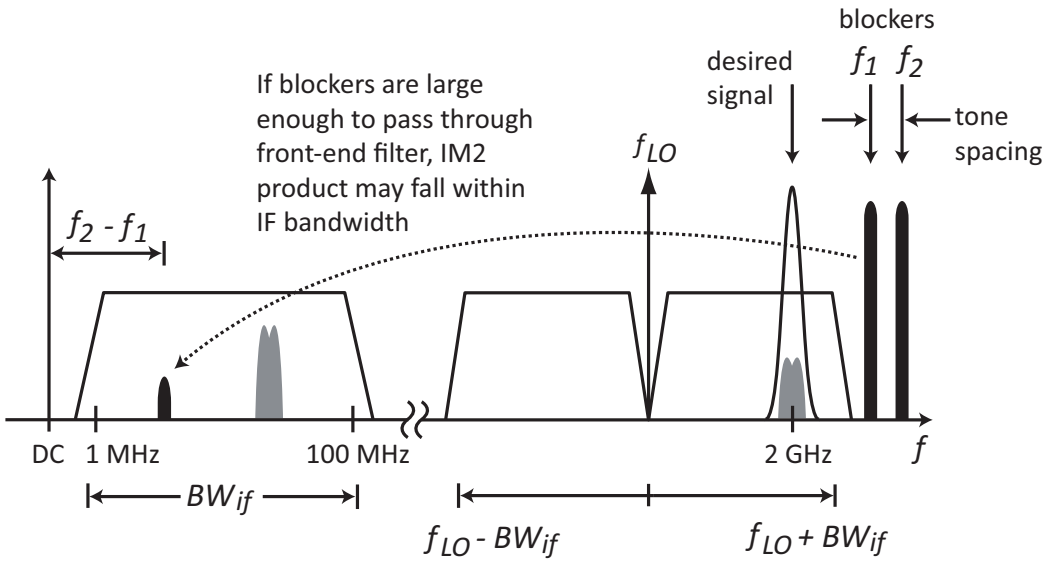

Figure 4.25: Two-tone blocker scenario

a concern. In contrast to the previous discussion on blockers inside the IF bandwidth, these cases occur when the undesired signals lie more than $\pm BW_{if}$ from the LO frequency. Figure 4.25 shows a two-tone blocker scenario where the blockers are outside the IF bandwidth from the LO, but close enough to the desired channel that they may not be completely filtered out by the resonator. Although the blockers themselves will not fall inside the IF band after mixing, the blockers will generate IM_2 products in the mixer due to the single-ended topology. Depending on the tone spacing, this distortion product could fall in the IF band and be detected by the envelope detector. This is the same problem that occurs with direct conversion receivers, which do not employ the wide IF bandwidth used in this receiver.

