Design Techniques for Robust and Reliable Ultra-Low Power Radios

Jesse Aaron Richmond

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2014-192
http://www.eecs.berkeley.edu/Pubs/TechRpts/2014/EECS-2014-192.html

December 1, 2014

Design Techniques for Robust and Reliable Ultra-Low Power Radios

by

Jesse Aaron Richmond

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

in

Engineering-Electrical Engineering and Computer Science

in the

Graduate Division
of the
University of California, Berkeley

Committee in charge:

Professor Jan Rabaey, Chair
Professor Ali Niknejad
Professor Paul Wright

Spring 2013

The dissertation of Jesse Aaron Richmond, titled Design Techniques for Robust and Reliable Ultra-Low Power Radios, is approved:

Chair Date

Date

Date

University of California, Berkeley

Design Techniques for Robust and Reliable Ultra-Low Power Radios

Abstract

Design Techniques for Robust and Reliable Ultra-Low Power Radios

by

Jesse Aaron Richmond
Doctor of Philosophy in Engineering-Electrical Engineering and Computer Science

University of California, Berkeley

Professor Jan Rabaey, Chair

A popular vision in the field of wireless communications has been to achieve ubiquitous communications, in which nearly everything we interact with is connected wirelessly to a network. When combined with sensors, such a system could have applications ranging from biomedical ones, such as health monitoring sensors, to smart houses that can track their occupants' presence and automatically adjust such things as lighting, temperature, and music or video to suit the individual's preferences. In both instances, the radios that connect the sensors to the network must consume very little energy to allow long battery life or operation with energy scavenged from the environment, must be very small in size to permit them to disappear into the environment, and must be able to be created at a minimum cost to enable a large number of devices to be used.

This thesis addresses some of the design concerns facing the receivers for such sensing systems, with a particular focus on improving the sensitivity, reliability, and robustness of ultra-low power receivers. Two prototype radios have been implemented exploring different applications. The first focuses on a general purpose low power receiver for an Active RFID tag, which provides a highly integrated wireless sensing platform including a full transceiver along with power conditioning and a DC-DC converter to interface with an energy harvester. This receiver consumes 48uW of analog power, along with 61uW of digital power for demodulation and synchronization and achieves a sensitivity of -66dBm for 100kbps of data while being powered by the on-chip power supply. The second implements a wakeup radio, which operates along with a more powerful, full-featured radio such as a WiFi or Bluetooth and listens to the channel to determine if signals are present. This design focuses on minimizing the active power consumption while improving the sensitivity, and achieves a -90dBm sensitivity for a 10kbps OOK wakeup signal while consuming only 37.5uW of power.

To my wife, Sam, for helping me believe this would be possible, and to my grandfather, William Richmond, who planted the spark that encouraged me to dream big.

Contents

List of Figures

List of Tables

Acknowledgments

My time at Berkeley has changed my life in ways I would not have dreamed possible when I started thinking about applying to grad school so many years ago. The opportunity to be surrounded by so many brilliant minds has always been inspiring. I would first like to thank my advisor, Jan Rabaey, who has supported me for all these years, both giving encouragement when things have gone well and support when things have seemed off track. I am also grateful to my committee, Ali Niknejad, who helped so much to open the world of circuits that I know inhabit through courses and conversation, and Paul Wright, whose encouragement has helped keep me driven to continue even at times when I felt so stuck.

My time working on this dissertation has taken me through many twists and turns, but through it all from the beginning, I thank Cristian Marcu, Michael Mark, and Chris Clifford. From commiserating about temporary being split from our relationships through figuring out our way through classes and research, I've been ever grateful to have friends like you. I also thank Nate Pletcher and Louis Alarcon for helping to welcome me into Berkeley and grad school life. I thank my fellow groupmates David Chen, Tsung-Te Liu, Simone Gambini, Rikky Muller, Wenting Zhou, and Wen Li for keeping conversations interesting. I thank Mervin John, both for the many days and nights spent in front of computers and on lab benches trying to make our Active RFID project finally come together and for becoming such a good friend away from the lab.

I thank the many staff who have made life so much easier over the years. The Berkeley Wireless Research Center really leaves us all spoiled by how well things work and how much support we have, and much of the credit goes to the people behind the scenes who keep things going, especially Tom Boot, Gary Kelson, Deirdre McAuliffe, Brian Richards, Fred Burghardt, Olivia Nolan, and Leslie Nishiyama.

I give special thanks to two of my professors from my days at the University of Tennessee, Ben Blalock and Paul Crilly. Without your courses, encouragement, and support, I never would have embarked on the journey that has led me to here.

Finally, I am forever grateful to my loving and supportive family. My parents have always been my biggest champions, believing in whatever I do and driving me to always aim high. Lastly, I thank my wife, Sam, who was willing to go along with my crazy plan to move across the country for this. Her support has kept me going when things seemed darkest, and she has always helped me recognize when things are at their brightest. I look forward to many more adventures to come.

Chapter 1

Introduction and Background

1.1 Wireless Sensing and the Sensory Swarm

In recent years, wireless communications have undergone a dramatic transformation. For much of the history of radio, wireless has been used as a broadcast medium, where a radio or television station uses a very high power transmitter to send information to many different receivers. While two-way communications were possible, they tended to be used for limited applications like emergency services. Over the last couple of decades, however, this has been changed dramatically. Innovations such as the use of cordless and cellular telephones, wireless internet, and other devices have led to an explosion in the number of radios we each interact with regularly.

As wireless devices have become more ubiquitous, the vision of what can be connected and networked together has grown. From the popular press and the consumer industry, there has been talk recently of the "Internet of Things," [19], in which everyday objects like our appliances are connected to a network performing functions like monitoring and control. This idea has been expanded in the technical press into the "Sensory Swarm" [41], illustrated in Fig. 1.1. Because we have reached the point that many people are carrying powerful computing platforms such as laptops, smart phones, and tablets on a regular basis, we can begin to think about expanding out the reach of the connected devices into many sensors, constantly monitoring and interacting with the world around us. These sensors can range from the mundane, like monitoring the temperature in a room and dynamically controlling the thermostat settings based on its use, to the profound, like recording brain activity and using the data to provide things like artificial limbs for amputees.

In order for this swarm of hundreds to thousands of radios per person to be practically realized, there are a few major technical concerns that need further work. We can think of a few major requirements for these systems, such as:

1. **Very low cost.** This vision involves orders of magnitude more wireless devices for each of us than we currently possess. While people are buying more communications technologies all the time, in order for this leap to be practical, the price per radio also needs to fall considerably.

2. **Very low maintenance.** It is not realistic to expect people to be willing to expend much time and effort on things like replacing batteries, troubleshooting wireless connections, or correcting settings. This points to the use of things like energy scavenging to extend usable lifespan, which also requires very low power sensors, very simple protocols to prevent communications problems, and easy to use software applications.

3. **Very small size.** In order for the vision of the sensory swarm to be fully realized, these wireless sensors should blend seamlessly into their environment, with the people around them not even noticing their presence. This means that the sizes must

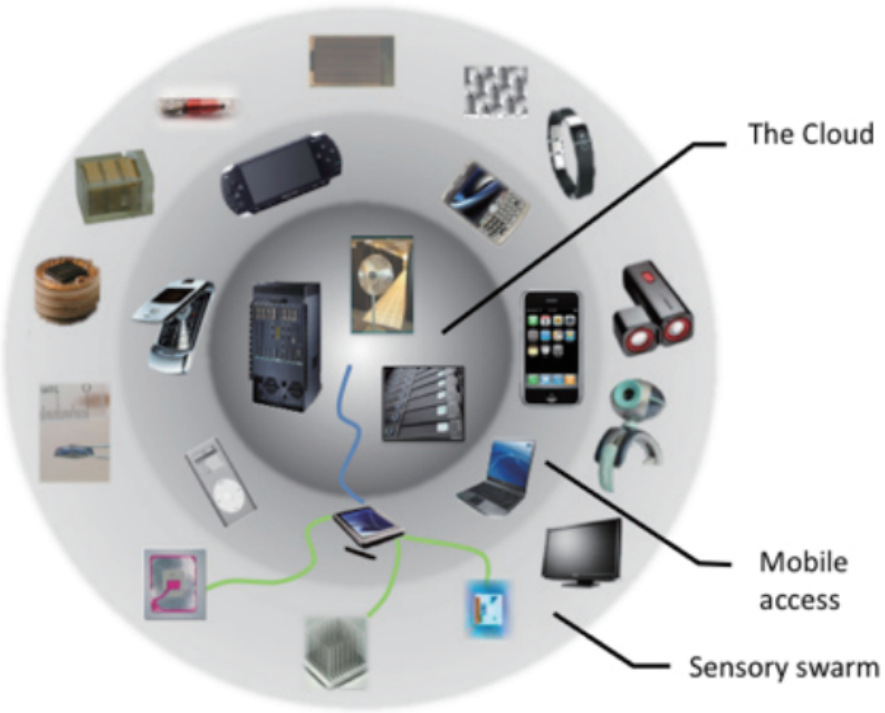

Figure 1.1: Concentric rings of communication (Image courtesy of Jan Rabaey in [41]).

be kept to a minimum and aggressive form factors should be explored to make their use as easy as possible.

Achieving the full vision for such devices will require efforts from across a wide variety of disciplines, including things like energy harvesting and storage, packaging techniques, software and application development, and many others, much of which falls outside the scope of this work. The primary goals here will be to investigate techniques that will allow the receivers for such systems to be realized and to test these techniques in two proof-of-concept designs implemented and tested in silicon.

1.2 Thesis Organization

This chapter has given a brief overview of the background of wireless sensor networks and some of the current trends in their design and implementation. The remainder of this thesis will provide a look at the important design issues that need to be considered for the design of low power radios, as well as two prototype implementations of low power radios.

In particular, Chapter 2 begins with a brief history of the design of receivers, with a particular focus on designs used for low power radios. It then discusses the impact of CMOS technology, with some thoughts on the direction of future scaling and its impacts on low power design. Finally, it concludes with the current state of MEMS resonators, which provide many opportunities for low power circuit design.

The first design example, discussed in Chapter 3, implements a low power radio as a part of a bigger system for wireless sensing. Based on an Active RFID protocol, this design uses duty cycling and a large amount of integration to bring together a complete wireless communications platform for sensor networks.

The second example, in Chapter 4, explores the trade-offs required for a wakeup radio, which is a specialized form of receiver used for minimizing both the active power consumption and the latency required for communication. The radio implements energy detection for a coarse wake-up command and points to a path to implement more complex modulation schemes within the same architecture.

Finally, Chapter 5 provides a brief conclusion and some future directions for research in the area of ultra-low power radios.

Chapter 2

Architecture and Technology Overview

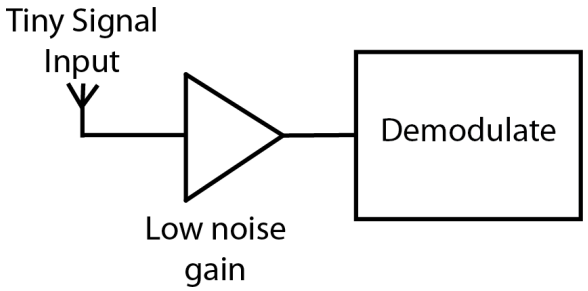

Figure 2.1: General digital receiver architecture

Achieving ultra-low power consumption for reliable wireless communications requires careful evaluation of many different design parameters. Because of the relentless pace of innovation in the technologies available, it is important for a circuit designer working on low power designs to keep up to date with the changes in technology, and careful thought should be given if possible to ensure the best technology for the task.

This chapter discusses the impact of some of these decisions. It begins with an overview of receiver architectures that have been used, and discusses the merits of each. It then gives some general guidelines for the architectures chosen for the design implementations described in the following chapters of this work. The next sections cover recent directions in CMOS and MEMS technologies.

2.1 A Brief History of Receivers

The basic architecture of most digital wireless receivers follows the same general principle, illustrated in 2.1: a small input signal is given to the system, which first has to somehow increase the signal level to something that can be further processed, without adding too much excess noise. The radio then demodulates the data from the signal, gaining back the original digital bits that were sent. Of course, for a real system, there are a number of additional factors that make the architectures typically used more complicated, and it is useful to explore the traditional designs of architectures to determine the most important factors to consider for creating low power designs.

A major factor to note is that the general principle illustrated above does not consider that the purpose of modulation is, in part, to allow multiple streams of data to be communicated simultaneously without interfering with each other. One of the primary ways this goal is accomplished is by using many different frequencies, and relying on the radios to filter out the frequencies outside the band of interest. A challenge involved in designing these filters is that it is very difficult in a typical radio design to get the bandwidth narrow enough to select only the channel of interest.

In the early days of radio design, the active devices used were vacuum tubes, which were bulky and relatively expensive, so it was desirable to minimize the complexity of

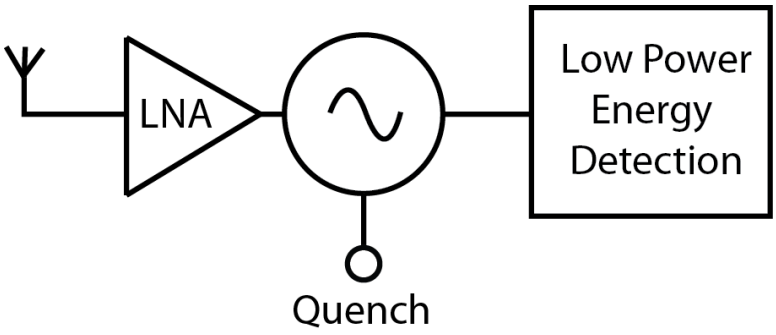

Figure 2.2: Super-regenerative Receiver Architecture

the radio design as much as possible to keep the cost and size down. This drove the development of the super-regenerative radio [7], illustrated in figure 2.2, which uses an oscillator to both produce a large amount of effective gain as well as a relatively sharp frequency response. More recently, the relative simplicity and small number of components in the super-regenerate architecture has been used to create low power radios, particularly based around low power MEMS-based oscillators [35]. This architecture presents two main drawbacks limiting its ultimate effectiveness in ultra-low power design. The first is that the power consumption of the receiver is limited by the oscillator itself, which in published works has required at least 89μW simply to start oscillation [12]. In addition, the topology inherently requires an on-off keyed modulation, limiting the flexibility in network standards that the receivers can operate.

The two basic challenges of receiver design - achieving high gain while selecting only a portion of the total spectrum - are both easier to accomplish at lower frequencies than at high frequencies. Rather than trying to process the signal at RF, the high frequency signal can first be mixed to a lower frequency, where it is filtered and amplified, then mixed down to baseband for demodulation. This technique, illustrated by the block diagram in 2.3 and a basic frequency plan in 2.4, is known as super-heterodyne, and dates far back in the history of radio design [8]. The mixer is driven by an oscillator referred to as a local oscillator (LO), which provides a reference frequency, ensuring that the desired signal falls in the correct location for the channel selection filter. The basic concept of this architecture, in which the intended received signal is translated to a lower frequency for processing, has become the dominant way to design receivers since its introduction and still forms the basis for most architectures today [26].

The standard super-heterodyne architecture poses some significant challenges for low power design, however. Because the circuits preceding the channel-selection filter experience a wide range of possible input frequencies, the design of each of these elements must accommodate a wide dynamic range. This means that they must provide both a high amount of gain at RF frequencies to minimize noise contributions as well as high linearity to avoid interference from unwanted frequencies creating tones at the same frequency as the desired signal, corrupting the ability to detect it. In addition to the linearity

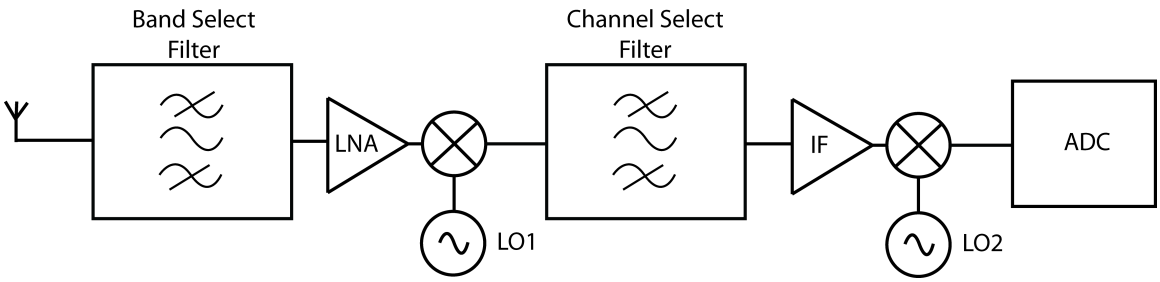

Figure 2.3: Traditional super-heterodyne architecture

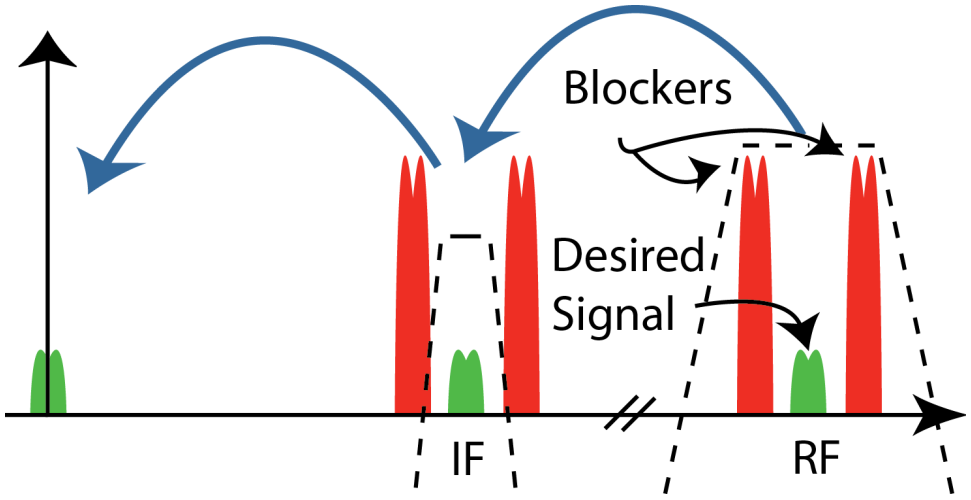

Figure 2.4: Frequency plan for a super-heterodyne receiver.

requirements, the super-heterodyne architecture presents a relatively complex architecture, involving a large number of circuit blocks. These requirements limit the ability of the architecture to achieve low power consumption, and as a result, even less stringently power-constrained radio designs, such as for mobile phone and laptop computer applications, have tended to move away from its use in recent years.

Many recent receivers have used a slightly different version of the super-heterodyne architecture known as low-IF [16]. In a traditional super-heterodyne radio, the intermediate frequency is placed well above DC and the second downconversion is performed in the analog domain using a second set of mixers along with a second LO. In contrast, a low-IF receiver places the signal close to DC, such that it becomes feasible to directly sample the IF signal and perform the final downconversion in the digital domain. This removes some of the complexity required in the super-heterodyne architecture, potentially allowing for a lower power consumption receiver. A version of the same basic concept has been recently implemented for a wakeup receiver [21], which uses an envelope detector and a clocked input to perform the first conversion, and then uses the same clock to sample the IF for the final conversion to baseband. This allows for a lower power consumption than a traditional super-heterodyne or low-IF architecture, but follows the same basic concept. Low-IF receivers have many of the same drawbacks and limitations as more traditional super-heterodyne designs, but also introduce an additional challenge resulting from the close placement of the image frequency to the desired band, requiring sharp filtering or image-rejection circuits and increasing their complexity.

More recently, many radio designs have been moving towards an architecture called direct conversion, illustrated in 2.5 [4]. This architecture was developed in part to reduce the complexity of receivers as a way of minimizing the power consumption required, which is a key driver of circuit designs for mobile applications. The basic difference compared to the super-heterodyne architecture is that the receiver only includes a single mixer, which takes the incoming signal directly from RF down to baseband so that the receiver can directly sample the signal and perform the demodulation. The key advantage of this architecture compared to the super-heterodyne architecture is that it requires fewer precision components; in particular, it only requires a single precise frequency for the local oscillator, tuned to the same frequency as the incoming RF signal. There are some drawbacks to the direct conversion receiver. One notable issue is that some of the LO energy can leak back through the mixer to the front end, which causes two major issues. First, it can become an interferer to other nearby radios by emitting a tone in their signal bands, and second, it can mix with the desired incoming signal, which causes time-varying DC offsets to appear. In addition, because the energy is brought directly to baseband, it is more challenging to deal with flicker noise in the baseband. Despite these drawbacks, direct conversion has become one of the most common design techniques for recent radios.

The direct conversion architecture looks like a good possibility for low power radio designs, but it is worth looking more closely at the individual components to identify areas where the power consumption can be further limited. In particular, for typical radios, the front-end filtering is still limited to a relatively wide-band channel selection filter, because

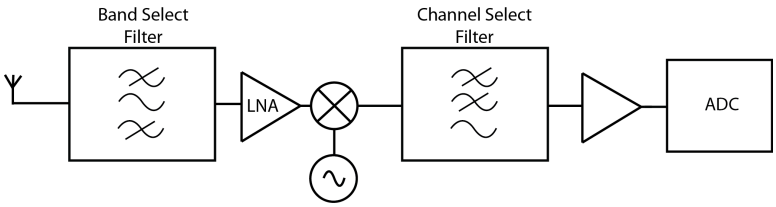

Figure 2.5: Direct conversion or low-IF receiver block diagram.

current devices limit the quality factor of the filters. This means that additional signals aside from the desired channel can enter the front end of the receiver, leaving the same stringent linearity requirements as with super-heterodyne radios. In addition, because energy from both above and below the LO frequency is translated down to baseband, direct conversion receivers must operate in quadrature, as both sidebands will fold on top of each other at baseband. This requires more complexity in the signal path, with two mixers and a second baseband chain, which will substantially increase the power consumption required by the receiver.

To improve performance at low power levels, there have been a number of approaches to improving the linearity and filtering performance of receivers with low power consumption. In [15], the radio makes use of low Q passive elements at the front end for image rejection, and a narrow bandpass filter shape at baseband for its channel selection. The drawback of this approach is that it requires a precision frequency from the local oscillator, because this directly will set the effective frequency of the bandpass filter at the RF frequencies.

Another approach used by many low power radios is to use high-Q passive elements such as MEMS resonators (which will be discussed in more detail in 2.4) to directly filter away interference after the antenna. Then, the receiver can either provide gain directly at RF to feed the demodulator, as demonstrated in [38], or can follow the technique of mixing the signal down to a lower frequency for further processing. An advantage of performing the filtering first is that it is less important to limit the bandwidth of the following gain stages, because there is little danger of interfering with the signal. This led to the uncertain IF architecture proposed in [39] and illustrated in 2.6 and 2.7, in which the baseband is a relatively wide bandwidth to allow for the local oscillator driving the mixer to be less accurate, which permits the use of such low power oscillator designs as digital ring oscillators to be used, rather than more power hungry PLL-based LO designs or resonant oscillators. This architecture can be thought of as a modification to the super-heterodyne architecture that removes the final channel-selection filter that precedes the baseband amplification chain. In addition, this receiver does not use a mixer to perform the second downconversion, but instead relies on the simplicity of the on-off keyed modulation to allow a simple envelope detector to perform the downconversion.