Reciprocal mixing with a large interferer can also occur if the blocker is

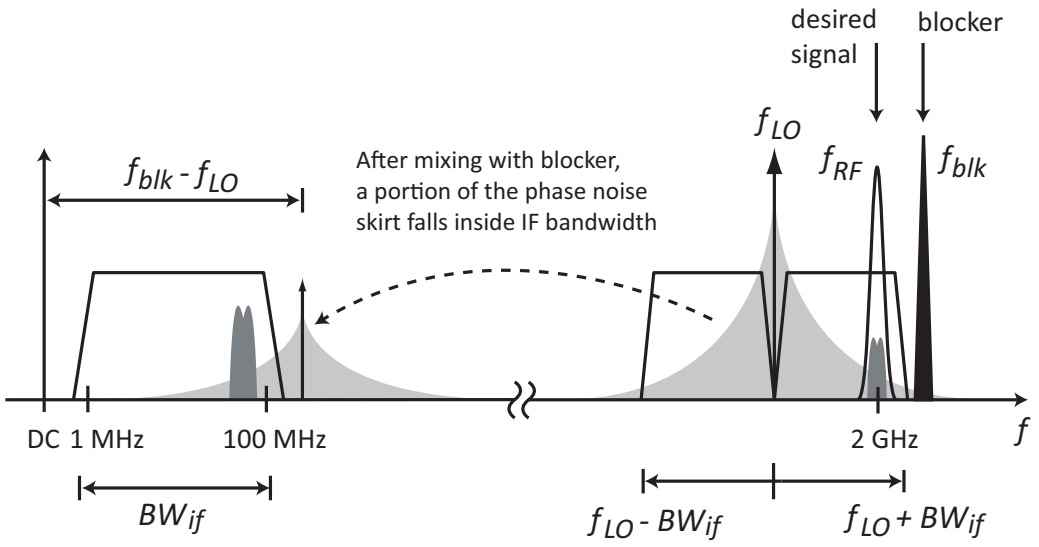

Figure 4.26: Reciprocal mixing

close to the channel frequency and the phase noise skirts of the LO extend to the blocker offset. Figure 4.26 illustrates this situation in the frequency domain. If the blocking signal is very large and not sufficiently attenuated by the front-end filter, it will mix with the LO phase noise skirts and some portion will fall into the IF band.

Unfortunately, lack of robustness to interferers is the price paid for low power consumption in this architecture. However, the preceding special cases are expected to be transient phenomena and statistically unlikely. If problems arise, one way to combat interferers is to overdesign the LO calibration steps so that it is always possible to re-tune the LO to another frequency. Not surprisingly, all of the preceding issues with selectivity and robustness are improved by a sharper *front-end* filter with improved selectivity of the RF

channel. The single resonator used in this implementation can only provide about 20 dB of out-of-band attenuation before the mixer. A better filter implementation using higher Q resonators or a filter structure with multiple resonators could provide better robustness to interfering signals.

4.3.5 Performance Summary

The total power consumption of the receiver is 52 µW from the 0.5 V supply. The LO generation and IF amplifiers draw about 80% of the total, with 20 µW and 22 µW, respectively. The mixer consumes 8 µW, while the envelope detector accounts for the remaining 2 µW. The measured power of the ring oscillator alone is 6 µW at 2 GHz. Although this figure increases to 20 µW when the LO buffers are included, this total still represents an order-of-magnitude reduction compared with the integrated LC oscillator in [13] and more than a factor of 4 improvement over the BAW-based oscillator in [46]. Table 4.3.5 compares the ring oscillator LO used in this receiver to other published LO generation circuits. Of course, the other oscillators listed in the table deliver much lower phase noise and better frequency stability than the ring oscillator. The key innovation of this receiver is the use of an architecture that is tolerant to high LO phase noise and frequency variability, while capitalizing on the drastically reduced power consumption. Otherwise, Table 4.3.5 clearly illustrates that LO generation alone would consume the entire WuRx power budget.

The complete receiver performance is summarized in Table 4.3 and the power consumption breakdown is shown graphically in Figure 4.27. In

Table 4.2: Comparison with published LO generation circuits

Reference	Type	Frequency (GHz)	Power (µW)
This work	CMOS ring	2	20[a]
[46]	BAW	1.9	89
[13]	integrated LC	2.4	160[b]
[47]	LC and PLL	2.6	2000[c]
[48]	bondwire LC	1.9	100

[a]Includes LO buffers
[b]Single-phase operation
[c]VCO alone

Table 4.3: Uncertain-IF receiver performance summary

Parameter	Measurement
Global supply voltage (V)	0.5
Carrier frequency/modulation	2 GHz / OOK
Power dissipation (µW)	
Mixer	8
IF amplifiers	22
LO + buffers	20
Envelope detector	2
Total	52
Date rate (kbps)	100/200 (nom/max)
Energy per received bit (nJ)	< 0.5
Raw sensitivity for 10^{-3} BER (dBm)	-72/-70 (100/200 kbps)

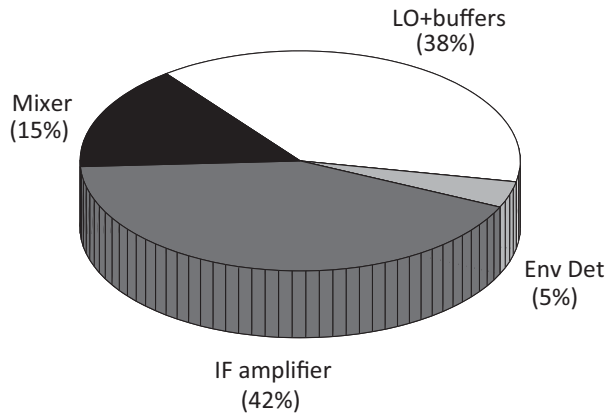

Figure 4.27: Uncertain-IF receiver power breakdown

contrast to the first prototype from Chapter 3, no single block dominates the receiver power consumption. The mixing function allows the signal gain to be realized at a lower frequency, significantly reducing the portion of system power devoted to amplification. Because the uncertain-IF receiver is still fundamentally based on envelope detection, the same coding and digital baseband techniques from Chapter 3 can also be applied to this receiver.

It is also important to emphasize that while frequency calibration is still required, this uncertain-IF architecture also guarantees a high tolerance to reference frequency inaccuracy. For example, the 100 MHz IF bandwidth chosen for this implementation corresponds to approximately 5% of the 2 GHz carrier frequency. A frequency reference that guarantees 2.5% accuracy over process and temperature variations is sufficient. This requirement is over 100 times less stringent than the performance of typical communication-grade quartz crystals, and can be obtained with a fully integrated *LC* or CMOS oscilla-

tor [49, 50], potentially reducing cost and increasing integration when compared to more conventional solutions.

4.4 Conclusion

This chapter presented a receiver implementation using an uncertain-IF architecture, designed specifically for the ultra-low power wake-up application. The significant power reduction is made possible through the combination of a CMOS ring LO and RF-MEMS resonator technology, breaking the power floor that arises using traditional high performance oscillators. Compared to the first prototype in Chapter 3, this receiver improves sensitivity by 200x (23 dBm) while maintaining similar power dissipation and data rate. Although communication efficiency[2] is not a goal for the wake-up receiver design, the energy efficiency performance of this prototype is still very good. The receiver achieves 500 nJ/bit at the nominal 100 kbps data rate and down to 250 pJ/bit at 200 kbps with slightly degraded sensitivity. This energy per bit figure is only half that of the receiver in [13] and a factor of 5 lower than the UWB receiver published in [51], which does not include synchronization energy.

[2]Low energy per bit

Chapter 5

Conclusion

This thesis presented a comprehensive investigation of receiver design for the wake-up application and explored the limits of ultra-low power receiver design. This chapter summarizes the work and puts the results in perspective by comparing to previously published work in the area of wireless receivers for WSN. Finally, the work concludes with suggestions for future research on the subject of wake-up receivers.

5.1 Performance Comparison

Figure 5.1 compares the sensitivity and power consumption performance of the receiver prototypes presented in Chapters 3 and 4 with previously published work. The uncertain-IF receiver achieves a final power consumption of 52 µW, which is about an order-of-magnitude below all previously published

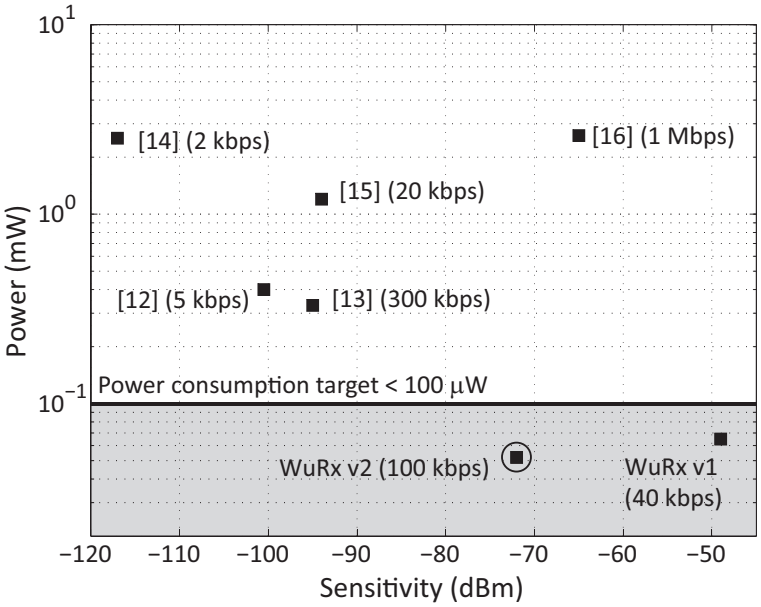

Figure 5.1: Comparison of WuRx with previously published receivers

receivers for WSN. The substantial power reduction is made possible through the combination of two technology factors:

1. **Scaled CMOS:** The high speed and low energy of modern CMOS transistors makes it possible to run a ring oscillator at 2 GHz with very little power.