In addition to the above architectures, an emerging class of radios known as ultra-wideband (UWB) radios has recently been used for low power applications [25]. The basic idea behind these radios is to use short bursts of radio power typically in the form

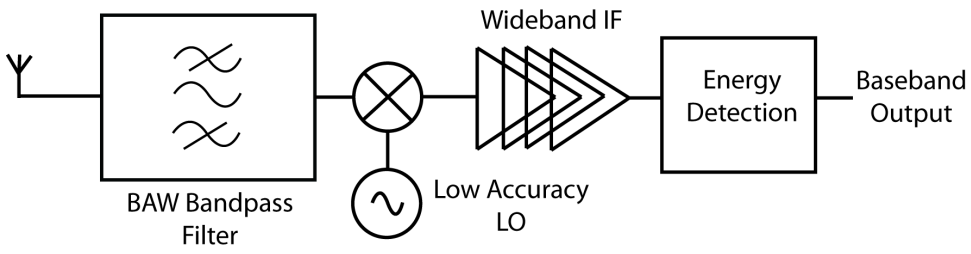

Figure 2.6: Uncertain IF architecture design

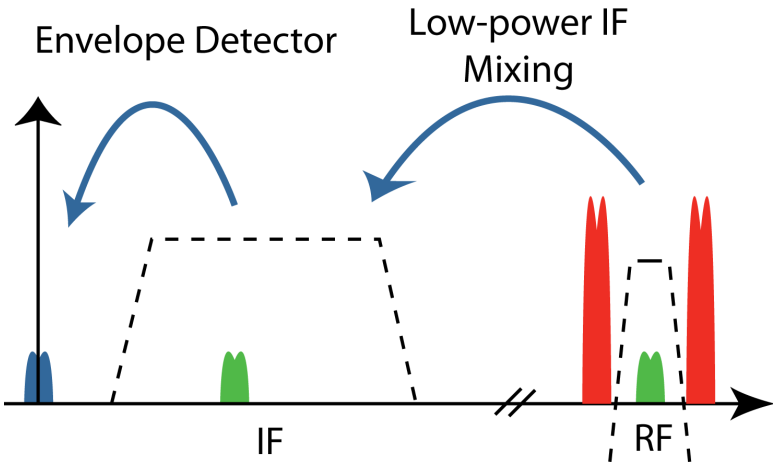

Figure 2.7: Uncertain IF Frequency Map

of pulses rather than a continuous sinusoidal signal. The idea for low power UWB radios is that the circuits need only operate during the periods of time when a burst of RF energy is present, and can be aggressively duty-cycled during other periods. This can allow for very energy-efficient designs when using certain metrics, such as the energy per bit of information [5]. While many applications can benefit greatly from these radio designs, such as applications featuring a steady stream of information transferred like streaming audio or video, for sensing applications that feature sporadic and sparse transmissions of small amounts of data, such as the wireless sensor networks targeted by the works in this dissertation, the overhead power required for timing and synchronization in these UWB systems makes them unattractive topologies. In particular, for the wake-up radio featured in Chapter 4, UWB becomes especially difficult to implement because by definition, UWB radios spread their information across a wide range of frequencies and require the receiver to collect this energy back together to demodulate the data. Without being precisely synchronized with the transmitter, these radios are highly susceptible to narrowband interference, making a wake-up radio implementation impractical. In addition, these radios tend to feature very high instantaneous power when operating, such as the 38mW required by [25], despite the low average power due to aggressive duty cycling. This fast cycling between low standby power and high on power also presents challenges to the DC-DC converter employed by the system, which has to cope quickly with large changes in power levels.

Section 2.1 gives an overview of the most promising recently published radios for achieving ultra-low power consumption. A few general trends can be discerned from these papers. For those designs using an oscillator based either around an LC tank as in [15] or around a PLL as in [25, 17], generating the LO consumes a large fraction of the overall power consumption. In addition, these oscillators act as lower bounds on the power consumption that is achievable from such architectures, which limits the attractiveness of such designs for ultra-low power application. The papers with the three lowest power consumptions, [38, 39, 21], all operate with an OOK modulation scheme, which allows for simple energy detection-based demodulation. Two of these, [38] and [21], are designed to perform a substantial portion of their total gain at RF, followed by energy detection operating at RF. The third, [39], first uses a mixer driven by a ring oscillator LO to mix the incoming signal down to an IF, where the receiver performs the majority of the gain before encountering the energy detector.

2.2 Architecture Exploration

2.2.1 Receiver Performance Metrics

There are many ways to characterize receiver performance. Some of the commonly used metrics include data rate, sensitivity, power consumption, linearity, and interference rejection. There also are many secondary measures, such as energy efficiency, which

	[15]	[25]	[38]	[39]	[21]	[17]
Technology	130nm	90nm	90nm	90nm	90nm	65nm
Architecture	Low-IF	Direct Conv.	Tuned RF	Uncertain IF	Tuned RF / Low-IF	Low-IF
Active Power	330μW	250μW Avg.	65μW	52μW	123μW	415μW
Data Rate	N/A	100kbps	100kbps	100kbps	10kbps	500kbps
Sensitivity	N/A	-99dBm	-56dBm	-72dBm	-86dBm	-82dBm
Frequency	2.4GHz	3-5GHz	1.9GHz	1.9GHz	915MHz	2.4GHz
Modulation	2-FSK	Pulsed UWB	OOK	OOK	OOK	PPM-IR

Table 2.1: Review of recently published low power radios.

attempts combine data rate and power consumption by quantifying the energy required per received bit. One of the keys to optimizing a design for ultra-low power applications is to appropriately choose the best metrics to target with the design.

As an example, many applications, such as streaming audio or video, require a large data rate and continuous transmission of data, so targeting a minimal energy efficiency without concern for overhead from synchronization or idle time may be the most appropriate metric. It is for this reason that often works focusing on UWB, 60GHz, and other high data rate applications prefer to focus on reducing the energy per bit, because most of their energy is expected to be spent on actually receiving data. In addition, the total active power in these radios may be quite high, simply because the amount of data being transferred is very large.

In contrast, a wireless temperature sensor may only need to collect data very occasionally and transmit in short packets, making the overhead to acquire signals and synchronize the radio much more important. Simply focusing on energy efficiency in this case may lead to a sub-optimal design, because there is a much smaller fraction of time that the receiver is expected to receive useful data. In addition, because of the limited total amount of data that may be transferred, it is much less critical to achieve a high data rate.

Metrics such as linearity and interference rejection are important to all wireless systems, and must be dealt with in some fashion. In architectures such as the superheterodyne and direct conversion receivers, the requirement to coexist with other wireless signals that may be more powerful requires that the first stages of the receiver be very linear to avoid distortion tones from falling on top of the desired signal. Another way to mitigate interfering signals is to filter them out as soon in the chain as possible, making the linearity of the active circuits less critical for proper operation. This will be explored further in Section 2.4.

The target application of this work is in wireless sensing platforms, with a particular focus on applications like building environmental controls, health monitoring, and smart energy systems. These applications often feature use cases where much of the time, the sensor is idle, so any power spent on keeping the wireless connection active is simply wasted, until a central controller needs to send information to the sensor to perform operations like taking a data point or issuing an emergency interrupt. Because these sensors are extremely limited in their power storage, either by the amount that can be scavenged from the environment or stored in small batteries, it is critical that the active power consumption be reduced as much as possible, even if this means that some other metrics like data rate must be traded off.

In many wireless sensor networks, the active power is managed through the use of duty cycling, in which the radios are turned off for periods of time, reducing their power consumption to a minimal standby level. Typical commercial receivers used today often have active power consumption levels on the order of 1-10mW, however, so to achieve total power consumption on the level of 10's of μWs as required by systems such as those described in [40], the duty cycle must be extremely low, such as 10^{-3}. This extreme duty cycling has a number of important consequences to the design of the network. One com-

plication is that keeping a large number of nodes in proper synchronization so that they all wake up during the very short time window available to them means that an accurate timing reference is required, which consumes additional wasted energy. Another complication is that there can be very long latency times between available cycles, which may not be acceptable in time-sensitive applications like healthcare monitoring or emergency control systems.

An alternative to duty cycling, presented in [27], is to use a wake-up radio with an extremely low active power consumption, to constantly monitor the channel. This radio listens for commands destined for its wireless node, and then wakes up a more powerful, general purpose radio when a signal is received. This work demonstrated that a wake-up radio can improve the overall system power consumption if it uses less than about 50μW of active power.

2.2.2 Sensitivity of Receivers

One of the key metrics of operation for any receiver is its sensitivity, which essentially is a measure of the smallest amount of input power the receiver requires in order to correctly demodulate a signal. Traditionally, this figure is determined to be the input power required for the bit error rate (BER) of the receiver to fall to 10^{-3}, in which one bit in a thousand is received incorrectly. There is a wide range of typical sensitivity ratings across different classes of radios, from cellular radios that often are better than -110 dBm up to passive RFID tags that require -15dBm. A receiver's sensitivity, coupled with the maximum transmit power of the sending end of the wireless link, determines the maximum range that the link can operate. The sensitivity of a receiver can be estimated in a very straightforward way from a few simple parameters as

$$Sens = 10 \cdot log(kT) + 10 \cdot log(BW) + NF + SNR_{min}\ dB \tag{2.1}$$

where the first term in the sum gives the thermal noise power, BW is the effective noise bandwidth of the system, NF is the noise figure of the receiver, and SNR_{min} is the minimum signal to noise ratio required to detect the given modulation scheme. In traditional radio designs, the only real aspect of this equation that can be directly affected by the circuit designer is the NF of the radio, which is a large reason that the NF of a receiver's front end receives such attention. The thermal noise power is a physical phenomenon, so there is no way to change this parameter. In a typical radio design, the BW is made as close to the required bandwidth for the design's modulation and data rate as possible through the channel select filter; however, this simply represents the minimum amount of bandwidth required by the receiver, or a lower bound on the amount of noise bandwidth that can be achieved. The SNR_{min} is a function of the modulation scheme used by the receiver, and can be readily calculated from communications theory.

The noise figure of a receiver can be calculated as

$$NF = 10 \cdot log\left(\frac{SNR_{in}}{SNR_{out}}\right) = 10 \cdot log\,(F) \tag{2.2}$$

in dB. This figure tells how much additional noise the receiver itself adds to the signal while processing it, so it is important to minimize it as much as possible in order to receive weak signals. Typical NF values vary significantly, from roughly 1dB for the most sensitive radios used in devices like cellular phones up to as high as 20dB for some low power radios such as [39]. The noise figure is largely dominated by the first stages of the receiver chain, because the contribution of following stages is reduced by the gain of the following amplifiers as

$$F_{total} = F_1 + \frac{F_2 - 1}{A_1} + \frac{F_3 - 1}{A_1 \cdot A_2} + ... \tag{2.3}$$

where $F_n$is the noise factor of each stage and $A_n$is the gain of each stage. The relationship in equation 2.3 explains why in most receiver designs, the LNA is the critical block that largely determines the overall noise of the receiver, because its gain will act to reduce the impact of the noise of the following stages. It also reinforces that in all receiver chains, the emphasis needs to be on the first few stages of the chain, because their noise performance will dominate the noise performance of the whole receiver.

With passive components like resistors, calculating the amount of noise that is generated is quite straightforward and depends only on the component value, but active components are more complex as their noise depends both on the size of the devices as well as their bias levels. To gain insight into the design tradeoffs that are available, it is useful to explore rough estimates of the amount of noise generated by simple active stages. We can calculate the noise of a single transistor as

$$\overline{v_{on}^2} = 4kT \cdot \frac{\gamma}{\alpha} \cdot g_m \cdot \Delta F \tag{2.4}$$

where k is Boltzmann's Constant, T is the temperature in Kelvin, γ is a technology-specific parameter ranging from approximately $^2/_3$ in long-channel devices to 2-3 in short-channel devices, and g_m is the transconductance of the device at its given bias point. While in general, active stages may use a large number of transistors and other components, it is useful to consider a very simple topology like the basic common-source amplifier, because this topology forms the basis for many other types of amplifiers [20]. We can estimate the noise figure of a common-source LNA as

$$F = 1 + \frac{\gamma}{g_m \cdot R} \tag{2.5}$$

where R represents the source impedance of the preceding stage. This equation demonstrates that the noise figure is inversely proportional to the transconductance of the transistor. Given that the transconductance is directly proportional to the drain current, we can see that the noise figure of an amplifier is also inversely proportional to the drain current. In addition to the dependence on the transconductance, we also can see that the source impedance has an impact on the noise figure of an amplifier. Eq. 2.5 also shows that increasing the source impedance will reduce the noise figure for a given power consumption.

2.2.3 Low Power Design Techniques

In many low power radio systems, the LNA is removed, with the antenna instead directly driving the mixer stage to reduce the power consumption [39, 15]. The danger of this approach is that the noise figure performance is typically degraded by a mixer, such as the roughly 20dB noise figure from [39]. In contrast, passive mixers have been shown to be able to produce quite low noise figure levels when used as the first stage in [6] and [15], reaching as low as 3dB in the former and 7dB in the latter. In addition, because passive mixers are essentially switches, they benefit from CMOS scaling, which will act to reduce the capacitive load they present to the oscillator, reducing the power required to drive the mixer. These characteristics make the passive mixer appealing for the purposes of an ultra-low power receiver, allowing for both good performance as well as low power consumption.

In traditional architectures, the frequencies of the local oscillators are precisely controlled to ensure that the RF signal is mixed down to the proper frequencies for the narrow filters that remove any excess noise and interference from the signal path. Typically, this is controlled with the use of a PLL, which locks the signal to a multiple of a precise reference frequency generated by a quartz crystal oscillator. The drawback of these systems is that they require a significant amount of power. In the design examples listed in 2.1, none of the narrow-band architectures make use of a PLL for their frequency generation because of the power overhead required for their operation. Alternatively, the LO can be generated by use of a high-Q component such as a MEMS resonator. Unfortunately, with current technology for these devices, which will be discussed in further detail in 2.4, the power consumption is still higher than the $50\mu W$ power limit desired for low power radios [12], limiting its usefulness for an always-on wakeup receiver. In contrast, the approach from [39], using a ring oscillator to generate the LO is particularly attractive given further device scaling in CMOS technologies. A ring oscillator's power consumption should fall as the technology scales due to lower device capacitances, plus the ring oscillator is capable of generating a full-rail swing with ease using CMOS inverters, which maximizes the drive strength on a passive mixer's switches.

The downside of using a ring oscillator is that the frequency accuracy and phase noise performance is known to be quite poor. The phase noise can be mitigated by the use of a high-Q filter preceding the mixer, which can eliminate interference that would otherwise be mixed by the phase noise into the desired signal band, but the frequency accuracy requires that the IF bandwidth be made much wider than the nominal bandwidth for the signal's information content. This directly impacts the sensitivity of the receiver as shown in equation 2.1, by increasing the bandwidth term in the calculation. The amount of degradation to the sensitivity due to excess bandwidth is shown in 2.8, calculated assuming a 100kbps data rate. For a 100MHz IF bandwidth, such as that used in [39], the loss in sensitivity approaches 30dB, presenting a substantial degradation due to the poor frequency accuracy of the oscillator. This can be mitigated by either improving the accuracy of the oscillator, thereby increasing the precision that the frequency can be

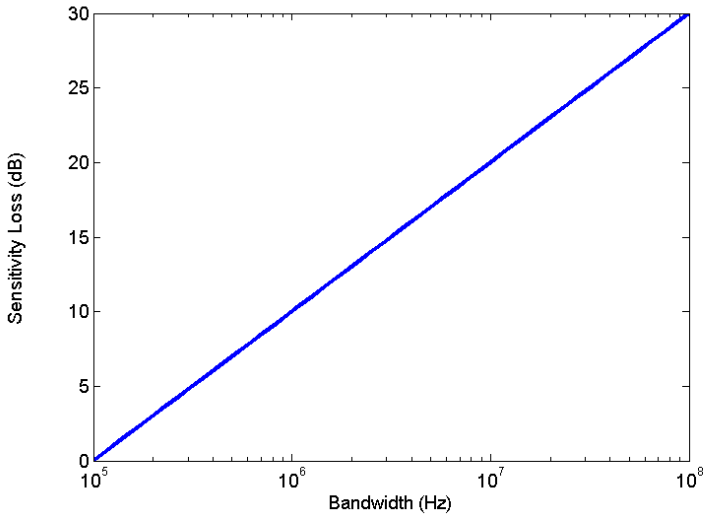

Figure 2.8: Sensitivity loss due to excess bandwidth, given a 100kbps data rate

selected, or by finding a way to filter out the excess noise without knowing the precise frequency the signal content occupies.

This work will present two approaches for each of these techniques. In the work presented in Chapter 3, the oscillator is implemented with a large number of bits in the DAC that controls the oscillator's frequency, allowing the LO to be tuned to within a narrow band. It also uses a digital oversampling technique from [45] that allows the receiver to partially filter out some of the excess noise in the digital domain. The second prototype, presented in Chapter 4, uses a temperature compensation circuit to remove some of the source of uncertainty in the oscillator's frequency, which allows the bandwidth to be narrowed, along with an analog integration-based technique to reduce the impact of excess noise in the IF band.

2.3 CMOS Technology Directions

A fundamental driving principle of the entire electronics industry is that over time, the transistors in CMOS technologies get smaller [32]. Traditionally, this has meant that the devices also get higher performance, reaching higher possible operating frequencies and requiring lower power consumption for digital logic as discussed in [42]. As the devices have continued to shrink, however, this traditional scaling has become more difficult to achieve, and a divergence in the scaling has occurred. Instead of all processes following the same general trends identified by the International Technology Roadmap for Semiconductors [1] as occurred for technologies as recently as 90nm, newer processes starting with the 65nm generation have diverged into multiple paths, most notably for either

Node	GP VDD	GP VT	LP VDD	LP VT
29nm	1V	285mV	0.95V	428mV
27nm	0.97V	289mV	0.95V	436mV
24nm	0.93V	294mV	0.85V	407mV
22nm	0.9V	291mV	0.85V	419mV
20nm	0.87V	221mV*	0.8V	311mV*
18nm	0.84V	221mV*	0.8V	317mV*
17nm	0.81V	229mV*	0.75V	320mV*
15.3nm	0.78V	228mV*	0.75V	323mV*
14nm	0.76V	232mV*	0.7V	327mV*
12.8nm	0.73V	236mV*	0.7V	299mV**

Notes: Values marked with * indicate the use of fully-depleted SOI technology, while those marked ** indicate the use of multi-gate devices.

Table 2.2: 2009 ITRS Roadmap projections for general purpose and low operating power technologies

low power applications or for general purpose and high performance applications. The primary difference in these two technology options is that the low power versions focus heavily on reducing the leakage power that tends to limit the power consumption for low power digital logic [29] while shrinking the size of the devices, reducing die sizes, while the general purpose processes continue to push the performance of the devices along with shrinking their physical dimensions. The projected trends from the ITRS roadmap for two pertinent values for circuit designers, the nominal supply voltage and the nominal threshold voltage of an NMOS device, are illustrated in Table 2.2.

This trend means that designers face a choice in their technology beyond simply the size of the devices. For ultra-low power designs, it would seem most obvious to choose the low power technology, given its name and target application; however, this choice has more subtle implications than are first seen. To truly optimize for power consumption, it is useful to remember the most basic calculation of the power required by a circuit as

$$P_{dissipated} = V_{DD} \cdot I_{total} \tag{2.6}$$

where the V_{DD} is the supply voltage and I_{total} is the total current required, or $I_{active} + I_{leakage} + I_{analog}$. This is, of course, an extremely simplified and high-level view of the power, but it illustrates an important point: designers can minimize the power consumption required by either reducing the supply voltage or by reducing the current consumed by each block. It is useful, then, to determine what are the limits and trade-offs presented by focusing on either the supply or the current and how these trade-offs will be impacted by the choice of device technology.

To begin, we can consider the impact of digital logic. The active power consumed by

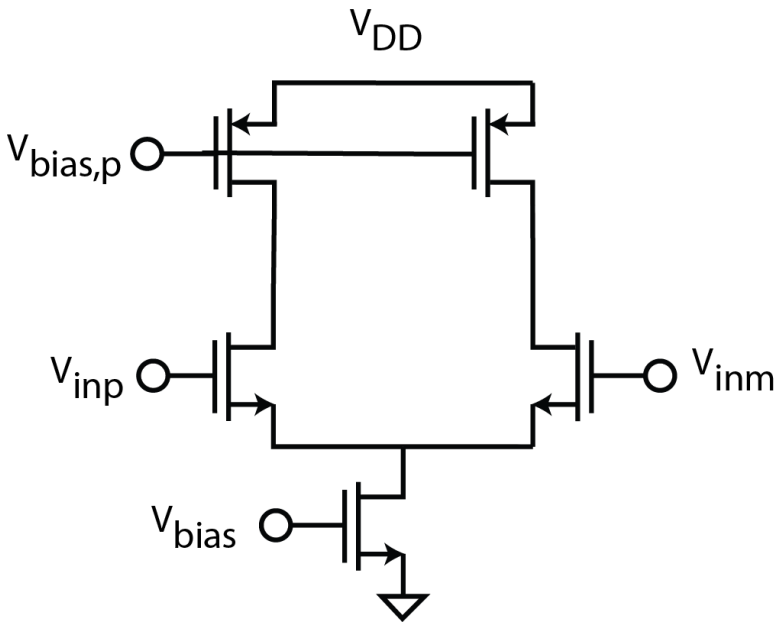

Figure 2.9: Typical analog gain stage.

a digital CMOS circuit is found by

$$P_{active} = \frac{1}{2} V_{DD}^2 \cdot C_{eff} \cdot f_{switch} \tag{2.7}$$

where C_{eff} is the amount of capacitance that is switched in each cycle on average and f_{switch} is the frequency that the logic is operating. This means that reducing the supply voltage for digital logic has a quadratic impact on the power consumed by the logic, so reducing the supply voltage significantly reduces the power consumed. The leakage power can be found by

$$P_{leakage} = V_{dd} \cdot I_s \cdot e^{\frac{\sigma \cdot V_{DD} - V_T}{n \cdot \phi_t}} \tag{2.8}$$

where I_s is a technology-specific parameter representing the beginning of inversion, σ accounts for the DIBL effect of short-channel devices, n represents the subthreshold slope, and ϕ_t represents the thermal voltage. This term has an exponential relationship with supply voltage, indicating that the power in absolute terms also falls with supply. For low supply voltages and digital-heavy circuit designs, however, the leakage current can come to dominate the overall power consumption because the speed of the logic also falls quickly with a reduction in supply [30]. For the purposes of low power radios, however, there is typically very little digital logic, so the digital tends to not dominate the power consumption.

Analog circuit designs have somewhat different constraints. In a digital design, reducing the supply voltage directly impacts the circuit's performance by reducing the speed at which operations can be performed. In contrast, in analog design, the requirement on

the supply voltage is simply to be high enough to support the bias points needed by each of the transistors. In Fig. 2.9, the minimum supply would be such that each of the three transistors remains in saturation, plus sufficient overhead to support the required gate voltages for each device. The drain current produced by a given gate bias can be found by the following generalized equation, which holds in all regions of operation from weak inversion through strong inversion [50]:

$$I_{DS} = \frac{W}{L}\mu C_{ox}(2n)\phi_t^2 \left\{ \left[\ln\left(1 + e^{(V_{GS}-V_T)/(2n\phi_t)}\right)\right]^2 - \left[\ln\left(1 + e^{(V_{GS}-V_T-nV_{DS})/(2n\phi_t)}\right)\right]^2 \right\}. \tag{2.9}$$

This shows that in all regions of operation, the gate voltage is a function of the threshold voltage; this means that a higher threshold voltage will directly translate into a higher minimum supply voltage. This means that in general, for low power analog design, it is advantageous to operate with devices having the lowest threshold voltage possible. In contrast, we have seen in equation 2.5 that the noise figure of a circuit is inversely proportional to its transconductance, so it is advantageous to increase the transconductance of at least the first stage of amplification as much as possible to minimize the noise of the receiver. Putting the two together, we can see that for analog circuits, a desirable approach is to minimize the voltage supply as much as possible while maximizing the current allotted to the first gain stages. We also can see that the analog devices encourage the use of the lowest thresholds possible from the technology, because the threshold essentially wastes voltage headroom, restricting the radio from operating at a lower voltage supply level, which therefore raises the required power. Reducing this voltage headroom reduces the available dynamic range to the receiver, but this presents a worthwhile tradeoff for ultra-low power designs.

This means that, for low power analog and RF design, it is preferable to use the CMOS technology flavor with the lowest threshold. Given the general purpose and low power technology split, this points towards the general purpose devices as actually yielding potentially lower power consumption than low power technologies, assuming that the power is dominated by analog components and that it is possible to use a low supply voltage. Because a wakeup receiver will also have to share space with a more traditional radio, there may be a design challenge presented because the traditional circuit blocks are often designed to operate with a higher nominal threshold due to their different tradeoffs in the design. This does mean that a holistic, system-wide view of power distribution needs to be taken into consideration, and may push towards the use of highly efficient, on-chip switched capacitor DC-DC converters such as found in [48], which allow the supply voltage of individual blocks to be tailored to an optimal level, which helps to optimize the overall system power consumption.

2.4 MEMS Technology Directions

2.4.1 Why MEMS?

One of the most critical components in practical radio designs is the filter that allows the receiver to select only the particular set of frequencies in which its information may be found. For a bandpass filter, a key metric is its quality factor, found as

$$Q = \frac{f_c}{BW} \tag{2.10}$$

where f_c is the center frequency of the filter and BW is its bandwidth. In order to filter just a specific, narrowband frequency, the Q of the filter needs to be quite high; for example, to achieve a 1 MHz bandwidth filter centered at 1GHz, the Q of the filter needs to be 1,000.

One of the most common ways to create such filters is by using a passive resonant structure, such as an RLC network made up of a resistor, inductor, and capacitor. The difficulty is that in order to achieve a filter without a large amount of loss, it is necessary to use components whose Q is also very high, but achieving high quality passive electrical components is very challenging to impossible. As an example, a typical on-chip inductor Q would be in the range of 10-20 [26].

Another common technique to achieve the frequency selectivity required for narrowband receivers is by utilizing the frequency translation provided by the super-heterodyne or direct conversion architectures to move the signal down to a lower frequency. As shown in Eq. 2.10, the required Q of the filter is proportional to the center frequency, so by first moving the frequency down to a lower frequency, the requirements on the passive components can be relaxed, making the filter more practical. In addition, the lower operating frequency makes it practical to implement the filters using active filter designs, which can offer better performance than purely passive designs. There are a number of drawbacks to this technique that limit its attractiveness for low power receivers, however. Because the filtering happens after the frequency conversion, the mixer and any preceding active devices must be made very linear, because they can be exposed to interference with a very high dynamic range, which tends to increase the required power consumption. In addition, in order to select a very accurate frequency at either IF or baseband, the oscillator frequency must be very well controlled, which does not allow for low power ring oscillators to be used.

In contrast with electrical components, mechanical vibrations can achieve very high quality factors. An example of a high-Q mechanical vibration is actually a common metal bell, which, when struck, can continue to vibrate for many seconds, indicating a large amount of energy is being retained in the system. While large size mechanical resonators tend to have low resonant frequencies, in recent years, a substantial amount of research has been conducted into the use of micro-electromechanical systems (MEMS) that use small sizes to achieve resonant frequencies as high as several gigahertz, with Q of over

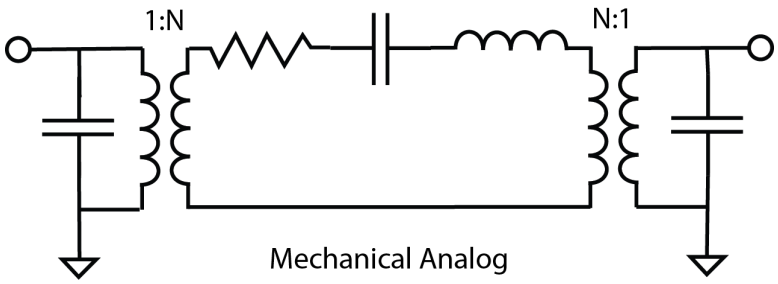

Figure 2.10: General model of a MEMS resonator.

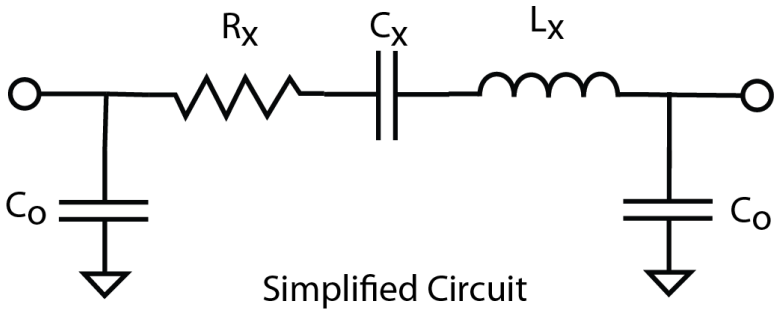

Figure 2.11: Simplified MEMS model

1,000 [34, 47]. This very high Q can allow the filtering to take place directly following the antenna, removing any interference before it has the opportunity to encounter nonlinear elements. In addition, removing interference early in the chain relaxes requirements on the stability and accuracy of the LO generation, allowing for much lower power consumption.

2.4.2 Types of MEMS Resonators

There are a number of different shapes of MEMS resonators in use, ranging from beams to discs to stacks of materials, but all resonators follow the same basic pattern of operation: an electrical signal is converted into a mechanical force, either through electrostatics or piezoelectrics, and this force is applied to a mechanical structure. At the proper frequency, the structure itself begins to vibrate and forms a strong resonance, which can then be converted back into an electrical signal through the same mechanism that generated the original force.

At a high level, all MEMS resonators can be modeled by using the generalized model in 2.10. The electrical elements forming the RLC network directly correlate to mechanical properties of the resonator [49]. The resistor represents the loss due in the system, due to such things as energy transfer to the substrate or compression of the gasses surrounding the resonator. The capacitor models the inverse of the spring constant of the material as it is deformed from the vibration, and the inductor represents the effective mass that

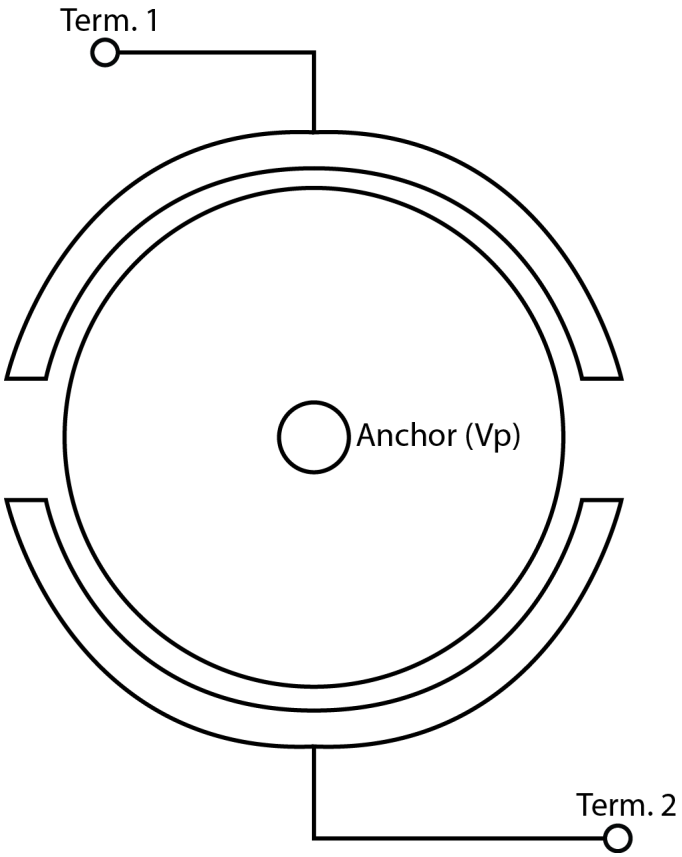

Figure 2.12: Basic design of a vibrating radial disc MEMS resonator

is in motion. The transformer on both ends represents the transduction between the electrical and mechanical domains, and the capacitors represent parasitic capacitances associated with the input and output terminals of the device. Because this model is rather cumbersome for use in electronics design, typically the transduction elements are omitted and applied directly to the RLC network, transforming the model into the simplified circuit shown in 2.11.

There are two broad categories of MEMS resonators, defined by the type of transduction between the electrical and mechanical energy domains: electrostatic and piezoelectric transducers. Electrostatic resonators utilize the attraction formed by a separation of positive and negative charges. The transduction is formed by a capacitor whose plates have a gap of d_o, and as charge is manipulated on the two plates of the capacitor, a mechanical force is exerted on the plates. The other method of transduction is formed by the use of piezoelectric materials, which deform in the presence of an electric field. These structures also utilize a capacitor, but rather than an open gap between the plates, the piezoelectric material such as aluminum nitride forms the dielectric.

From a circuit designer's standpoint, the voltage applied to a resonator translates

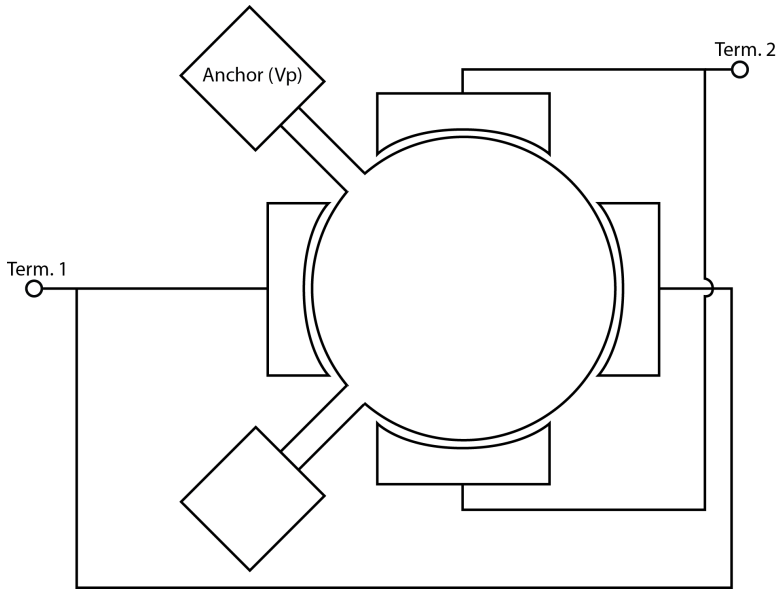

Figure 2.13: Basic design of a wineglass-mode vibrating disc resonator

into the force that is applied to the mass of the resonator, so the movement of the mass is similar to the charging of the capacitor in the model in Fig. 2.10. The momentum of the mass, then, translates into the current moving through the inductor, which is being damped by the various loss mechanisms modeled by the resistor. The transduction determines how strongly the electrical voltage is converted into mechanical force. This means that a stronger transduction results in a lower impedance for the resonator, with a smaller capacitance and a larger inductance. Note that the transduction, in itself, does not contribute loss to the system, so changes in the transduction strength do not, by themselves, affect the quality factor of the resonator.

There are a large number of different shapes that can be used to implement resonators, both with electrostatic and piezoelectric designs. Previously implemented designs have included vibrating beams, which act similarly to guitar strings, vibrating discs, which can exhibit a number of different modes of operation depending on the design of the transducer, and structures in which the entire bulk of the resonator vibrates, which is commonly used with piezoelectric devices. These structures each have a number of advantages and disadvantages, yielding trade-offs the designer must make between things like quality factor, operating impedance, frequency control, and fabrication ability.

In terms of maximizing Q, some of the most effective resonators that have been demonstrated include vibrating disc resonators, illustrated by both a radial mode, in which the vibration happens uniformly around the radius of the disc in 2.12, and wineglass-mode resonators, which expand and contract in opposite directions along the X and Y axes, shown in 2.13. These resonators are able to achieve a very high quality factor by anchoring the resonators at points with minimal movement, which reduces the amount of

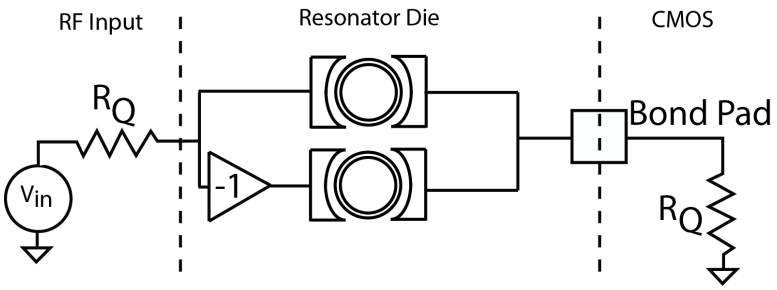

Figure 2.14: MEMS-based

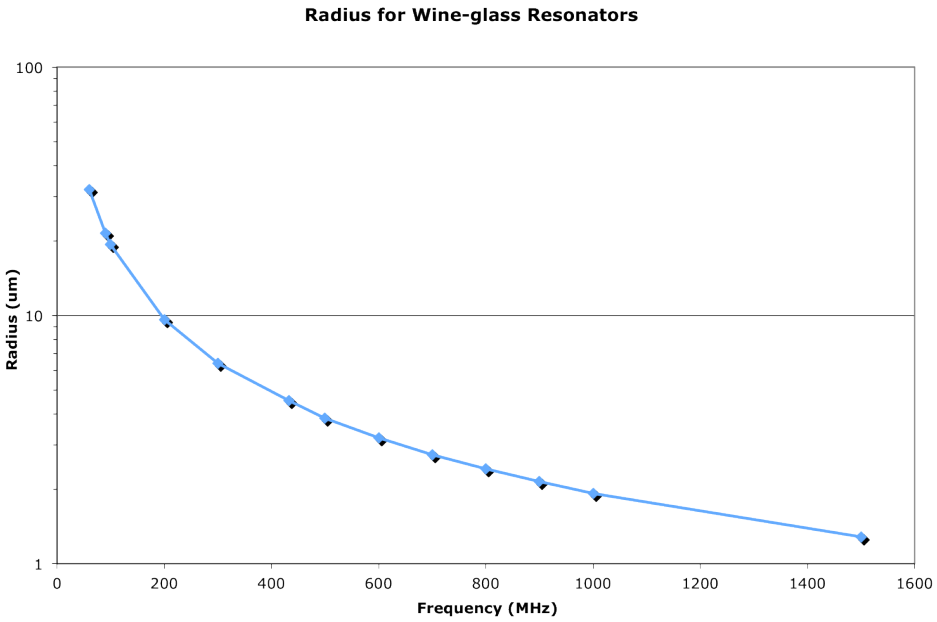

Figure 2.15: Radius required for a given resonant frequency with a wineglass disc resonator

vibrational energy lost to the support structures. In the case of the radial disc, the resonators have been shown to be capable of achieving a Q greater than 11,000 at 1.5GHz [10], while the wineglass mode can reach 98,000 at 73.4MHz [3]. With such a high quality factor, it would be possible to implement a very low loss, very small bandwidth filter using the structure shown in 2.14. An additional beneficial feature of this style of resonator is that its frequency is determined by the radius of the device, as shown in 2.15, making it straightforward to implement multiple frequencies in a single lithographic step. This is important for minimizing the number of components needing and therefore minimizing the cost of the additional devices.

While these resonators are attractive for their ability to create high quality bandpass filters, they have a number of drawbacks making them currently impractical for imple-

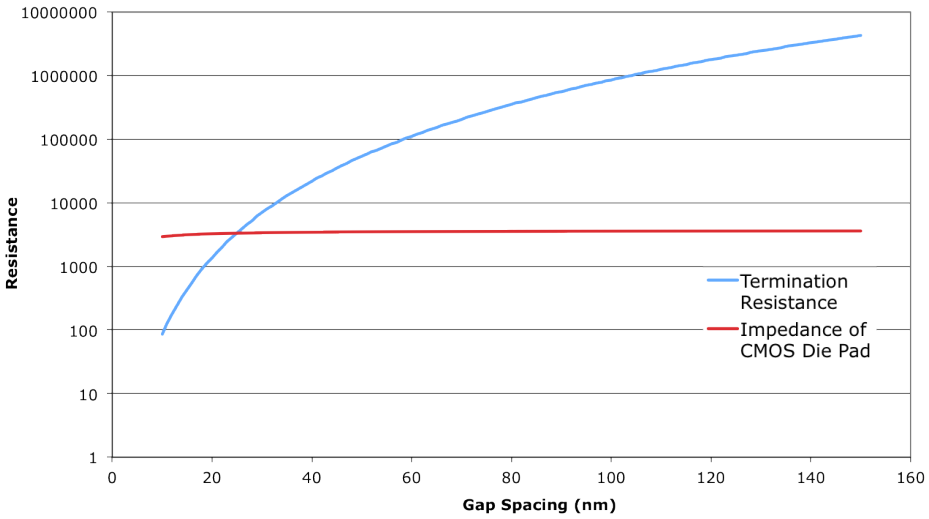

Figure 2.16: Termination impedance required for a filter with electrostatic MEMS resonators

mentation in integrated systems. One issue is that they typically have a much lower quality factor in air than in a vacuum [3], resulting in the need for costlier vacuum packaging to guarantee good performance. More problematic for integration with CMOS is that they require very high bias voltages and have very high impedances. Because the electrostatic transducer operates using a capacitor, its strength is strongly dependent on the gap spacing between the plates. In [28], it is shown that the motional resistance, R_x, is proportional to d_o^4, so it is very important to reduce the gap spacing as much as possible. For the high frequency resonators most applicable to radio technologies from [10], the series resistance at resonance with currently achievable gap spacings is on the order of 1MΩ, which is much higher than typical impedances achieved by active circuits in CMOS technologies and much higher than the impedances created by the bond pads and wire traces between the resonator and the CMOS die. Fig. 2.16 compares the required termination resistance to properly terminate the filter structure from Fig. 2.14 with the typical impedance found by a CMOS bond-pad presenting a parasitic capacitance of 100fF to ground. In order for this type of filter to be practical, it would be necessary to achieve impedances much lower than the impedance caused by the parasitic capacitance of the pad in order to prevent distorting the filter shape. As can be seen, the gap spacing would need to reach down to less than 20nm, which presents significant challenges for the manufacturability of the devices with current technologies.

MEMS resonators based on piezoelectric structures have seen much more success in industry applications. In particular, the thin-film bulk acoustic wave resonator, or FBAR, has been utilized recently to provide low-loss filters for cellular phone applications [47, 46],

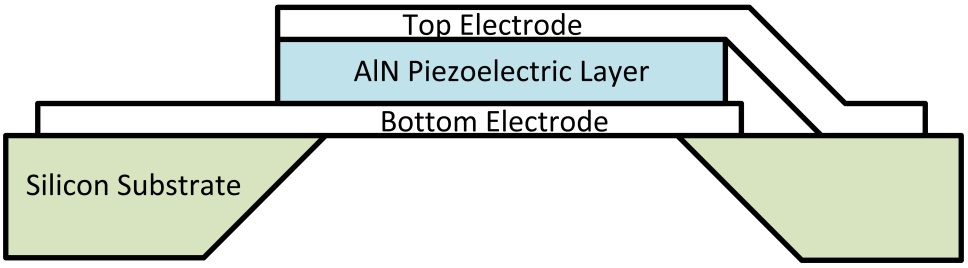

Figure 2.17: Diagram of the structure of an FBAR resonator

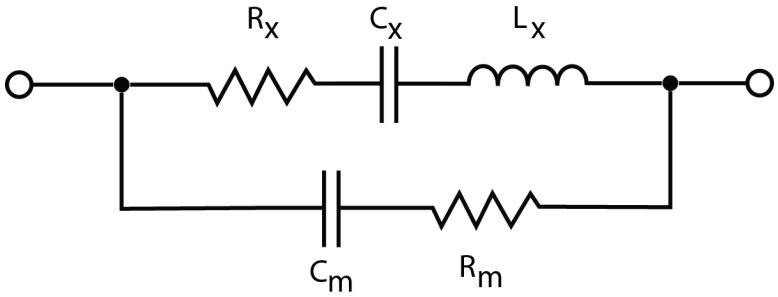

Figure 2.18: Electrical model of the FBAR resonator

where they are used as band-selection filters and duplexers to assist with the co-existance between the transmit and receive chain. These resonators use a piezoelectric transduction method, which provides a much stronger transduction than electrostatic, but comes at the penalty of increasing the loss in the structure and therefore reducing the quality factor. The basic structure of the FBAR, shown in 2.17, is two metal conductors above and below a piezoelectric insulating material, much like a parallel plate capacitor. The piezoelectric insulator begins to vibrate when an electrical signal is applied across the terminals, causing the entire stack of the metals and insulator to vibrate. The structure is isolated from the substrate by either removing the material from underneath the resonator via an isotropic etch process [47] or by building a stackup of materials to act as a Bragg reflector to impedance match the resonator to the substrate [23, 31].

The electrical model of the FBAR, shown in 2.18, is similar to the model of the generic MEMS resonators shown above, but with the capacitors on the terminals merged into a single shunt capacitor representing the physical capacitance between the two terminals of the device. This means that the resonator exhibits two separate resonant modes: a series mode, when the series RLC branch reaches its resonance, as well as a parallel mode where the C_o capacitance resonates with the effective inductance of the series branch. The measured response of a typical 987MHz FBAR is shown in 2.19, and a table of the values of the equivalent circuit model is shown in 2.3.