2. **MEMS technology:** High Q micromechanical resonators allow the use of a ring oscillator in spite of poor frequency stability.

Even with the ultra-low power dissipation, the uncertain-IF receiver maintains sufficient sensitivity for 10 meter communication range and meets the functional specifications outlined in Chapter 1.

5.2 Future Research Directions

5.2.1 Improving the RF Front-end

First, it is important to note that the design choices made for this receiver were strongly dictated by the very low power budget and were optimized for the minimum possible power consumption. A wide IF bandwidth was selected to minimize required LO tuning accuracy and lower complexity. The resulting design has high tolerance to LO variability and re-calibration is rarely required. However, different design choices could be optimal in other network environments. For instance, the interference scenario described in Section 4.3.4, where second-order intermodulation distortion in the mixer corrupts the IF signal, could be avoided through the use of a differential LO and balanced mixer design. A differential LO path was avoided for this design due to the power penalty, but may be a viable alternative for other applications where the power specification is less stringent. In addition, the choice of carrier frequency for these designs (and the PicoRadio network in general) was restricted by the availability of BAW resonators, most notably lacking in the 2.4 GHz band. Both receivers developed here could easily be re-designed for 2.4 GHz operation if resonators were available.

Finally, one area for future improvement is in the differential envelope detector. The detector structure shown in Section 4.2.4 was selected as a simple extension of the single-ended detector in Section 3.2.1 and its convenient coupling to the IF amplifier. For the same input amplitude V_{if}, however, the

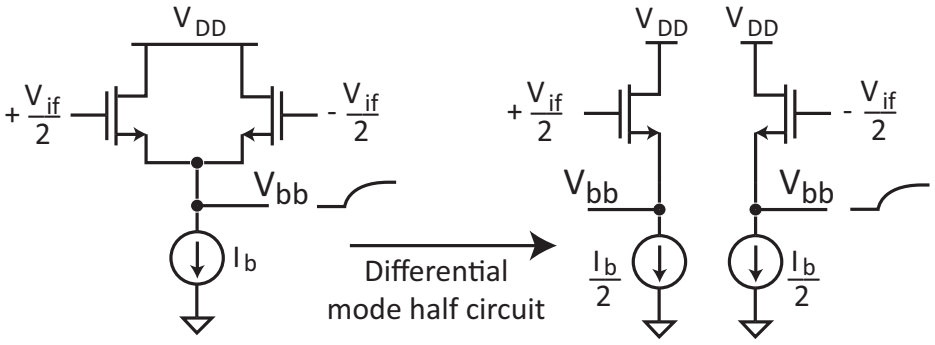

Figure 5.2: Differential half-circuit representation of detector

differential-pair envelope detector has lower conversion gain than the single-ended version. This can be seen by considering the differential half circuit shown in Figure 5.2. Each side of the differential pair operates with an input amplitude of $V_{if}/2$ instead of the full V_{if}. On top of the reduction in input signal, the detector gain itself is proportional to the input amplitude. The final result is that the differential pair detector has a conversion gain 4 times lower than the single-ended version.

A better differential detector structure is shown in Figure 5.3, which is basically a self-driven passive mixer. The schematic is drawn without biasing for clarity, but the IF signal could be AC coupled to the NFET gates to allow DC biasing of the transistors around the threshold voltage. The advantage of this structure is that it combines the outputs of both halves of the differential input constructively at the output. Preliminary simulations indicate that the detector shown in Figure 5.3 can improve the overall receiver sensitivity by about 4 dBm.