In addition to their use in traditional wireless radios, FBARs have seen extensive use in recent years in low-power applications, where they are used for both filtering elements [39,

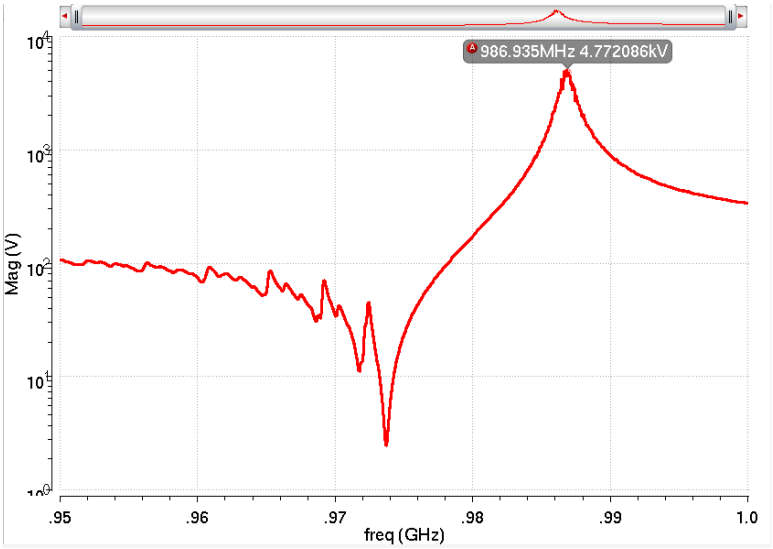

Figure 2.19: Measured impedance of a 987MHz FBAR resonator

Component	Circuit Value
R_x	2.6Ω
L_x	1.378μH
C_x	19.38fF
R_m	4Ω
C_m	730fF

Table 2.3: Component values for a 987MHz FBAR resonator

44] as well as elements for low power oscillators [12, 13, 35, 37]. The primary advantage of the FBAR compared to electrostatic based resonators lies in its easily interfaced electrical values. The series impedance is only 1-4Ω, while the parallel impedance is on the order of 2-5kΩ, making it easy to integrate with typical CMOS impedance levels. While the Q of FBARs has not been able to reach the levels shown by electrostatic resonators, they still can reach as much as 5,000, which is much higher than can be achieved with electrical passives such as inductors. This allows them to achieve both very narrowband filtering as well as low phase noise, low power oscillators.

In addition to their ability to act as high quality filters, FBARs have also been used to create low power oscillators that still have excellent phase noise and accuracy performance [36, 12, 33], but the power consumption that has been achieved does not allow for operation with a complete radio consuming less than 50μW. In any oscillator, the negative resistance provided by the active circuits must be enough to overcome the loss from the resonator, so for a resonator with a 2kΩ impedance, which is fairly typical for a resonator such as that shown in Figure 2.19 after including the parasitic impacts of the bonding, and a transistor biased in moderate inversion, providing a good trade-off between parasitic capacitance and efficiency, with a $g_m/I_D = 20$, the minimum current that can possibly start an oscillator can be found as:

$$I_{min} = \frac{2}{R_p \cdot g_m/I_D} = \frac{2}{2k \cdot 20} = 50\mu A \tag{2.11}$$

In practice, the current bias needs to be somewhat higher, both to ensure reliable startup as well as to overcome any loading effects from the surrounding circuits. While the lowest reported power consumption FBAR oscillator, from [33], only requires 22μW for startup, making it nominally able to support the target power consumption for the receiver in this work, this value does not include the buffering that would be necessary to drive the mixer and it only supports a single-ended oscillation, limiting its applicability in many architectures. For this reason, the works presented here focus on the use of a ring oscillator to generate the LO, rather than a MEMS oscillator. With further improvements in MEMS technologies for oscillators, however, it may become practical to consider their use again.

For the purpose of this work, it was decided that it was more practical to use the mature and reliable FBAR technology rather than the more experimental electrostatic resonators. Research work on MEMS resonators is a very active field that has been changing quickly in recent years, so for future work, alternative resonator structures could be useful.

Chapter 3

Example Project 1: Active RFID

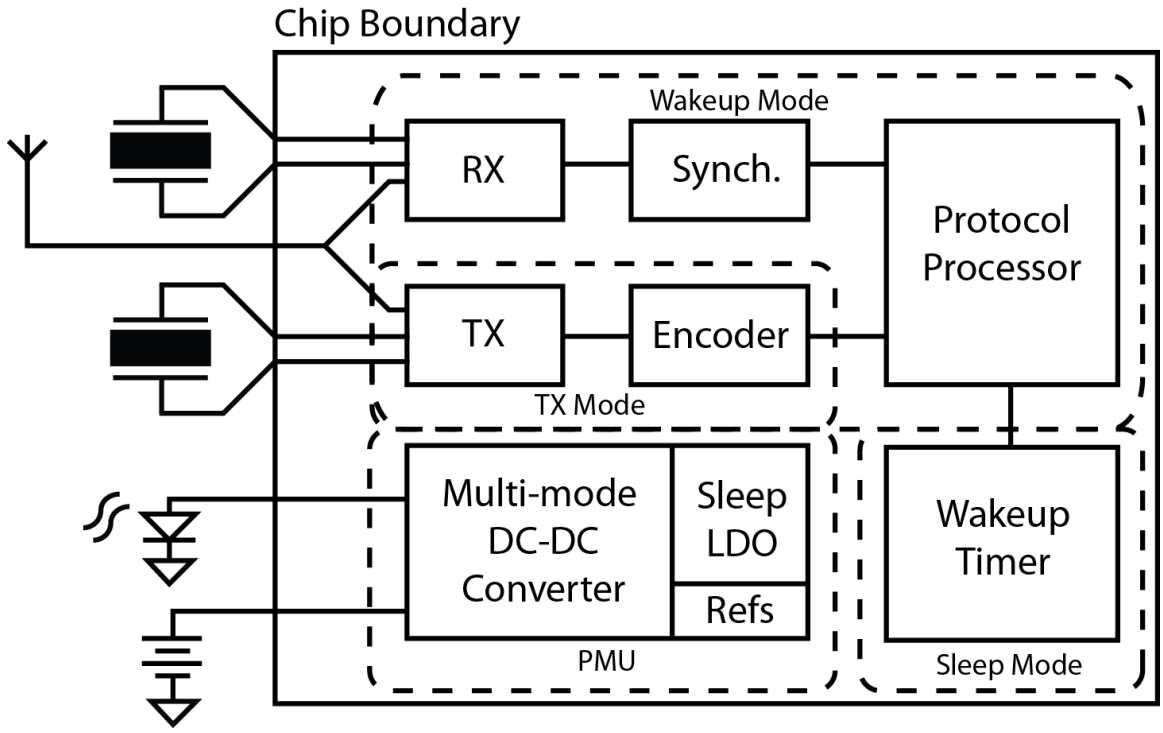

Figure 3.1: Block diagram of the Active RFID system

Chapter 2 gives an overview of some of the issues and trends impacting ultra-low power receiver design. The next two chapters explore the impact of these trends in prototypes of receivers, for two different applications.

3.1 Background and Goals

With the recent push for energy efficiency in buildings, transportation and logistics, and industrial processes, wireless sensor networks (WSN) have become a promising technique to help monitor the world and prevent waste. In order to be most effective, these sensors need to be inexpensive, easy to deploy, and require little to no maintenance after installation. The work in this chapter proposes one way to achieve these goals by utilizing highly scaled CMOS, recent advances in MEMS components, and energy harvesting. Careful system design with multiple modes of operation allows a wireless communications platform to operate perpetually by harvesting ambient energy from its environment through solar cells to target indoor sensing applications that require a moderate range of roughly 10 meters and moderate data rates.

Active RFIDs [2] supplement the low cost and ease of use of traditional RFIDs with an enhanced interrogation range and superior processing capabilities, as essential for usage as wireless sensor nodes. Active RFID links combined with advanced sensing platforms

Location	Power Available in μW
Collaborative area	1.89-3.67
Window on cloudy day	32
Desk in main work area	10
Classroom	23
Front office	18
Bathroom	3-22

Table 3.1: Available solar power in the Berkeley Wireless Research Center

[14] can substantially extend the applicability of wireless sensing networks. This chapter presents a highly integrated prototype of an active RFID platform, combining a full wireless platform for sensing, consisting of a very low power transceiver, a full power conditioning system for interfacing with both a battery for energy storage along with a solar panel for energy harvesting, plus the digital control circuitry to implement the standard and maintain the system's energy consumption at an acceptable level.

In order for sensing platforms such as this to be practical for widespread deployment, it is critical that the cost of the system be minimized. In addition, for these types of sensors to be accepted into the home for uses such as home automation, they need to be made small enough that they can blend in with their environment such that the sensors themselves can be installed once and then forgotten by the end user. One of the most important strategies to minimize cost of a system such as this is to integrate as many components into a single die as possible to reduce the bill of materials that the end user would need to use. For this prototype, the only external components required are a 25x10mm solar panel, a small NiMH battery for temporary energy storage, a printed antenna, and two FBAR resonators. The total design, including its power harvesting and antenna, are intended to take up roughly the space of a postage stamp, allowing the devices to be placed throughout the environment without being obtrusive to the user. The wireless link supports 100kbps of on-off keying (OOK)-modulated Manchester-coded data operating in the well-known 2.4GHz ISM band to aid in compliance with wireless standards around the globe. The block diagram of the chip is shown in 3.1.

In order to ensure that the tag can operate perpetually from the solar energy it can harvest, it is first necessary to get a realistic estimate of the amount of energy available in a variety of settings. To facilitate this, the solar panel was placed in a number of locations within the indoor office environment of the Berkeley Wireless Research Center. 3.1 shows the results of this experiment, demonstrating that most locations can provide between 3 and 30μW of power during the daytime, depending on the amount of light available.

The amount of available energy is too low to provide power to the full transceiver and digital baseband at all times, so the system is designed to provide duty cycling to reduce the average power draw. For typical sensing applications, it is not necessary for the circuitry to all be powered on continuously, because the data need only be collected peri-

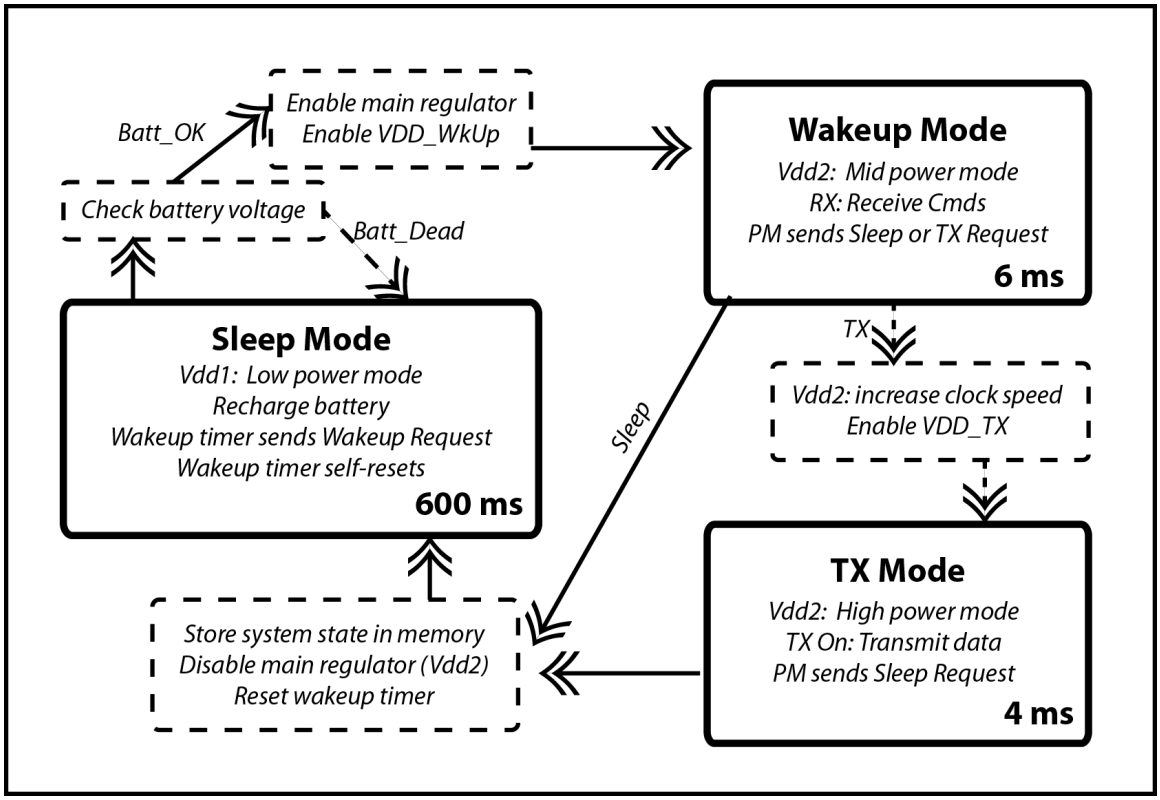

Figure 3.2: Operational modes of the Active RFID system

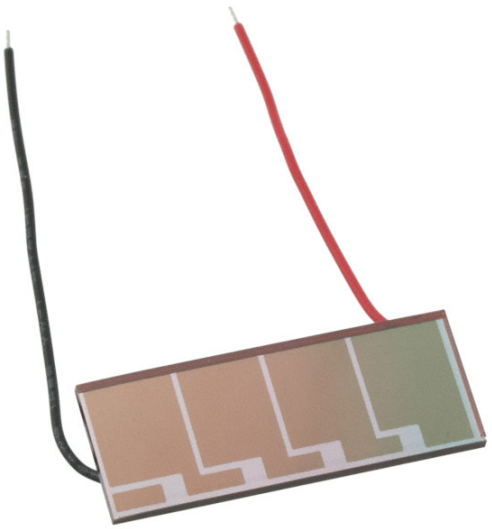

Figure 3.3: General example of solar panels intended to power the Active RFID chip

odically. The system requires an average power of 3μW under the duty cycle conditions between the three modes of operation specified in 3.2. During the sleep mode, the only active components are a wakeup timer (controlling the period of the sleep cycle), system state registers, and the power management unit (PMU). Upon wakeup, the receiver and synchronizer listen to the channel for Manchester-encoded data and wake the protocol processor upon detecting synchronization. If a valid tag request is received, the protocol processor prepares a packet acknowledging the reader and transmits its unique identification code during TX mode. Otherwise, the system reenters sleep mode and resets the wakeup timer.

The solar panel shown in 3.3, consisting of 4 cells in series, provides an open-circuit voltage up to 2.8V. The harvested energy is directly stored on the battery to ensure that operation can be maintained even during the nighttime hours when the lights in an indoor environment are typically turned off. To minimize active power, circuit blocks operate at reduced supply voltages ranging from 450mV to 650mV. Because the different modes of operation require a very large range of active power, it is necessary for the power management unit to be able to efficiently generate the proper supply for each mode. Switched-capacitor circuits offer the possibility to provide this energy both in a highly integrated implementation as well as with very high levels of efficiency [48, 24].

Although the full Active RFID implementation covers a wide range of circuit blocks and functionality [44], the scope of this work covers only the receiver design. The following sections detail the implementation and performance of the receiver and the important trade-offs that must be balanced for its operation.

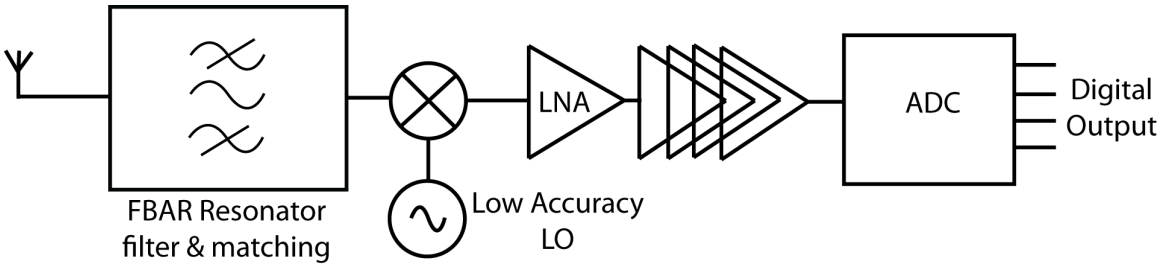

Figure 3.4: Block diagram of the Active RFID receiver

3.2 Circuit Implementation

Because the primary goal of the Active RFID project was to demonstrate the high level of integration required to have an entire sensing platform on a single die, it was decided that the basic architecture would follow the same principle as the uncertain-IF architecture from [39], due to its demonstrated performance and low power consumption, shown in Fig. 3.4. The primary goals for this implementation were:

1. Identify challenges presented by the integration of other components such as the transmitter and on-chip voltage conversion.

2. Implement a complete receiver solution from the analog input to the antenna through digital bits to the protocol processor, including the digitization of the output.

3. Provide sufficient resolution to allow the digital baseband to acquire synchronization and detect data.

The on-chip DC-DC converter inherently will produce a voltage ripple, caused by its attempt to keep the supply voltage regulated to the desired level. Minimizing this supply ripple requires either a larger decoupling capacitor, at the penalty of an increase in die area, or a faster clock frequency on the main capacitors, which results in a lower efficiency. For this implementation, the converter is designed to yield a roughly ±25mV output ripple on the main 550mV supply rails. To prevent this ripple from causing a degradation to the sensitive receiver components, a low drop-out linear regulator was implemented that was designed to reduce the supply ripple by an order of magnitude. This regulated supply is used for the analog front-end to the receiver, with the digital for the demodulation and protocol implementation running off the main DC-DC converter rail.

3.2.1 Antenna Interface and Matching Network

One important consequence of integrating all the components together for the Active RFID design is that the antenna interface becomes a critical element to the design that must be carefully considered to avoid causing problems for either the transmitter or

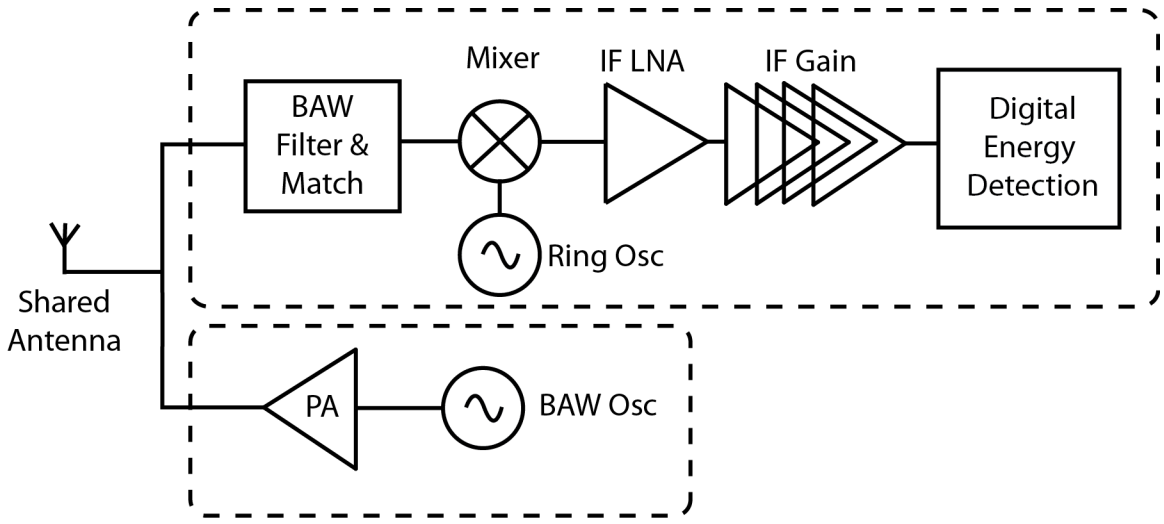

Figure 3.5: Block diagram of the complete transceiver

receiver. In order to practically implement a network consisting of dozens or more sensors, the physical size of the device is very important [11]. The antenna is one of the larger physical components to a wireless device operating in the 2.4GHz band, which means that to meet the small size requirements, it becomes necessary to share the antenna between the transmitter and the receiver, as illustrated by the complete transceiver block diagram in 3.5.

In a typical wireless system, the antenna is designed to be a standard 50Ω impedance, which has become the standard impedance for interconnections between RF components. The choice of impedance originates in a trade-off between power handling ability and lossiness of the coaxial cabling [26], but presents some challenges for low power design particularly for the transmit side. For a given target power output, and a fixed supply voltage, we can calculate the optimal resistance for the power amplifier load as

$$R_{opt} = \frac{4V_{DD}^2}{P_{out,pk}} \tag{3.1}$$

because the maximum output voltage swing is twice the maximum supply voltage, due to the use of an inductive load from the power amplifier. Using the target numbers of a 500mV supply and a 0dBm output power, we can calculate that the optimal impedance for the output of the power amplifier is roughly 140Ω. Traditionally, this impedance is created by the use of a matching network between the output of the PA and the antenna; however, due to the low Q achievable in CMOS processes, the matching network leads to a loss of power and, therefore, a lower efficiency from the transmitter. For this reason, the transmitter uses the topology from [13], illustrated in Fig. 3.6. The value of R_{ant} is set equal to the impedance calculated by equation 3.1, while the value of L_{ant} is chosen to resonate out the parasitic capacitance presented by the combination of the power

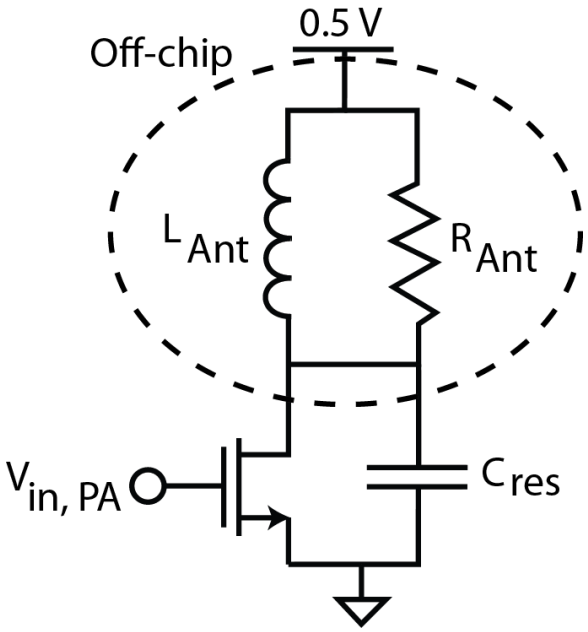

Figure 3.6: Power amplifier schematic, illustrating antenna impedance

amplifier's output device and the loading of the receiver.

In contrast with the PA, the receiver does not have as strict of a requirement on the optimal antenna impedance, because it is necessary to use a matching network to present the best noise and power transfer performance in general anyway. For this work, it was decided to have the matching network operate with the optimal impedance for the power amplifier to assist in maximizing the transmitter's efficiency. The receiver matching network uses a tunable capacitive transformer network as shown in 3.7 that resonates with the effective inductance of the FBAR resonator to provide a match between the antenna and the parallel resonance of the FBAR, which can be calculated as

$$R_p = \frac{1}{\omega_o^2 C_{eff}^2 \left(R_x + R_{cap}\right)} \tag{3.2}$$

where ω_o is the resonant frequency of the FBAR, C_{eff} is the total amount of capacitance in shunt with the resonator, taking into account both the parasitic capacitors of the receiver input, pads, and PCB along with the designed capacitors for the transformer, R_x is the motional resistance of the FBAR, described in 2.4, and R_{cap} represents the effective resistance due to the finite Q of the capacitors in parallel with the FBAR.

As shown in equation 3.2, the design of the matching network is complicated somewhat by the fact that the parallel resistance of an FBAR is dependent upon the amount of capacitance that is loading it. Fortunately, because the resonator's parallel resistance is much higher than the typical resistance of an antenna, often reaching the range of 500-1500Ωs, the value of C_1 needs to be higher than C_2, and since the capacitors appear

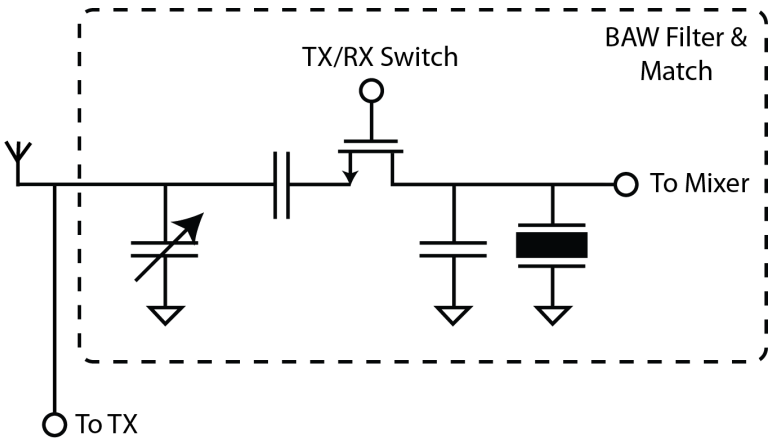

Figure 3.7: Matching network and TX/RX switch

roughly in series in the circuit, the total value is not impacted by the variations in C_1 as badly as might be expected. In a traditional radio design with a standard 50Ω antenna impedance, the capacitor C_1 value needs to be large enough to absorb the parasitic capacitances associated with the bond pads and the PCB trace; in contrast, for this application, part of the capacitance at the antenna connection resonates with the inductive component of the antenna, reducing the amount available for the matching network. The capacitors were implemented using MIM capacitors, which provide both high quality factors, that do not degrade the Q of the resonator significantly along with high density, so that they take up a small amount of die area. The capacitor C_2 was designed for a value of 120fF, while C_1 is expected to be 400fF and implemented as a 4-bit binary weighted switchable capacitor array with a 45fF LSB size.

The final major consideration for the design is that the antenna is shared between the power amplifier and the receiver. This necessitates some kind of a switch to accommodate the two modes of operation, because if the receiver's matching network remained in place, roughly half the power from the transmitter would be lost through the receiver's matching network, severely impacting the transmitter's efficiency. Implementing a switch on the antenna side would require a very large transistor, because the impedance is relatively low and the available supply voltage limits the on resistance of the device. This would lead to a large loss in the matching network, which would act to harm the noise figure performance of the receiver's front end. In contrast, the impedance after the matching network is high, due to the large value of R_p from the FBAR, allowing the switch transistor to be much smaller while still leading to a low loss in the matching network. When the transceiver is configured for receive mode, this switch is turned on and presents a small load of roughly 50Ω, while the power amplifier is turned off and appears mostly as additional parasitic capacitance. During transmit mode, the switch is turned off, and the receiver's matching network appears only as an additional capacitance due to the C_1 variable capacitor, which can be switched to its lowest setting to limit its impact on the transmitter.

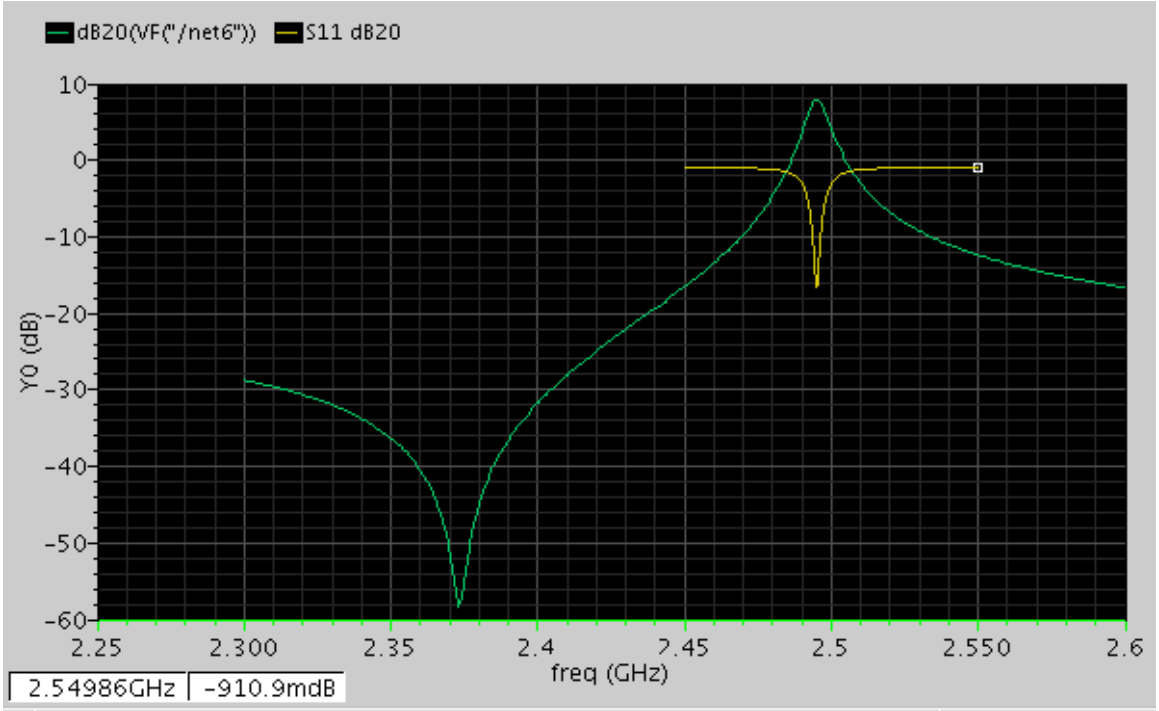

Figure 3.8: Simulated S11 and AC performance

In addition to providing a match to the antenna impedance, the high quality factor of the FBAR yields a narrow filter passband, helping to eliminate interference at the front end. Because the impedance of the FBAR is fairly high, it also provides roughly 8dB of passive voltage gain, relaxing some of the requirements on the active circuits that follow. 3.8 shows the simulated performance of the matching network when it is enabled, demonstrating that a good match is obtained along with its bandwidth performance.

A consequence of using the real impedance of the FBAR as as the load of the matching network is that it contributes 3dB to the noise figure of the front end. This is in addition to the noise that is contributed from the following stages, but is seen as a worthwhile trade-off in the context of an ultra-low power receiver due to the above advantages of the FBAR. In particular, the ability to raise the impedance to such a relatively high level allows the noise figure of the following stages to be substantially lower than if the impedance were close to the 50Ω standard antenna impedance.

3.2.2 Mixer and LO Generation

One of the most challenging aspects of the high level of integration targeted in this work is that the receiver front end must implement a large amount of gain, of more than 60dB, in order to bring the small signal at the input up to a level sufficient to be detected and processed, while also rejecting the supply ripple that results from the switching DC-

DC converter to ensure that it does not overload the analog front end. In many low-power radio designs, such as [39, 35], a single-ended front end is used to minimize the power consumption of the analog front end. If a single ended topology is used with a switching power supply producing ripple in the signal band, which difficult to avoid given the wide bandwidth required due to the ring oscillator LO, it would require an excessively large amount of decoupling capacitance to prevent the supply ripple from entering the signal path.

This can be avoided through the of fully differential analog circuitry in the baseband, which rejects the supply ripple from entering the baseband chain. This has an important consequence for the mixer operation, which now must be, at a minimum, a single-balanced topology. From [51], we can see that a single-balanced passive mixer, shown in Fig. 3.9, actually offers more gain than a double-balanced passive mixer, while also allowing the total amount of capacitance that needs to be switched to be reduced by half, allowing the power required for the mixer, calculated as

$$P_{mix} = \frac{1}{2} C_{mix} V_{DD}^2 f_{switch} \tag{3.3}$$

to be half as well.

Because the passive mixer essentially operates as a sample-and-hold operation, the gain of the mixer is not strongly dependent on its sizing, but the noise contribution of the mixer is essentially caused by the sampling of the voltage noise due to the on-resistance of the mixer's switches. Hence, the mixer's noise contribution can be minimized by making the size of the switches larger, thereby decreasing the on-resistance. For this design, a transistor width of 2.5μm was chosen for each device in the mixer, along with a minimum length of 60nm to minimize parasitic effects. The mixer was AC-coupled into the following gain stage through a 2pF capacitor on each side in order to allow the source and drain nodes of the mixer devices to remain at ground. The relatively small size of these switches will harm the linearity performance of the front-end, particularly for second-order nonlinearities; however, this is partially mitigated through the use of the very sharp filtering provided by the front-end.

In addition, we can calculate the approximate impedance that the passive mixer will add to the matching network circuit by viewing the circuit similarly to a switched capacitor network, translating the input capacitance of the following gain stage after the mixer into an effective resistance at the mixer input. This places an upper bound on the size of input transistor that is acceptable for the following stage to avoid adding too much resistance to the node, which would de-Q the resonator and harm the performance.

As discussed in Chapter 2, one way to improve the sensitivity of this architecture is to minimize the bandwidth required in the IF amplifiers, which requires either an accurate LO tuning. For this design, the target IF bandwidth was 20MHz, which meant that the LO must be tunable to within 20MHz of the carrier frequency to ensure that the signal can always be acquired. Through corner simulations, it was determined that for an oscillator with a target frequency of 2.4GHz, the impact of process variations would

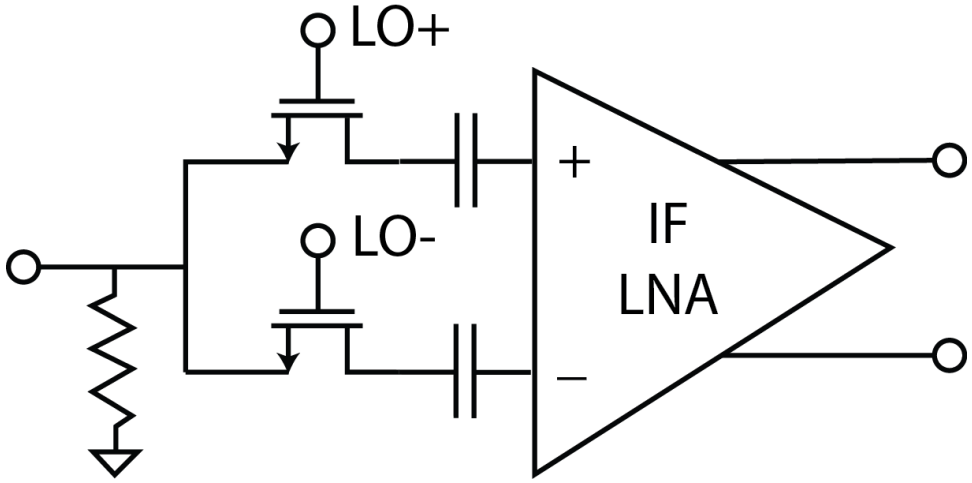

Figure 3.9: Passive mixer schematic

move the frequency by roughly a factor of 2 in each direction, so the oscillator's frequency across process variations for a nominal supply voltage would range from 1.2GHz up to roughly 5GHz, yielding nearly 4GHz of tuning range. Assuming a roughly linear tuning curve, this would then require

$$N_{bits} = \log_2\left(\frac{4GHz}{20MHz}\right) \approx 7.6bits \tag{3.4}$$

so it was chosen to implement 8 bits of tuning for the oscillator. The topology of the tuning followed the same strategy as [39], using a resistively current-starved topology with switchable resistors to tune the oscillator frequency as illustrated in 3.10, using a tunable resistor on both the V_{DD} and the ground sides of the oscillator to help ensure that the inverters stay roughly balanced in the middle of the supply range so that they can be further regenerated back to a full swing by the following buffer stages. The relatively large number of bits required for each of the resistor DACs leads to a somewhat large area in implementation, which is shown in Fig. 3.11 and occupies roughly 55µm by 135µm.

The single-balanced mixer topology requires a differential LO signal, which is achieved by using two 3-stage ring oscillator cores coupled lightly coupled together through transistors, as illustrated in Fig. 3.12. In effect, this coupling produces two oscillators who are locked to each other's frequency, but out of phase, ensuring that the oscillator produces the proper differential output to drive the mixer. The output of each side of the oscillator is then buffered through a series of two inverters with the full supply voltage rail to regenerate the output back to a full-swing logic signal and properly drive the mixer load.

One of the most critical challenges in implementing the LO core is ensuring that parasitic effects from the layout do not harm the operating frequency too severely. Because of the reduced operating voltage of the oscillator, each inverter has a very limited drive strength, which is then further reduced by the current starving resistors that provide the

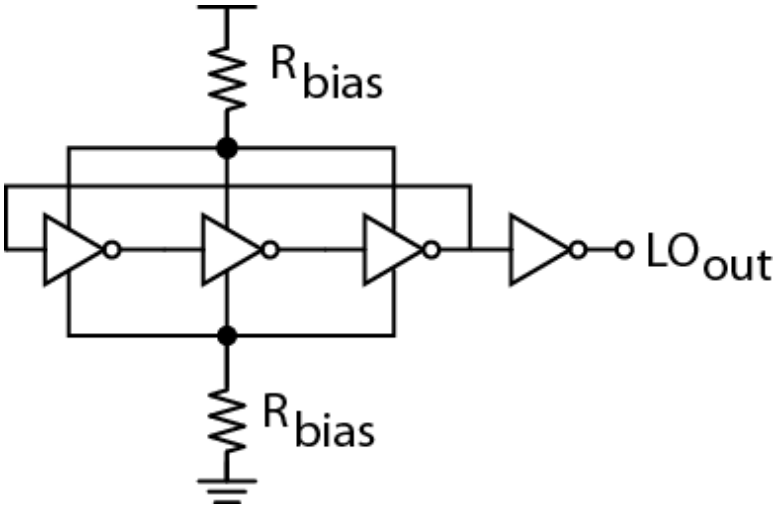

Figure 3.10: Ring oscillator for LO generation

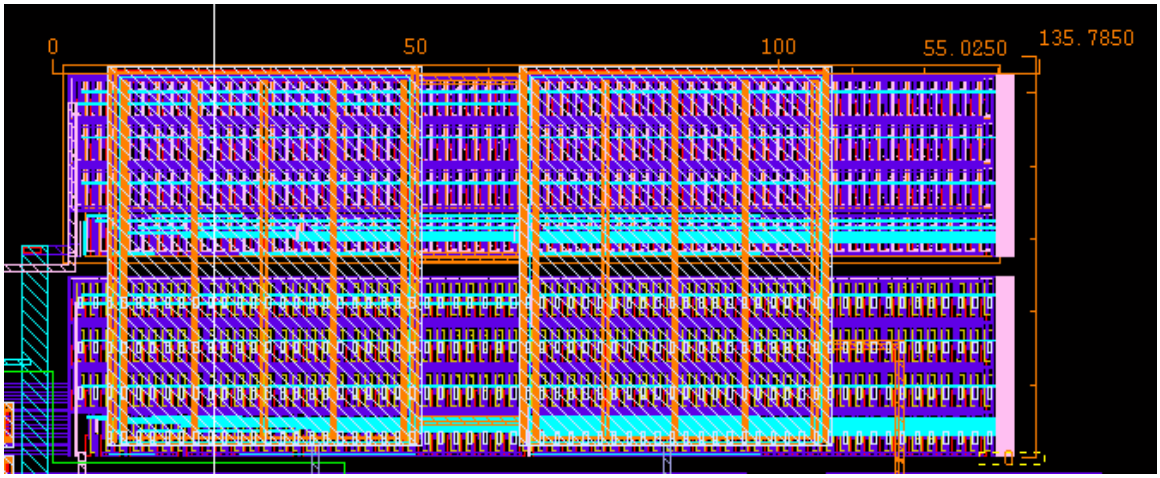

Figure 3.11: Resistor DACs used to tune the ring oscillator LO

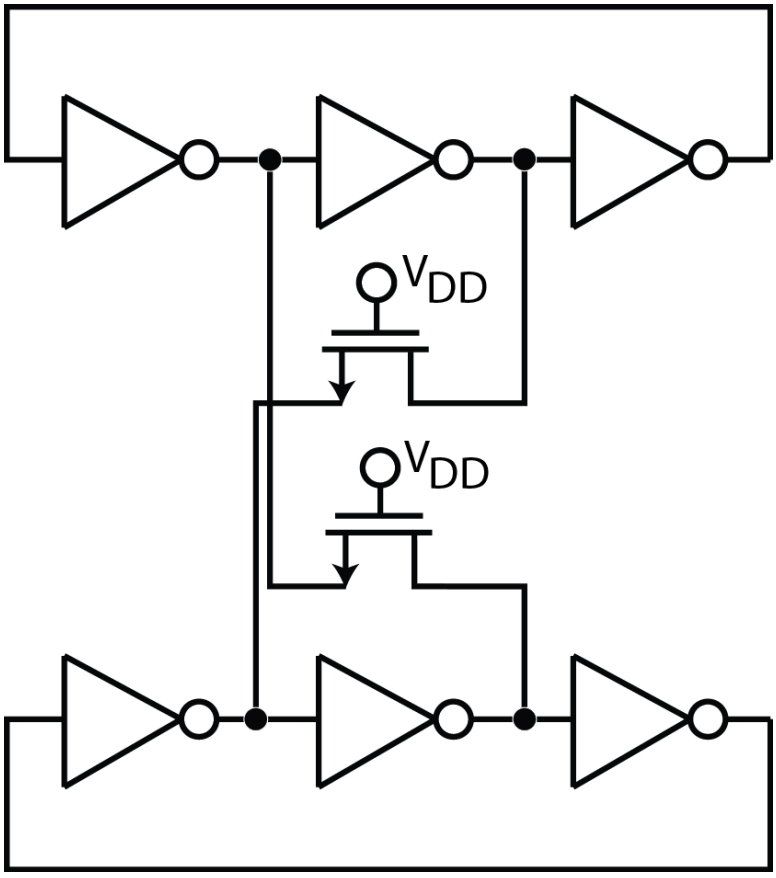

Figure 3.12: Ring oscillator core schematic, showing differential coupling

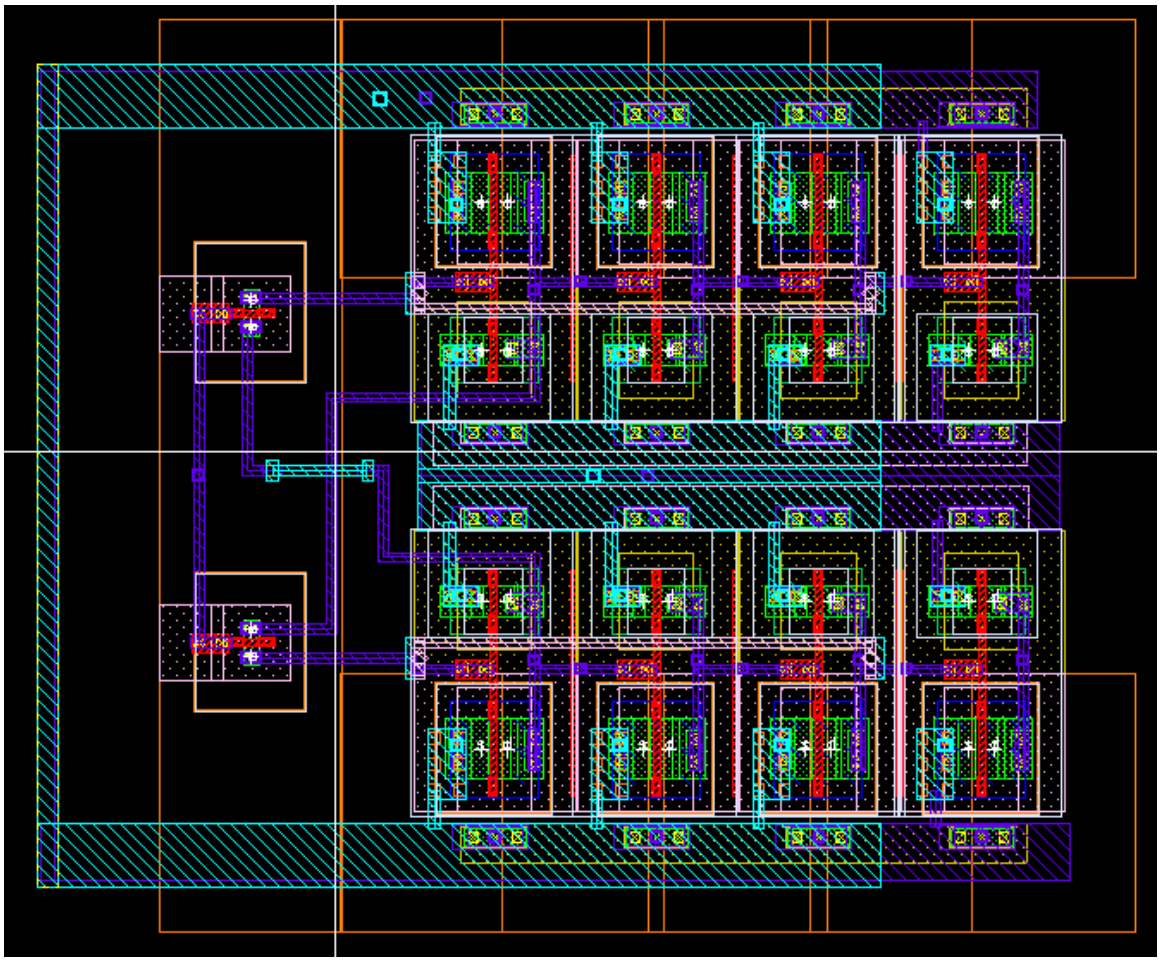

Figure 3.13: Oscillator core layout

tuning. Without carefully balancing between parasitic resistance due to wiring and vias and parasitic capacitance, it is very easy for the frequency of the oscillator to drop by more than a factor of two between schematic and layout design. The final layout, shown in Fig. 3.13, was extracted and simulated extensively to ensure that the oscillator would have sufficient tuning range to be able to meet the desired frequency even across process and temperature variations.

3.2.3 Gain Stages

The first amplifier stage following the passive mixer can be thought of as an IF LNA, because its noise contribution dominates the overall noise figure of the receiver. To maximize the headroom available for the amplifier, it uses a pseudo-differential topology loaded by an active inductor formed by a PMOS transistor that is diode connected through a resistor and bypassed through a capacitor at AC, as shown in Fig. 3.14. This has two

benefits: it gives a relatively high output resistance at AC frequencies due to a longer than minimum channel length for the devices, but it also gives a low impedance of roughly $1/g_m$ at DC, which gives a gain of approximately 1. This property means that the amplifier inherently avoids accumulated DC offset, preventing the output from being railed without needing any additional circuitry to perform offset correction. The IF LNA stage consumes approximately 10μW of power and provides a noise figure of 18dB.