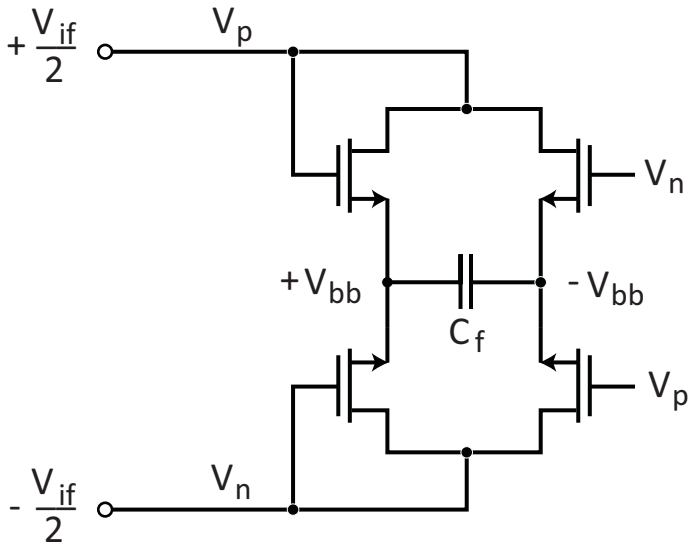

Figure 5.3: Passive mixer-based detector, biasing not shown

5.2.2 MEMS-based Front-ends

The uncertain-IF receiver saves power by shifting the burden of selectivity from an accurate local oscillator to the front-end filter. MEMS-based radio architectures continue to gain popularity and advances in resonator technology can open up new opportunities in circuit design. The uncertain-IF receiver is a good example of such a MEMS-enabled architecture. On the resonator research side, one recently published filter uses two coupled electrostatic resonators fabricated in bulk silicon to implement a complete channel selection filter at 425 MHz [52]. The filter achieves a narrow 1 MHz passband at the center frequency with almost 50 dB of stop band rejection, which is a significant improvement over the single BAW resonator filter used in this research.

This is just one example, but continued improvements in MEMS resonator technology may eventually enable RF channel selection in the front-end filter. This research shows that performing channel selection with high quality RF filters allows significant simplification of receiver circuitry and substantial power savings.

5.2.3 Wake-up Receiver Applications

This thesis focused on the wake-up receiver at the circuit level in order to meet the power consumption target. The next step is a more thorough exploration of potential applications for the wake-up receiver. Transceiver architectures that offer high efficiency communication but suffer from long synchronization time are good candidates for wake-up based synchronization. UWB transceiver architectures are becoming popular in WSN due to their low energy per bit requirements, but they may consume significant energy during acquisition as the receiver synchronizes to the incoming pulses. Instead, the transceiver could first use a low power wake-up receiver for coarse synchronization and then activate the UWB transceiver for efficient data transfer.

In addition to peer-to-peer sensor network applications discussed in Chapter 1, another interesting application is in active RFID tags. Unlike passive RFID, which requires a high power reader to provide power to the tags, an active tag contains a small power source and some active circuitry to extend the range or provide extra functionality. Active tags would have less longevity than their passive counterparts, but the increased functionality can enable new

applications. One possible example is a remote sensor that can be queried, not by a dedicated reader, but instead with a general-purpose wireless device like a mobile phone. For this application, the tag should be able to listen for queries and respond quickly, while still maintaining sufficient battery lifetime. Such a system is an excellent candidate for a low power listening receiver like the WuRx. With continued progress in device technology and low power wireless communication, additional applications are certain to emerge.

Bibliography

[1] N. Ota, S. Ahrens, A. Redfern, P. Wright, and X. Yang, "An application-driven architecture for residential energy management with wireless sensor networks," in *IEEE International Conference on Mobile Adhoc and Sensor Systems*, Oct. 2006, pp. 639–644.

[2] J. Rabaey, J. Ammer, T. Karalar, S. Li, B. Otis, M. Sheets, and T. Tuan, "PicoRadios for wireless sensor networks: the next challenge in ultra-low power design," in *IEEE ISSCC Digest of Technical Papers*, Feb. 2001, pp. 200–201.

[3] D. Steingart, S. Roundy, P. Wright, and J. W. Evans, "Micropower materials development for wireless sensor networks," *Materials Research Society (MRS) Bulletin*, vol. 33, no. 4, pp. 408–409, Apr. 2008.