The following stages are designed with a similar topology, but using a tail current source to ease biasing. In addition, they consume significantly less power than the first stage, at about 1μW per amplifier, because their noise contribution is much lower than the first stage, allowing the stages to be scaled for lower power. In addition to the power benefit, this also allows the devices to be smaller, which presents less of a capacitive load to the preceding stages, which allows the bandwidth to be extended.

The final stage needs to cope with the charge kickback due to the comparators used in the energy detector circuit, so instead of a diode connected output, it uses a simple differential pair loaded by resistors as shown in Fig. 3.15. This stage is sized for little gain to prevent DC offset, as well as to minimize the impact of the charge kickback. The other primary advantage of the resistively loaded device is that it allows more careful control over the DC levels at the output, ensuring that the levels are roughly correct for the comparator to operate properly.

Altogether, the IF gain lineup consists of four stages, as shown in Fig. 3.16. The amplifiers have a total bandwidth of roughly 25MHz including the layout parasitics while producing 60dB of gain. The total gain was designed to be about 60dB, and the overall noise figure was roughly 20dB. The first stage is the pseudo-differential IF LNA, followed by two identical gain stages with active inductor loads, and finally followed by a resistively loaded differential stage to drive the energy detector. The simulated performance of the IF amplifer chain is shown in Fig. 3.17. To provide a coarse gain control, there are switches that enable the signal path to bypass the final two gain stages, which will reduce the gain in the presence of large inputs. For a more robust design, one or more of the gain stages can be made to provide a variable gain by, for example, shunting some of the signal to the supply rather than to the output. With the simple modulation schemes used in this design, however, the signal gain does not need to be precisely controlled.

3.2.4 Energy Detection

The downside to this wideband IF gain, as discussed in Chapter 2, is that it allows a significantly larger noise bandwidth than is required to carry the data in the system. Because the sensitivity lost to this excess bandwidth can easily be many dB more than the noise figure of the front-end, it is worth exploring the method of performing the demodulation to see if there are areas in which some of the performance can be regained.

Because of the simplicity of the OOK modulation, all that is required to demodulate the signal is to perform energy detection, which determines the difference between a zero, when no energy is present, and a one, when the signal power is present. Traditionally, this

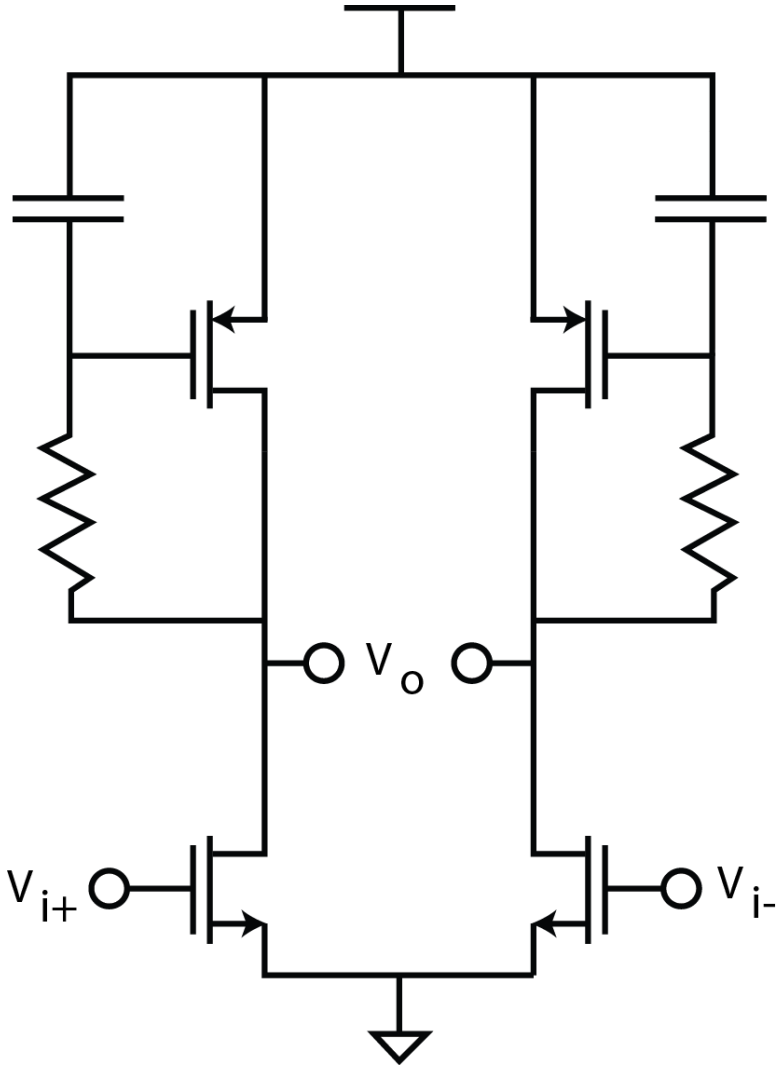

Figure 3.14: IF LNA schematic

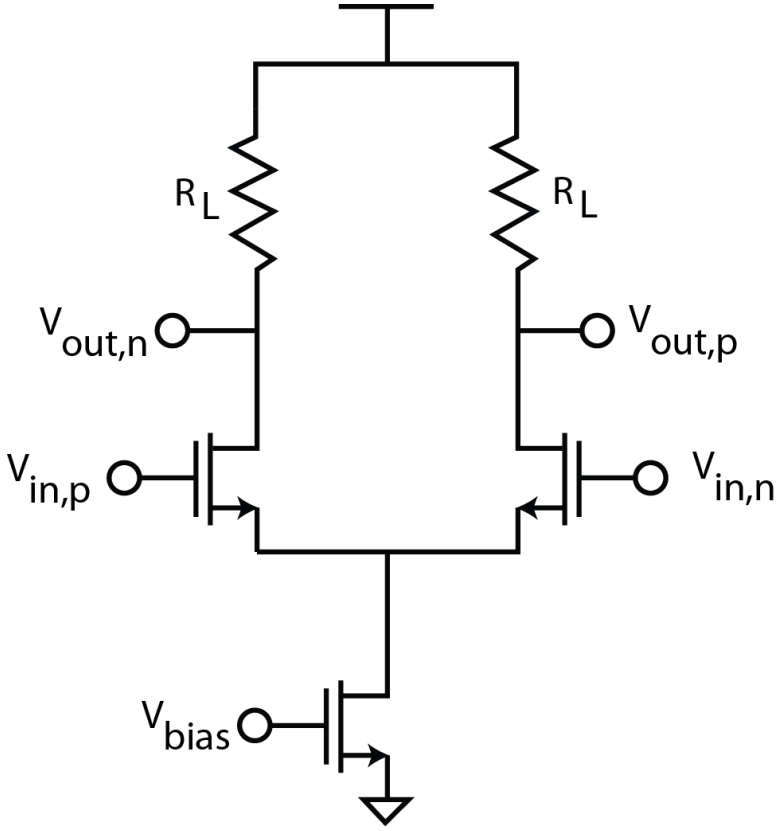

Figure 3.15: Final IF amplifier to interface with the energy detector

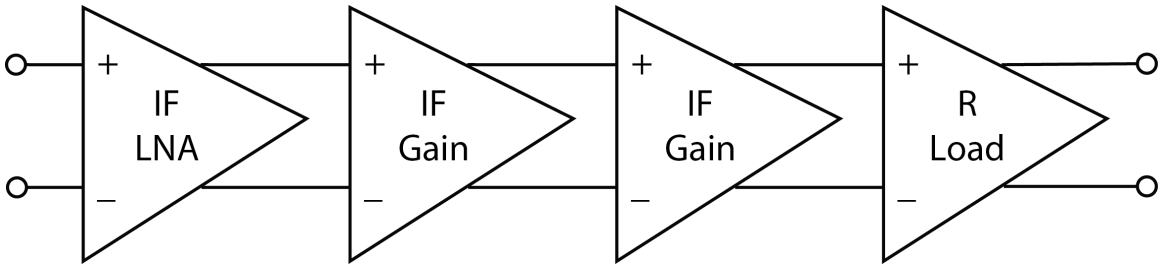

Figure 3.16: IF amplifier lineup diagram

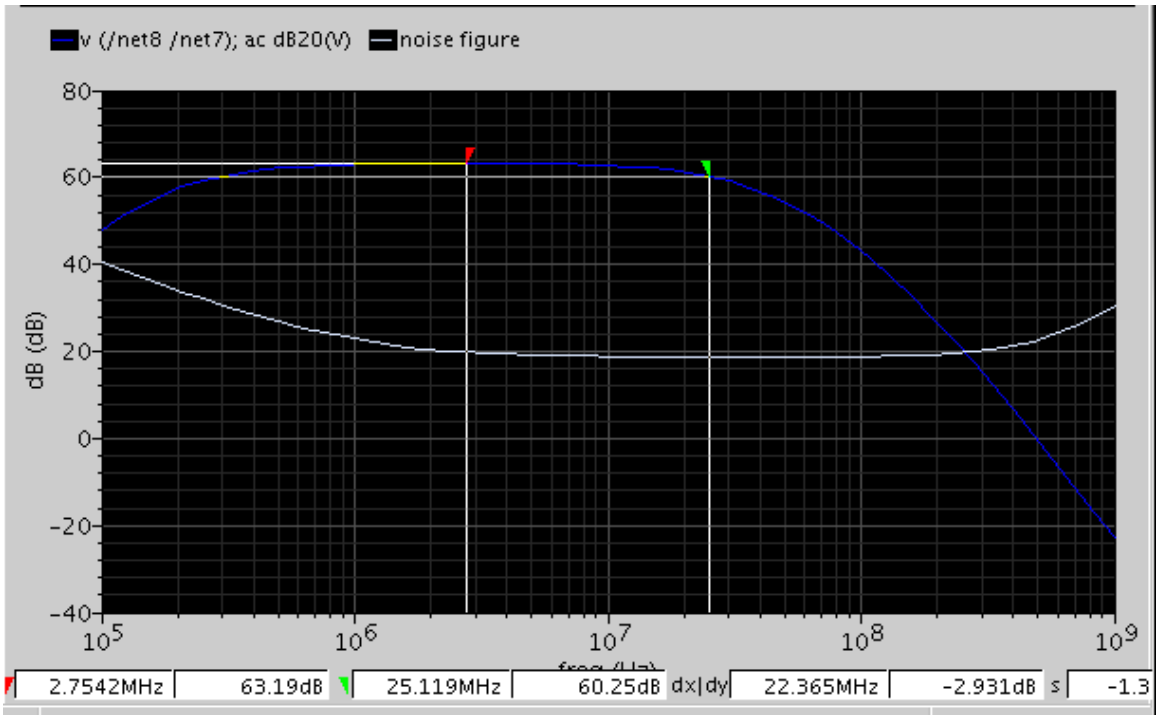

Figure 3.17: Simulated gain and noise figure for the IF amplifier chain

energy detection has been performed using a structure very similar to a diode AM radio, using the nonlinearity of the diode as an envelope detector. This presents a number of important drawbacks, however. A diode-based energy detector is inherently a wideband structure, which means that all the wideband noise present at the input will also appear after at its output. In addition, the gain of this nonlinear behavior falls rapidly as the signal level decreases, presenting a threshold effect that sharply limits the performance with small signals. The proposed energy detector addresses some of the shortcomings of these diode detectors by using the low power digital logic available in a modern, highly scaled CMOS process.

Rather than using an analog domain energy detector, the proposed technique directly samples the IF signal and performs the energy detection in the digital domain. This technique is inspired by two key insights about the signal at the IF band. First, detecting an OOK signal should in principle require only a single bit sampler, to determine whether energy is present. Second, because the IF bandwidth is much wider than the desired signal, it contains a significant amount of excess noise. This noise is uncorrelated with the signal, while the signal will provide a strong correlation with itself, so averaging can help to reduce the excess noise and detect the desired signal.

The digital energy detector, shown in Fig. 3.18, uses two comparators to perform an absolute value operation on the IF signal, detecting when the signal either rises above or below a given set of thresholds. Because the data uses OOK modulation, when the data

represents a zero, the signal only includes noise, so most of the samples will lie inside these thresholds, while a one will also include a sinusoid that rises above and below the thresholds. These samples are then fed into a shift register, which performs a majority vote operation on the detected bits, much like the use of a repetition code. We can estimate the BER of this signal in the same manner as a repetition code, as

$$P_e = \sum_{k=n/2}^{n} \varepsilon^n \cdot (1-\varepsilon)^{(n-k)} \tag{3.5}$$

where n is the number of effective bits in the repetition code, and ε is the error probability of each sample at a given SNR. Because the sampler is operating directly on a sinusoid, whose value is time varying, the effective number of bits is less than the actual number of bits used in the shift register by approximately 25%, due to some samples falling near the zero crossings.

Equation 3.5 shows that the error probability will fall rapidly as the sampling rate increases as long as the error probability of a single bit is less than 0.5. The reduction in error probability shows that the proposed detector is effectively able to remove some of the excess noise presented by the wideband IF, so the system will act like a digital filter, bringing the noise bandwidth closer to the desired bandwidth. Equivalently, this means the SNR_{min} at the energy detector has been reduced. It also shows that there is a limit to the amount of averaging that oversampling can perform, because it is not possible to gain any further improvement when the error probability of each sample is too high.

Increasing the sampling rate of the sampler provides more bits, and thus more possibility of correcting errors, but comes at the price of requiring a faster sampling clock and extra registers for storing the IF samples. This presents a trade-off for the receiver designer to choose how much improvement in SNR is achievable for a given amount of power. If a further reduction in the noise of the IF band is desired, it may be possible to implement a multi-bit sampling approach to implement a better digital filter; however, increasing the sampling resolution will result in a significantly more complex and power-hungry system, making it inappropriate for very energy-constrained architectures.

The digital technique provides a number of other advantages to the broader system in addition to reducing the noise impact of using the wideband IF. Because the sampler is operating at a much higher frequency than the desired data rate, synchronization is easier to implement since the resolution of timing information, such as the edges of data periods will be significantly increased over a traditional ADC sampling at the data rate frequency. Additionally, because there is no information present in the phase of the IF signal, the accuracy required in the frequency of the sampler is much reduced, and needs to only be fast enough to provide the desired data rate after the majority vote.

The digital energy detection system has three major variables for its design: the number of samples that are stored in the shift register, the sampling rate, and the threshold. These variables were studied more fully in [45, 43], and yielded the design chosen for this work. The shift register stores 100 samples, and is designed to be sampled at a mini-

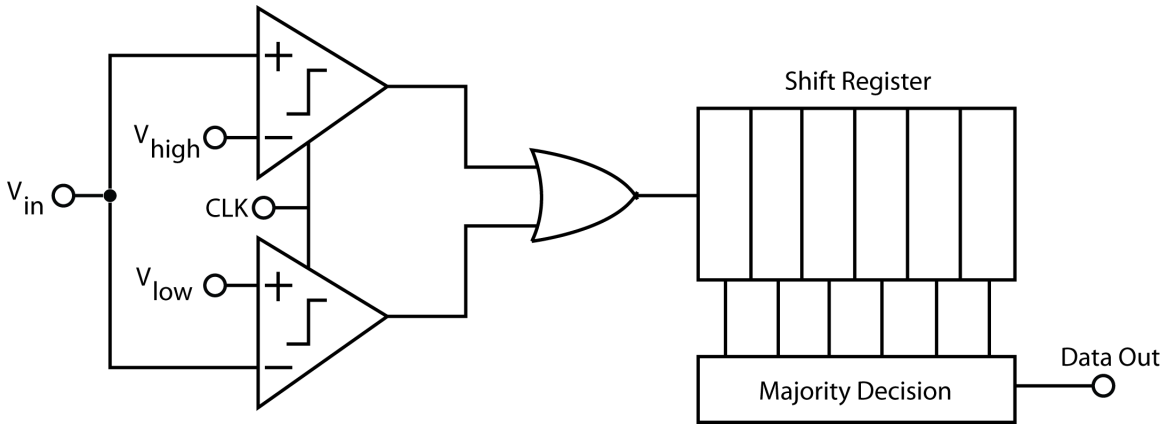

Figure 3.18: Digital Energy Detection

mum rate of at least 20MHz to ensure that it is able to properly demodulate a 100kbps Manchester-encoded signal as required by this work. The threshold is digitally tunable to allow the radio to find the optimal location to maximize its sensitivity. Fig. 3.19 shows the impact of varying this threshold on the bit error rate of the receiver, simulated using a MATLAB model. The maximum sensitivity is obtained by setting the threshold level at roughly the same level as the noise, and using the digital averaging to pull the signal back out.

The comparator used for each of the samplers is shown in Fig. 3.20. Its transistor sizes were chosen to bring the estimated offset due to mismatch to a small level of less than approximately 40mV through Monte Carlo simulations. The comparator itself was found to consume roughly 1μW of power in simulation, so that it does not consume a significant amount of power compared to the total system.

The sampling clock for the digital energy detector is generated using an on-chip relaxation oscillator, as shown in Fig. 3.21. This oscillator was chosen because it provides a relatively stable output frequency with a minimal amount of power consumption. The sampling frequency is required to be at least 20MHz to ensure proper demodulation with the 100 bit shift register, so this oscillator is designed for a nominal sampling rate of 30MHz, to allow for process variations in the resistor and capacitor, and consumes roughly 1μW. A Monte Carlo simulation confirms that the clock frequency will remain above the minimum, as shown in Fig. 3.22.

The shift register and majority vote logic were implemented by a colleague, Louis Alarcon, using Verilog code and implemented using standard logic synthesis tools. It was then laid out using standard cells and an automated place and route tool. The logic was simulated and found to consume an estimated 8μW. In addition, a synchronizer and digital baseband circuitry were implemented by another colleague, Wenting Zhou, also using Verilog and synthesized with standard cell logic. The synchronizer utilizes the oversampled nature of the energy detector to obtain the bit period for each bit with

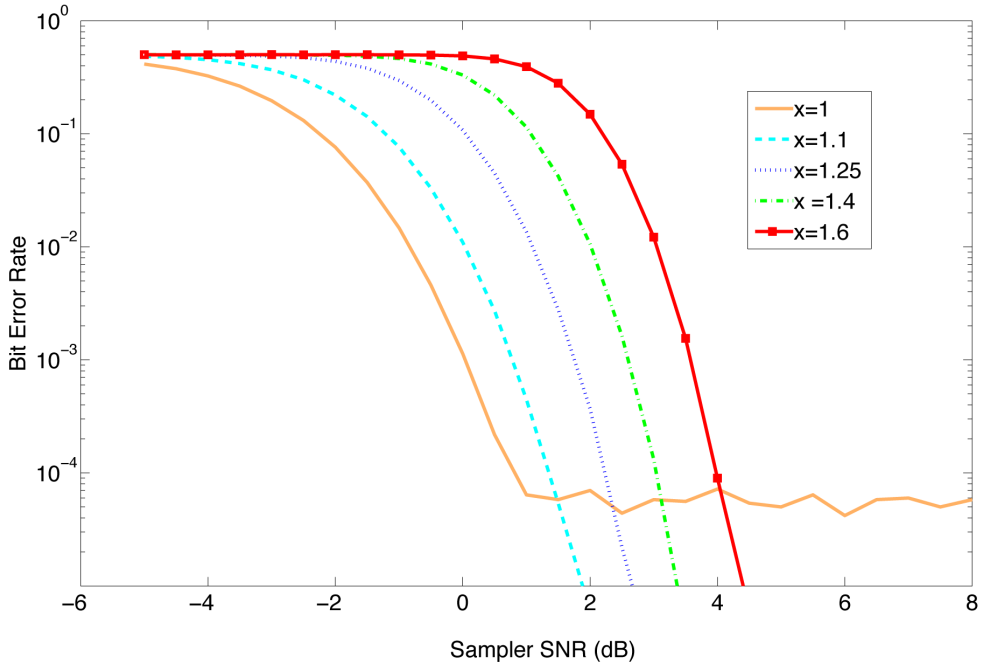

Figure 3.19: Required minimum SNR for demodulation

enough accuracy to ensure that a complete packet can be received correctly.

3.3 Sensitivity Analysis

We can now revisit the expected sensitivity level that this receiver should be able to achieve. Recall from Chapter 2 that we can estimate the minimum detectable signal as

$$P_{sens} = -174dBm + 10 \cdot \log(BW) + NF + SNR_{min} \tag{3.6}$$

where the -174dBm term is calculating the noise power at room temperature. Because the detection is being provided by the oversampling topology, we can split this calculation into two parts, first finding the expected noise power level, and then using the plots in Fig. 3.19 to gauge the SNR_{min} required.

Because the primary filtering in the receiver is performed at the antenna, all the additional noise added by the wideband IF reaches the input to the energy detector. This means that the effective BW for this receiver is 25MHz, despite the relatively low data rates that are used. In addition, the receiver achieves an NF of 20dB, so the total noise power reaching the energy detector is

$$P_{noise} = -174dBm + 10 \cdot \log(BW) + NF = -78dBm \tag{3.7}$$

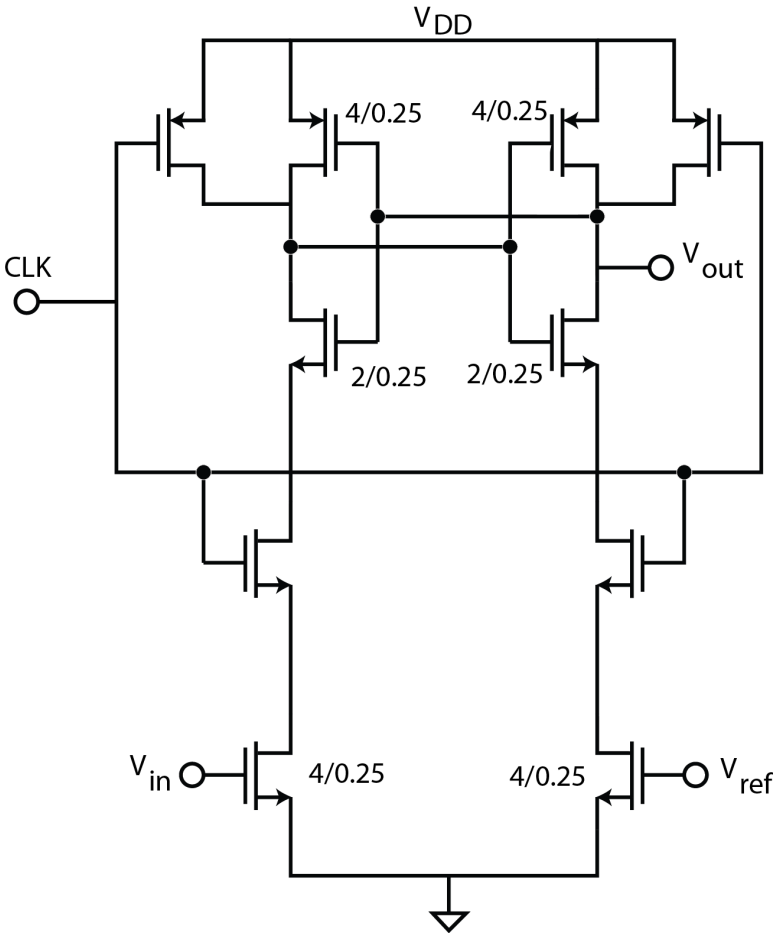

Figure 3.20: Comparator schematic

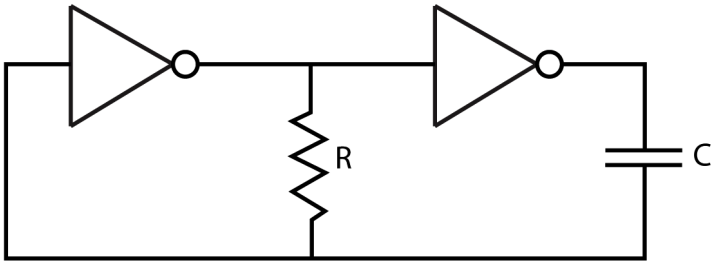

Figure 3.21: Relaxation oscillator for sampling clock

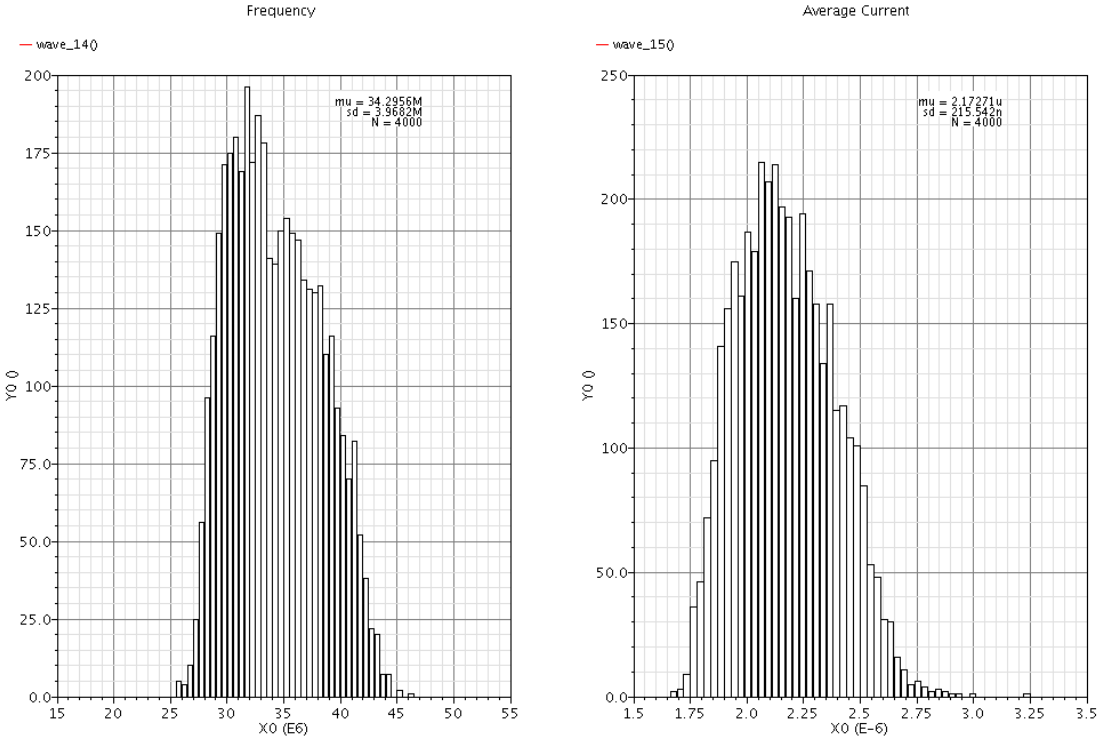

Figure 3.22: Monte Carlo simulation of the sampling clock

The minimum SNR required for reducing BER to below 10^{-3} varies depending on how aggressively the thresholds are set. When the threshold is set too low, the probability of receiving an incorrect "1" when a "0" was transmitted increases, ultimately limiting the achievable BER, even if the SNR presented to the sampler increases. For this reason, it is desirable to leave some additional margin in the choice of threshold, so a reasonable minimum SNR can be set at approximately 2dB, which yields at least a BER floor below 10^{-5}.

Putting this together with the total noise power, we can estimate that the receiver should be able to achieve a sensitivity of roughly -76dBm. This sensitivity would allow the receiver to compare favorably with those shown in Table 2.1, both in terms of its sensitivity as well as its total power consumption.

3.4 Measured Results

The complete receiver design was fabricated in a 65nm digital CMOS process with MIM capacitors, occupying a total area of 2mm by 2mm. The chip was assembled using a chip-on-board bonding to allow for easy characterization, and a photograph of the bonded chip is shown in Fig. 3.23. To ease testing, this board is implemented as a daughter board that interfaces with a larger board containing a series of linear regulators to generate the necessary supply voltages and ease the connectivity for testing. The chip includes two serial-to-parallel interfaces to allow the programming of the control bits for the different blocks, which are programmed using an Opal Kelly XEM3001 FPGA board. Because the receiver occupies a very small portion of the overall chip, its layout is shown in Fig. 3.24, and occupies an active area of 180µm by 110µm.

The power management circuit performs as expected from simulations, generating the correct output voltages for each supply voltage. Because of the nature of the switched capacitor DC-DC converter, the supply voltage contains a ripple as the control circuitry holds the voltage within 25mV of the nominal supply voltage. The linear regulator used to reduce the supple ripple on the sensitive analog components reduces this voltage ripple down to approximately 5mV of ripple. The measured waveforms showing the on chip supply voltage are shown in Fig. 3.25.

It was discovered while testing the receiver that the logic circuitry for the majority vote contains a flaw, which results in the output getting stuck high when random data is presented to the energy detector. When short sequences of data, such as expected from random noise presented to the receiver, are presented to the shift register and majority vote, this flaw is triggered and results in the receiver being unable to be properly tested with its digital output. To verify whether this bug is a result of a problem with the physical implementation of the design, the Verilog code was implemented using the Opal Kelly FPGA board, which exhibited the same behavior. The code was then modified to a simpler form of the energy detector, and the output of the chip's sampler was used to bypass the on-chip shift register and majority vote.

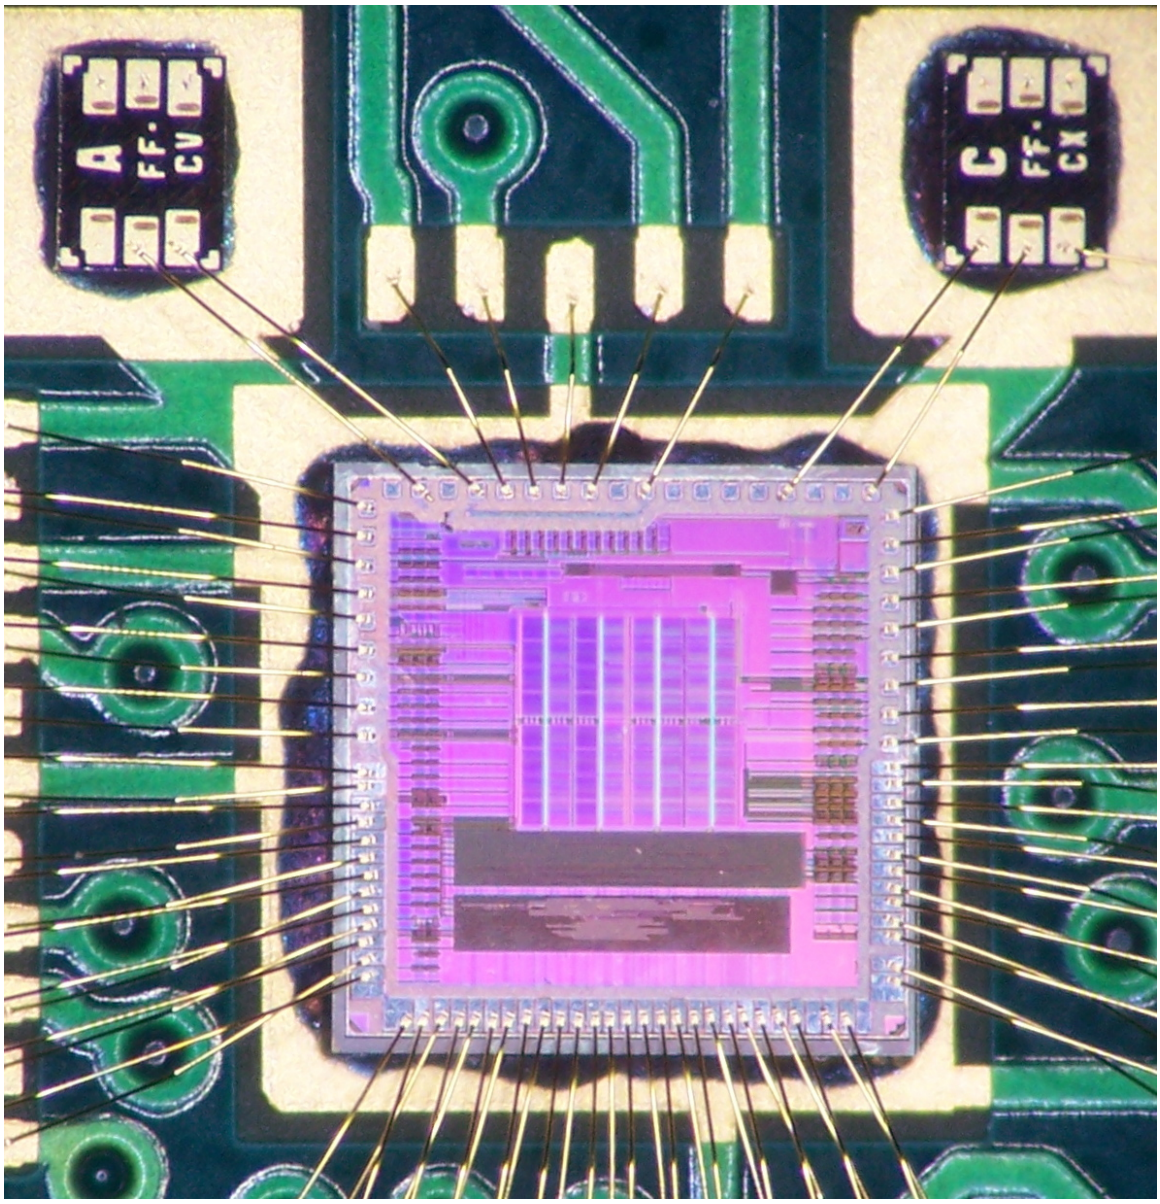

Figure 3.23: Die photo of bonded chip

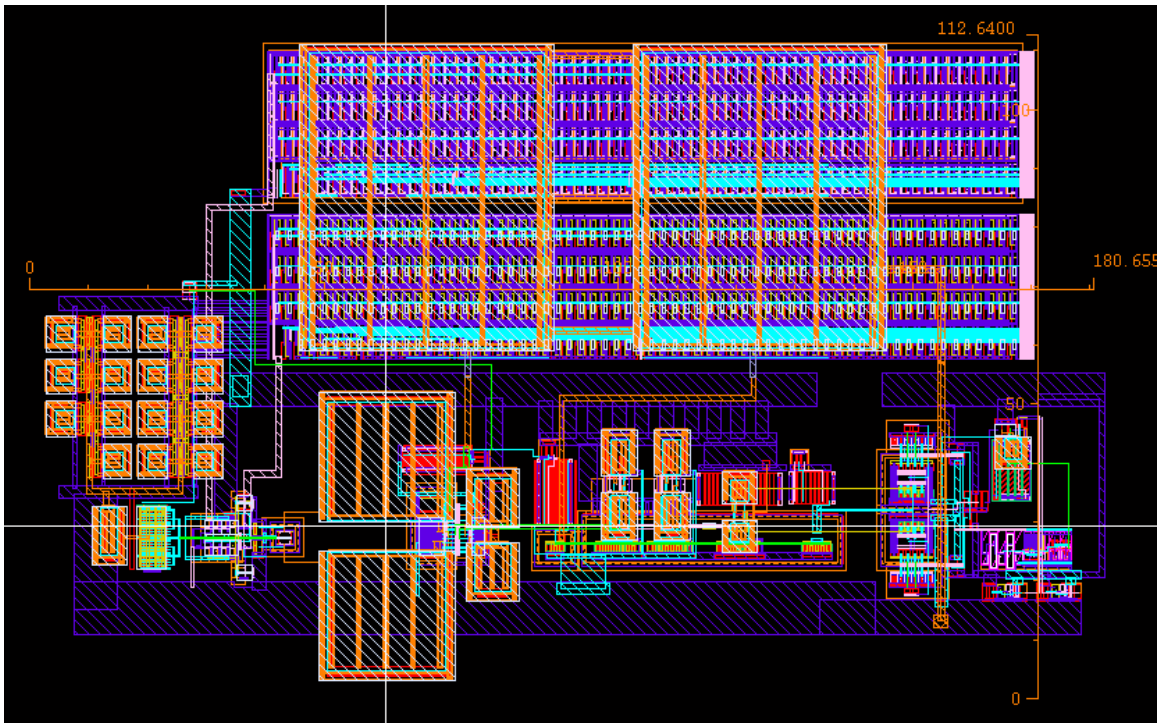

Figure 3.24: Layout of the receiver core

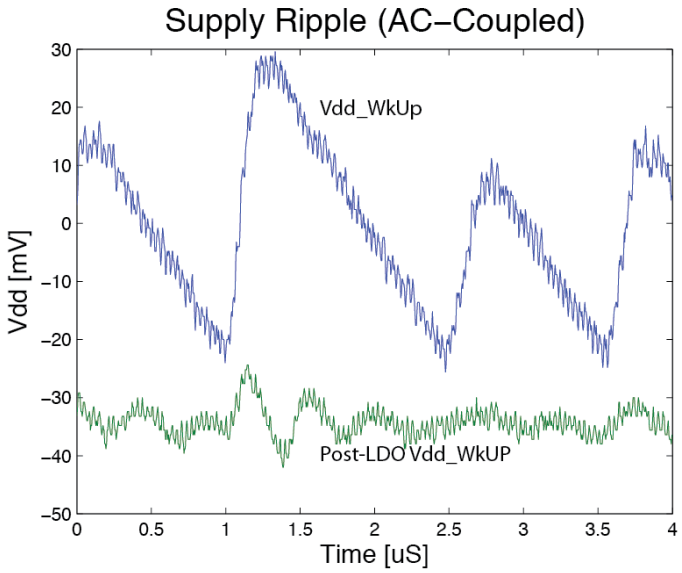

Figure 3.25: Voltage supply ripple due to integrated DC-DC converter

The chip exhibited another unexpected bug which limited its performance. While the calculations showed that the receiver should be capable of achieving approximately -76dBm sensitivity, it was discovered during the testing that the sampling clock along with the digital logic leaked into the supply for the LO. Due to the simple resistive current starving used to tune the ring oscillator, the LO is extremely sensitive to changes in the supply voltage, so this coupling resulted in a modulation of the LO generation, which then was translated into an interference at the baseband. This effect resulted in the sensitivity of the receiver being severely limited compared to the calculation, as shown in the measured result in Fig. 3.26. The receiver ultimately resulted in a measured sensitivity of -66dBm.

Unfortunately, the self-interference produced by the sampling clock prevented the full characterization and quantification of the performance of the oversampled energy detector. Due to the high level of integration, there was not an available output pin to directly measure the analog signals reaching the energy detector, and the single-bit nature of the energy detector limits the ability to precisely determine the effectiveness of the technique. Qualitatively, it could be seen on the oscilloscope traces that the sampler was producing a very noisy output directly, with many incorrect samples as expected by the MATLAB simulations, which were then corrected through the majority vote logic.

The analog front-end circuitry of the receiver, consisting of the mixer, LO generation, and IF amplifiers consumed a total of 48μW, including its biasing. The sampler, majority vote, synchronizer, and protocol processor consume an additional 61μW, resulting in a total power consumption of 109μW during the receive mode. The analog power consumption of this work compares favorably with the power required by [39, 21], while the digital power is providing additional functionality not supported by these references.

A summary of the specifications are given in Table 3.2. While the primary focus of this chapter has been on the receiver design and performance, it is important to note that the complete chip integrates substantially more functionality and provides a complete platform for a low power wireless sensor node, including both the self-sustaining energy management, a full transceiver, and a complete digital baseband implementing the protocol.

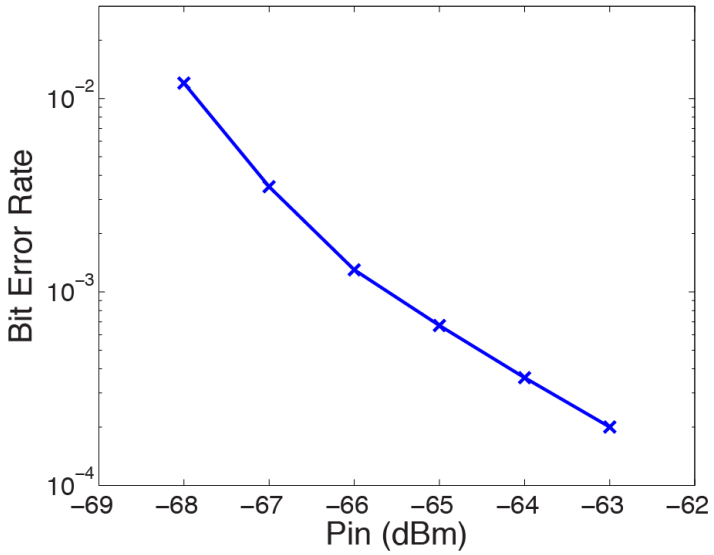

Figure 3.26: Measured BER versus input power

Battery Voltage	1.2-1.8V
Die Size	2mm x 2mm
Process Node	65nm CMOS
Core Supply Voltage	550mV
RF Supply Voltage	500mV
RX Sensitivity	-66dBm
Carrier Frequency	2.47GHz
Modulation Scheme	On-Off Keying
Data Rate	100kbps Manchester

Table 3.2: Specifications summary for the Active RFID design

Chapter 4

Example Project 2: Wake-up Receiver

While the performance the receiver in the Active RFID project discussed in Chapter 3 provided is somewhat lower than expected, it did yield some important insights into the design of low power radios. One of the biggest lessons was that the design and implementation of the LO generation is critically important and requires care both in its design and its physical implementation. The use of the resistive tuning circuit left the oscillator extremely vulnerable to interference from elsewhere on the chip, because any signal that manages to couple into the tuning network directly impacts the frequency of the oscillator. Part of this issue is attributable to a lack of sufficient local decoupling capacitors to isolate the design from the other circuits, but it also reinforces the importance of a clean environment on chip, so minimizing the amount of large, fast signals nearby to the receiver can help to improve its performance.

The single-bit sampler used for energy detection in the receiver has some clear limitations. Because of the very limited resolution provided by the sampler, the detector is only able to remove approximately 10dB of excess noise from the baseband, while the total amount of excess noise was shown in Chapter 2 to easily exceed 20dB, especially given the power restrictions on the LO generation. By using other techniques to reduce this excess noise, additional sensitivity performance can be found from the receiver front-end.

An additional issue that arose in the design relates more to the transmitter and voltage converter than the receiver itself. OOK-modulated data inherently involves very large swings in the current draw requirements on the transmitter, which places strict requirements on the DC-DC converter. This is caused by the need for the entire power amplifier to switch, at the data rate, between very low current (roughly 200μW in the Active RFID design) and the high current output when a sinusoid is being sent out (roughly 6mW, assuming a 30% efficient transmitter). Given that the data rate was designed to be a 200kbps raw switch rate, the current draw in transmit mode switches as quickly as once every 5μs across more than an order of magnitude, so the voltage regulation must be able to react quickly and accurately, adding to the difficulty in the design of the converter and regulator. The design of the receiver for Active RFID, however, can only function with an OOK signal because the demodulation scheme was specifically designed for it. Therefore, a design that enables the use of constant-envelope modulation schemes can be beneficial from an overall system design standpoint.

4.1 Design Goals

Previous ultra-low power designs, including those destined for wakeup radios, have focused on both detecting the presence of a signal as well as demodulating an addressing signal directly to determine if the wakeup signal is intended for the designated radio. This means that the radio needs to achieve dual purposes: it needs to detect that there is energy in the signal band, but it also needs to determine the address of the intended target of the wakeup signal. These two goals have typically been wrapped together into a single receiver design, but fundamentally they have different needs and can be addressed

separately.

Compared to the work presented in Chapter 4, the Active RFID receiver, this work aims to improve the sensitivity of the receiver significantly as well as making the receiver more flexible so that it can accommodate a wider range of applications. Even if the design had worked perfectly in the Active RFID receiver, the simulated sensitivity limit was only about -80dBm, which compares favorably to other recent low power radio designs as showed in Table 2.1, but is still well below the typical sensitivities found in current commercial wireless receiver designs. This work aims to both improve the sensitivity of the receiver by roughly 15dB as well as further optimize and reduce the power consumption required of the radio design. In addition, it is more desirable for receivers to operate with a constant-envelope modulation than OOK, because it makes the design of the transmitter easier and is more immune to noise and interference issues than amplitude modulation. The receiver should be able to be configured for either energy detection or constant-envelope demodulation, depending on the needs of the system.

For a wakeup radio, one of the most important specifications is the active power draw, because this power is constantly consumed by the radio platform. Therefore, this project aims to produce an active power in the main energy detection mode to less than 50μW, and preferably well below this figure to account for the additional power draw for demodulation and synchronization with the wakeup signal. In addition to the energy detection mode, it is also desired to achieve a data rate of up to 1Mbps in order to make the receiver more usable in a wide range of applications, such as traditional 802.15.4-based wireless sensor networks.

In contrast with typical radio specifications, a wakeup radio needs to be measured by different metrics compared to typical radios. In particular, rather than focusing on the bit error rate, it instead becomes important to focus on the false alarm rate and the missed detection rate. A false alarm occurs when the radio tries to wake up the rest of the system, but it turns out that there was not actually an incoming signal. This results in a waste of power for the radio, which can deplete its energy storage faster than necessary. In contrast, a missed detection occurs when there is an incoming signal for the radio, but the wakeup radio fails to find it and the system stays asleep. In this scenario, instead of wasting power, the wakeup radio results in an increased latency in establishing communications with the node. At its most basic, a wakeup radio simply needs to listen to the channel for energy, and upon detecting a sufficient amount of energy, begins the process of waking up the system. This naive approach of simply using energy detection is extremely vulnerable to interference, however, and prevents the wakeup system from using addressing to target specific nodes in a wireless network. This can result in a significant waste of energy, as the entire neighborhood of wireless signals is awoken each time there is any energy in the channel.

Previous attempts at designing wakeup-style radios, such as [38, 39, 44, 21], have been designed with a simple demodulation scheme in order to reduce this effect and allow addressing. The strict limits on the active power of the radios means that the demodulator has extremely stringent power requirements, which has led these radios to only use OOK

modulation, with all the problems discussed at the beginning of this chapter for the broader system. Furthermore, the focus on receiving an address to wake up the node means that in a practical system, the synchronizer and sampler must be running at all times, which will result in an additional power draw that is typically not discussed in works focusing on the front-end of the radio. It is worth stepping back to a higher level and considering if receiving data is the best tool.