[4] D. Steingart, C. C. Ho, J. Salminen, J. W. Evans, and P. Wright, "Dispenser printing of solid polymer-ionic liquid electrolyte for lithium ion cells," in *IEEE International Conference on Polymers and Adhesives in*

Microelectronics and Photonics (Polytronics 2007), Jan. 2007, pp. 261–264.

[5] B. Cook, "Low energy RF transceiver design," Ph.D. dissertation, University of California, Berkeley, 2007.

[6] E.-Y. Lin, J. Rabaey, and A. Wolisz, "Power-efficient rendez-vous schemes for dense wireless sensor networks," *IEEE International Conference on Communications*, vol. 7, pp. 3769–3776, June 2004.

[7] E.-Y. Lin, "A comprehensive study of power-efficient rendezvous schemes wireless sensor networks," Ph.D. dissertation, University of California, Berkeley, 2005.

[8] E. Shih, P. Bahl, and M. Sinclair, "Wake on wireless: an event driven energy saving strategy for battery operated devices," in *Proceedings of the 8th Annual International Conference on Mobile Computing and Networking*, Sept. 2002, pp. 160–171.

[9] L. Gu and J. A. Stankovic, "Radio-triggered wake-up capability for sensor networks," in *Proceedings of the 10th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS 2004)*, May 2004, pp. 27–36.

[10] B. P. Otis, "Ultra-low power wireless technologies for sensor networks," Ph.D. dissertation, University of California, Berkeley, 2005.

[11] Y. H. Chee, "Ultra low power transmitters for wireless sensor networks," Ph.D. dissertation, University of California, Berkeley, 2006.

[12] B. Otis, Y. Chee, and J. Rabaey, "A 400µW-RX, 1.6mW-TX super-regenerative transceiver for wireless sensor networks," in *IEEE ISSCC Digest of Technical Papers*, Feb. 2005, pp. 396–397.

[13] B. Cook, A. Berny, S. Lanzisera, A. Molnar, and K. Pister, "Low-power 2.4-GHz transceiver with passive RX front-end and 400-mV supply," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2757–2766, Dec. 2006.

[14] V. Peiris, C. Arm, S. Bories, S. Cserveny, F. Giroud, P. Graber, S. Gyger, E. L. Roux, P.-D. Pfister, N. Raemy, A. Ribordy, P.-F. Ruedi, D. Ruffieux, L. Sumanen, S. Todeschini, and P. Volet, "A 1v 433/868mhz 25kb/s-fsk 2kb/s-ook RF transceiver SoC in standard digital 0.18µm CMOS," in *IEEE ISSCC Digest of Technical Papers*, Feb. 2005, pp. 258–259.

[15] A. Molnar, B. Lu, S. Lanzisera, B. Cook, and K. Pister, "An ultra-low power 900 MHz RF transceiver for wireless sensor networks," in *Proc. IEEE Custom Integrated Circuits Conf.*, Sept. 2004, pp. 401–404.

[16] D. Daly and A. Chandrakasan, "An energy-efficient OOK transceiver for wireless sensor networks," *IEEE J. Solid-State Circuits*, vol. 42, no. 5, pp. 1003–1011, May 2007.

[17] J.-P. Curty, N. Joehl, C. Dehollain, and M. Declercq, "Remotely powered addressable UHF RFID integrated system," *IEEE J. Solid-State Circuits*, vol. 40, no. 11, pp. 2193–2202, Nov. 2005.

[18] B. Otis, Y. Chee, R. Lu, N. Pletcher, and J. Rabaey, "An ultra-low power mems-based two-channel tranceiver for wireless sensor networks," in *Proceedings VLSI Circuits Symposium*, Honolulu, June 2004, pp. 20–23.

[19] T. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*, 2nd ed. Cambridge, UK: Cambridge University Press, 2004.

[20] J. Borremans, P. Wambacq, M. Kuijk, G. Carchon, and S. Decoutere, "A 400µW 4.7-to-6.4GHz VCO under an above-IC inductor in 45nm CMOS," in *IEEE ISSCC Digest of Technical Papers*, Feb. 2008, pp. 536–537.

[21] S.-S. Li, Y.-W. Lin, Y. Xie, Z. Ren, and C. T.-C. Nguyen, "Micromechanical "hollow-disk" ring resonators," in *Proceedings of the IEEE Micro Electro Mechanical Systems Conference*, Jan. 2004, pp. 821–824.

[22] D. Weinstein and S. Bhave, "Internal dielectric transduction of a 4.5 GHz silicon bar resonator," in *Proceedings of the IEEE International Electron Devices Meeting (IEDM)*, Dec. 2007, pp. 415–418.

[23] R. Ruby, P. Bradley, J. Larson, and Y. O. D. Figueredo, "Ultra-miniature high-Q filters and duplexers using FBAR technology," in *IEEE ISSCC Digest of Technical Papers*, Feb. 2001, pp. 120–121.

[24] W. Pang, R. Ruby, P. Fisher, J. Larson, K. Grannen, D. Lee, C. Feng, and L. Callaghan, "A thermally stable CMOS oscillator using temperature compensated FBAR," in *Proceedings of the IEEE Ultrasonics Symposium*, Oct. 2007, pp. 1041–1044.

[25] B. P. Otis, "The design and implementation of an ultra low power RF oscillator using micromachined resonators," Master's thesis, University of California, Berkeley, 2002.

[26] S. Rai and B. Otis, "A 600 µW BAW-tuned quadrature VCO using source degenerated coupling," *IEEE J. Solid-State Circuits*, vol. 43, no. 1, pp. 300–305, Jan. 2008.

[27] M.-A. Dubois, C. Billard, C. Muller, G. Parat, and P. Vincent, "Integration of high-Q BAW resonators and filters above IC," in *IEEE ISSCC Digest of Technical Papers*, Feb. 2005, pp. 392–393.

[28] J. Carpentier, A. Cathelin, C. Tilhac, P. Garcia, P. Persechini, P. Conti, P. Ancey, G. Bouche, G. Caruyer, D. Belot, C. Arnaud, C. Billard, G. Parat, J. David, P. Vincent, M. Dubois, and C. Enz, "A SiGe:C BiCMOS WCDMA zero-IF RF front-end using an above-IC BAW filter," in *IEEE ISSCC Digest of Technical Papers*, Feb. 2005, pp. 394–395.

[29] M. Aissi, E. Tournier, M.-A. Dubois, G. Parat, and R. Plana, "A 5.4GHz 0.35µm BiCMOS FBAR resonator oscillator in above-IC technology," in *IEEE ISSCC Digest of Technical Papers*, Feb. 2006, pp. 316–317.

[30] S. Razafimandimby, A. Cathelin, J. Lajoinie, A. Kaiser, and D. Belot, "A 2GHz 0.25µm SiGe BiCMOS oscillator with flip-chip mounted BAW resonator," in *IEEE ISSCC Digest of Technical Papers*, Feb. 2007, pp. 580–581.

[31] J. Proakis, *Digital Communications*, 4th ed. New York: McGraw-Hill, 2001.

[32] R. Meyer, "Low-power monolithic RF peak detector analysis," *IEEE J. Solid-State Circuits*, vol. 30, no. 1, pp. 65–67, Jan. 1995.

[33] P. Gray, P. Hurst, S. Lewis, and R. Meyer, *Analysis and Design of Analog Integrated Circuits*, 4th ed. New York: John Wiley & Sons, Inc., 2001.

[34] S. Gambini, N. Pletcher, and J. M. Rabaey, "Sensitivity analysis for AM detectors," EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2008-31, Apr 2008. [Online]. Available: http://www.eecs.berkeley.edu/Pubs/TechRpts/2008 /EECS-2008-31.html

[35] Y. Wu, M. Ismail, and H. Olsson, "A novel CMOS fully differential inductorless RF bandpass filter," in *Proceedings of the IEEE International Symposium on Circuits and Systems*, vol. 4, May 2000, pp. 149–152.

[36] S. Gambini and J. Rabaey, "Low-power successive approximation converter with 0.5 V supply in 90 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 42, no. 11, pp. 2348–2356, Nov. 2007.

[37] G. Giustolisi, G. Palumbo, M. Criscione, and F. Cutri, "A low-voltage low-power voltage reference based on subthreshold MOSFETs," *IEEE J. Solid-State Circuits*, vol. 38, no. 1, pp. 151–154, Jan. 2003.

[38] H. Banba, H. Shiga, A. Umezawa, T. Miyaba, T. Tanzawa, S. Atsumi, and K. Sakui, "A CMOS bandgap reference circuit with sub-1-V operation," *IEEE J. Solid-State Circuits*, vol. 34, no. 5, pp. 151–154, Jan. 2003.

[39] M. J. Ammer, "Low power synchronization for wireless communication," Ph.D. dissertation, University of California, Berkeley, 2004.

[40] D. Markovic, V. Stojanovic, B. Nikolic, M. Horowitz, and R. Brodersen, "Methods for true energy-performance optimization," *IEEE J. Solid-State Circuits*, vol. 39, no. 8, pp. 1282–1293, Aug. 2004.

[41] L. Alarcón, T.-T. Liu, D. Pierson, and J. Rabaey, "Exploring very low-energy logic: A case study," *Journal of Low Power Electronics*, vol. 3, no. 3, pp. 223–233, Dec. 2007.

[42] A. Hajimiri, S. Limotyrakis, and T. Lee, "Jitter and phase noise in ring oscillators," *IEEE J. Solid-State Circuits*, vol. 34, no. 6, pp. 790–804, June 1999.

[43] A. M. Niknejad, "Lecture 17: BJT/FET mixers/mixer noise," EECS 142 Course Notes, University of California, Berkeley, 2005.

[44] T. Toifl, C. Menolfi, M. Ruegg, R. Reutemann, P. Buchmann, M. Kossel, T. Morf, J. Weiss, and M. Schmatz, "A 22-Gb/s PAM-4 receiver in 90-nm CMOS SOI technology," *IEEE J. Solid-State Circuits*, vol. 41, no. 4, pp. 954–965, Apr. 2006.

[45] T.-H. Lin and Y.-J. Lai, "An agile VCO frequency calibration technique for a 10-GHz CMOS PLL," *IEEE J. Solid-State Circuits*, vol. 42, no. 5, pp. 1003–1011, May 2007.

[46] Y. Chee, A. Niknejad, and J. Rabaey, "A sub-100µW 1.9-GHz CMOS oscillator using FBAR resonator," in *IEEE Radio Frequency Integrated Circuits (RFIC) Symp. Dig. Papers*, June 2005, pp. 123–126.

[47] S. Yu and P. Kinget, "A 0.65V 2.5GHz fractional-N frequency synthesizer in 90nm CMOS," in *IEEE ISSCC Digest of Technical Papers*, Feb. 2007, pp. 304–305.

[48] N. Pletcher and J. Rabaey, "A 100µW, 1.9GHz oscillator with fully digital frequency tuning," in *Proc. IEEE ESSCIRC*, Sept. 2007, pp. 135–138.

[49] M. McCorquodale, S. Pernia, J. O'Day, G. Carichner, E. Marsman, N. Nguyen, S. Kubba, S. Nguyen, J. Kuhn, and R. Brown, "A 0.5-to-480MHz self-referenced CMOS clock generator with 90ppm total frequency error and spread-spectrum capability," in *IEEE ISSCC Digest of Technical Papers*, Feb. 2008, pp. 350–351.

[50] K. Sundaresan, P. Allen, and F. Ayazi, "Process and temperature compensation in a 7-MHz CMOS clock oscillator," *IEEE J. Solid-State Circuits*, vol. 41, no. 2, pp. 433–442, Feb. 2006.

[51] F. Lee and A. Chandrakasan, "A 2.5 nJ/bit 0.65 V pulsed UWB receiver in 90 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 42, no. 12, pp. 2851–2859, Dec. 2007.

[52] D. Weinstein, H. Chandrahalim, L. F. Cheow, and S. Bhave, "Dielectrically transduced single-ended to differential MEMS filter," in *IEEE ISSCC Digest of Technical Papers*, Feb. 2006, pp. 318–319.