Any wakeup operation essentially consists of three steps:

1. Is there any energy at all present in the frequency band of my interest?
2. If there is energy, is it a type of signal that I can understand (i.e. demodulate)?
3. If I can demodulate the signal, is it actually trying to talk to me?

Previous works have typically focused on answering all of these three questions simultaneously, by listening for a modulated signal, but they do not necessarily need to be approached in this way. This work focuses on minimizing the power required to operate a wakeup radio by separating these questions into distinct steps and optimizing the receiver separately for each portion.

4.2 Circuit Implementation

4.2.1 Overview and Receiver Partitioning

This wakeup receiver architecture uses two main partitions, illustrated in the block diagram in Fig. 4.1. The first, including the FBAR matching network and filtering, the passive mixer, the ring oscillator, IF gain stages, and energy detection, constitute the always-on portion of the receiver that performs the basic energy detection looking for the presence of a signal in the designated frequency band. The output of the energy detector is fed into an integrator, which then compares the integrated energy to a threshold to determine if there is energy present in the system. If energy is detected, the receiver is reconfigured to operate in a constant-envelope demodulation mode by turning on an FBAR oscillator, which provides a very accurate and low noise reference for the main LO, and the output of the gain stages is fed instead to a demodulator for detecting BPSK signals.

This two-stage approach to processing a wakeup command has two important results. Firstly, because the always-on wakeup portion no longer needs to scan for modulated signals, the design of these blocks can focus instead on achieving the lowest power energy detection possible. Secondly, because the demodulation does not need to occur all the time, the power consumed by the demodulator is duty cycled, and so they have the ability to include much less stringent controls on their power, such as being able to afford the use of an FBAR oscillator that consumes more power, but provides substantially better performance. In addition, the proposed BPSK demodulator that will be discussed later

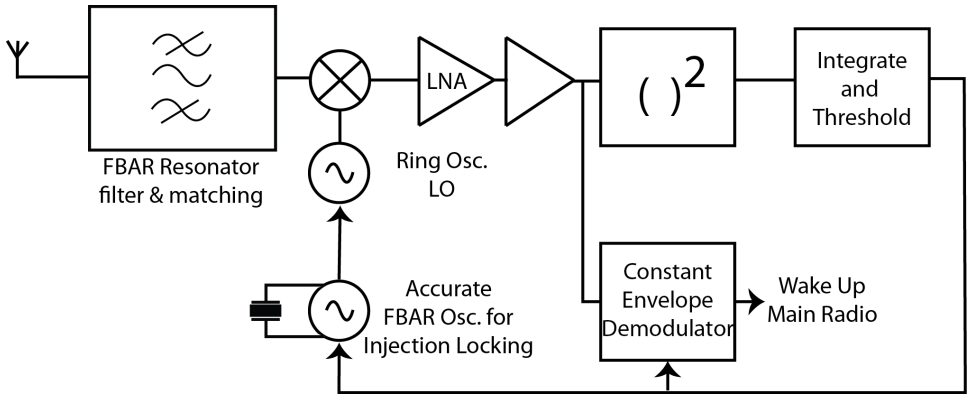

Figure 4.1: Block diagram of a complete wakeup receiver system

in this chapter is not the only possible demodulation that can be used in this place. A more traditional ADC and DSP combination can be used instead, enabling the front-end to work with a wider variety of possible standards.

The primary goals for this radio are to further improve the sensitivity compared to previous radios, targeting a -95dBm wakeup power, while maintaining total power consumption lower than 50μW. The receiver should also be highly reconfigurable and able to support both OOK modulation and energy detection as well as constant envelope modulations. As with the previous work, there should be no external components besides the MEMS resonator. The receiver is implemented using a standard 65nm bulk digital CMOS process with no additional process options.

4.2.2 Mixer and LO Generation

Based on their performance as discussed in Chapter 3, it was decided to again use a passive mixer at the front end of this receiver. Because the mixer is being driven by a switching source, in this case a CMOS inverter, one way to minimize the power consumption is to consider the choice of operating frequency. The power required from the LO chain and mixer can be found as

$$P_{mix} = \frac{1}{2} C_{total} V_{DD}^2 f_{rf} \tag{4.1}$$

where C_{total} represents the total capacitive load between the buffers and the mixer input and f_{rf} is the frequency of the switching, which, due to the low-IF nature of the architecture, will be approximately equal to the RF carrier frequency. This then demonstrates that there are two primary variables that can be controlled for minimizing the power consumption of the mixer: the size of the mixer, which will impact the amount of capacitance that needs to be switched, and the RF frequency, which will directly scale the power required for the mixer.

The choice of RF carrier frequency is typically dictated by regulations that determine what applications can be used in each frequency band, as well as the amount of regulatory burden required. The 2.4GHz is a well-known and highly popular choice because it is available around the world for unlicensed use, and so it tends to be crowded with signals for such devices as wireless LAN (802.11), Bluetooth, Zigbee (802.15.4), and others like cordless phones. There are other frequency bands that are available for unlicensed use, however, such as at 433MHz, 900MHz, and 5GHz. For the purposes of building a low power radio, the lower frequencies are best because they directly reduce the amount of power required for the mixer and LO generation. For this application, the 900MHz band was chosen as the carrier frequency, but because the only aspects of the radio that are specific to the operating frequency are the matching network and the ring oscillator, the design does not require many changes to operate in a different band.

The schematic of the mixer's operation is shown in Fig. 4.2. In the previous design, the capacitors used for AC coupling between the mixer and the baseband amplifier chain had to be very large in order to minimize the impact of the corner frequency on the noise performance of the mixer. For this design, the AC coupling was moved to the LO side of the mixer, allowing both smaller capacitors to be used and giving the ability to tune the bias point of the passive mixers, which improves noise performance slightly. To optimize the size of the mixer, simulations were run sweeping the number of 0.5µm fingers used in the mixer, with the result shown in Fig. 4.3. At small mixer sizes, the mixer itself dominates the noise figure of the front end, while at larger sizes, the IF amplifiers dominate. Because the mixer is driven by switching digital logic, the amount of capacitance directly determines the amount of power consumption from the LO. The mixer's size was set to 2.5µm width and minimum length for each side of the passive mixer, which provided a good trade-off between minimizing the noise contributions while keeping the power consumption low.

One of the major issues with using a ring oscillator to generate the LO for the mixer is that ring oscillators' frequencies vary significantly across process, supply voltage, and temperature. The resistor-based tuning structure used in the Active RFID was especially sensitive to each of these variables, and the supply voltage problem specifically ended up costing a significant amount of performance due to the issue of the supply being coupled with the digital clock signals. The purpose of the resistor string at the supply and ground nodes of the oscillator in Chapter 3 is to limit the amount of current available to the inverters, which then limits the rate at which its capacitors can be charged and discharged, which ultimately limits its speed, as can be seen by the equation

$$T_{osc} = \frac{N \cdot C_{eff} \cdot V_{DD}}{I_{bias}} \tag{4.2}$$

calculating the period of the oscillation. Resistors do not make the most effective means of controlling currents, however, because the current they allow depends on the voltage across them, so the speed depends heavily on the supply voltage as well. In contrast, a ring oscillator controlled by a real current source will be more stable across voltage. A

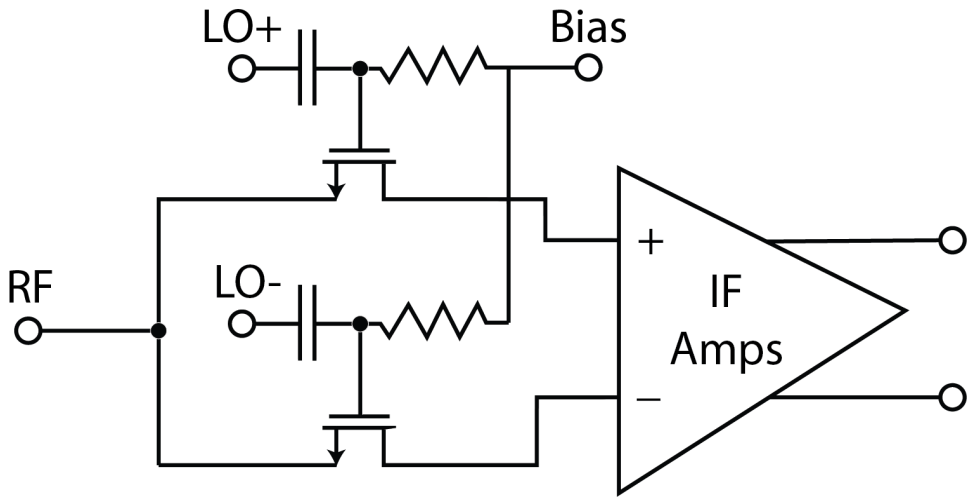

Figure 4.2: Mixer schematic

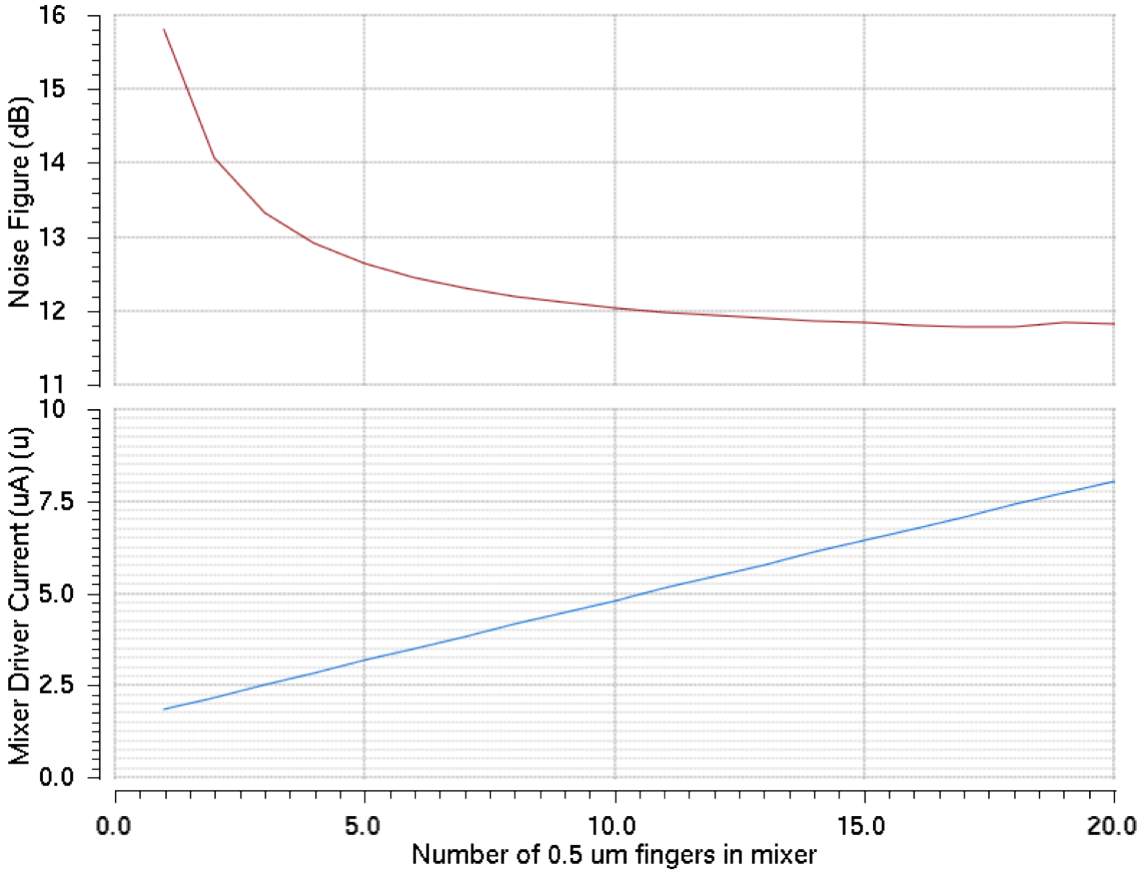

Figure 4.3: Noise figure and current versus the mixer size

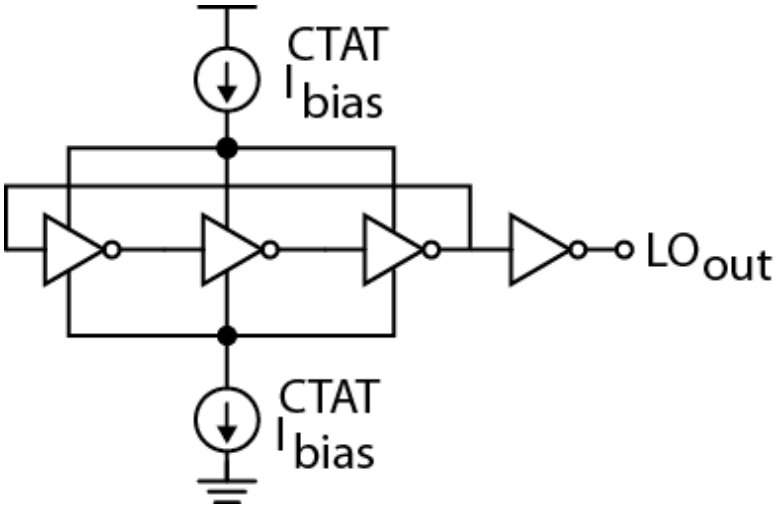

Figure 4.4: Temperature compensated ring oscillator LO

circuit version of this approach is shown in Fig. 4.4, which uses MOSFET current sources to control the ring oscillator instead.

Although a current source can reduce the impact of voltage variations, a ring oscillator is still very susceptible to process and temperature variations. Process variations can be dealt with by use of tuning circuits that operate at the receiver's start-up, measuring the operating frequency and comparing to a reference. Because process variations tend to be mostly stable across time, this tuning need not be performed very often, so its power overhead is negligible. Temperature variations are more problematic, as even in an indoor environment, temperature will vary by a few degrees through the day and can easily see sudden changes if, for example, sunlight periodically shines on a sensor or a heating and air conditioning vent is directed towards it. Simulations indicate that the frequency of oscillation increases roughly linearly with temperature, suggesting an approach to remove some of the temperature dependence of the LO. We can see from Eq. 4.2 that the period of the oscillator is inversely proportional to its bias current, meaning that its frequency is directly proportional to the bias current. By using a bias current generator complementary to temperature as the source of the oscillator's bias, the drift in frequency with temperature may be partially cancelled, in a manner similar to that used by a bandgap voltage reference to create a stable voltage.

The circuit for generating the bias for the oscillator is shown in Fig. 4.5, and operates by using a subthreshold MOSFET to emulate a classic bipolar and CMOS complementary-to-absolute-temperature (CTAT) reference, in which the amount of current falls as temperature increases. By controlling the value of the resistor, the level of the bias current can be adjusted and tuned, both for its absolute value as well as for its flatness across temperature. Fig. 4.6 shows the performance of the oscillator across temperature when using the simple temperature compensation scheme, indicating that it varies by only about

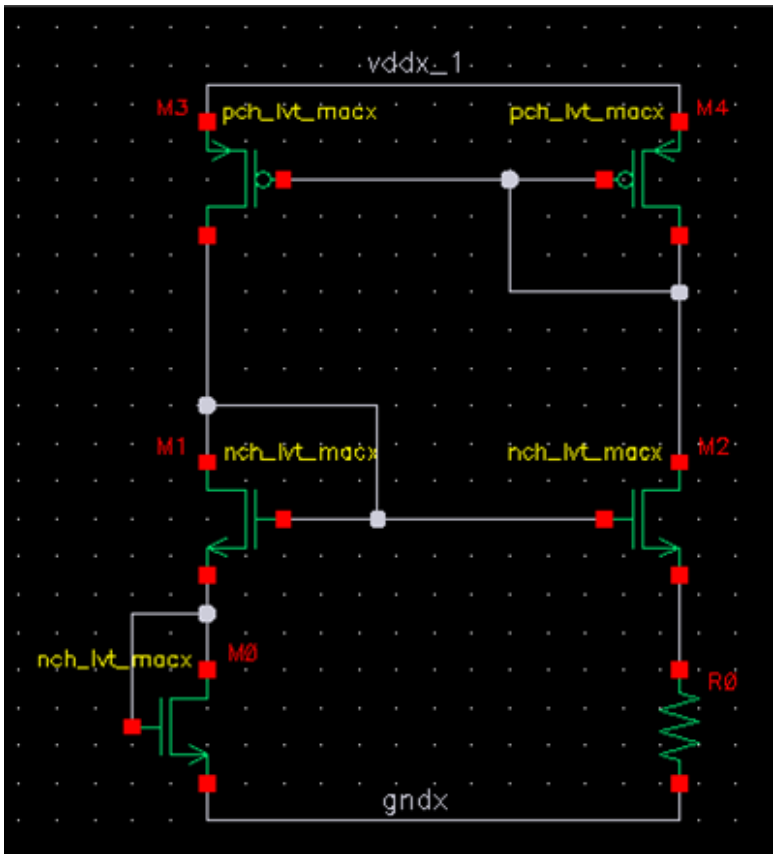

Figure 4.5: CTAT current reference used to bias the ring oscillator

18MHz across the range from 0-75^{o}C. The overhead power required for this compensation is only about 2µW, so it does not cost a lot of additional power to bring a substantial improvement in stability.

The input matching network that drives the mixer uses a similar topology as that from the Active RFID radio, except that it is designed for a 50Ω match this time to ease the testability of the system.

4.2.3 IF amplifiers

In the work in Chapter 3, each of the gain stages consisted of essentially a transconductance stage loaded by an active inductor; that is, the load had a low impedance at low frequencies, and a high impedance at high frequencies, but still functioned only as a load. This represents a wasted opportunity, because the PMOS load devices are effectively not contributing anything but noise to the performance of the amplifier, which reduces the noise figure. In contrast, by connecting the gate devices of the PMOS to the input of the amplifier, as shown in Fig. 4.7, the amplifier effectively becomes two amplifiers in parallel,

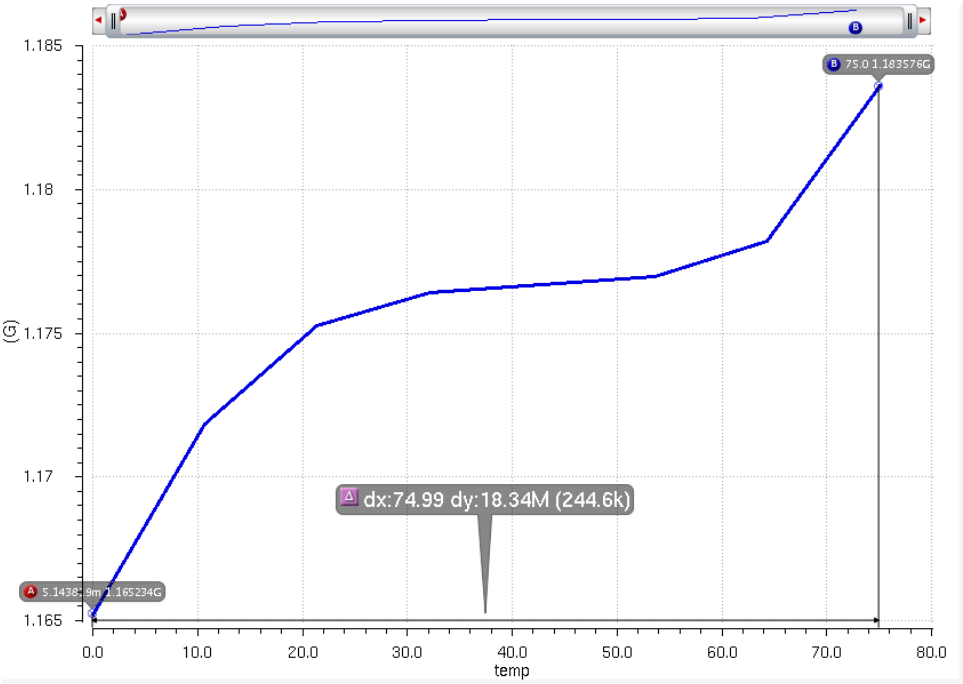

Figure 4.6: Simulation of the LO frequency across temperature

sharing a single current path [15]. This current-sharing technique roughly doubles the g_m of the amplifier, which in turn reduces the noise figure and increases its gain.

As before, we can think of this amplifier as acting like an IF LNA, because the noise performance of the entire amplifier chain is heavily determined by the performance of the first stage. Fig. 4.8 shows the noise figure of the front end and the IF LNA current consumption as the bias current is swept, with the mixer sized at 2.5µm per device as discussed in the previous section. At very low current levels, the noise figure falls rapidly as the bias current is increased, but eventually the noise figure begins to flatten out. This provides a trade-off between the current consumption of the amplifier and its noise performance, and suggests that the optimal point to operate is roughly in the range of 30-50µA, when the overall noise figure is approximately 12dB. Beyond this point, the noise again begins to be dominated by the passive mixer front end and the matching network, making an increase in the power consumption of the baseband unattractive as a means of improving performance.

To provide the remainder of the gain, an additional amplifier using a similar topology to the IF LNA is used, shown in Fig. 4.9. Because the preceding gain is quite large (~30dB), the noise performance of this additional gain stage is not critical, so its bias current can be significantly reduced. This lower current level allows for more headroom in the amplifier, so to simplify the biasing, a tail current source is used in addition to a similar current-reuse topology driving each of the transistors. In contrast to the IF LNA, the current re-use in this stage is primarily intended as a means to increase

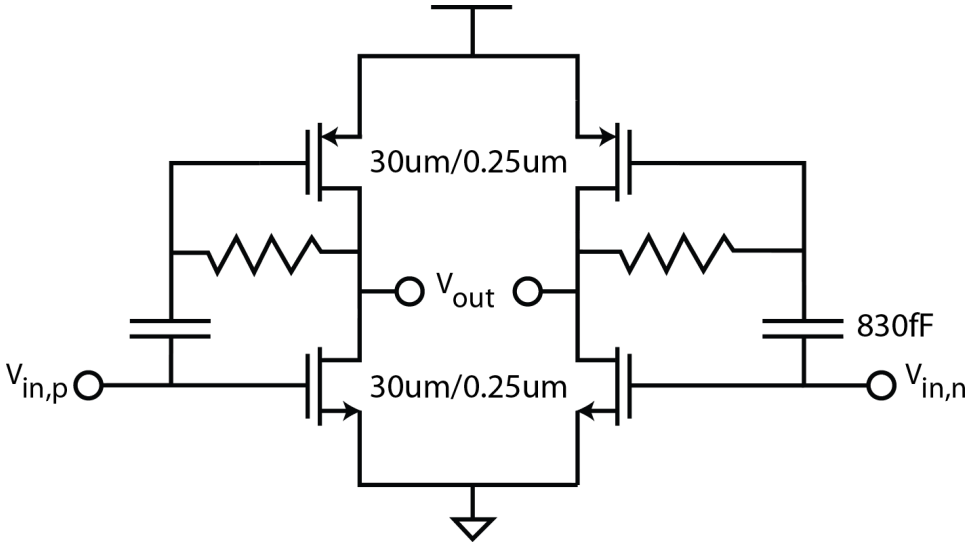

Figure 4.7: IF LNA schematic diagram

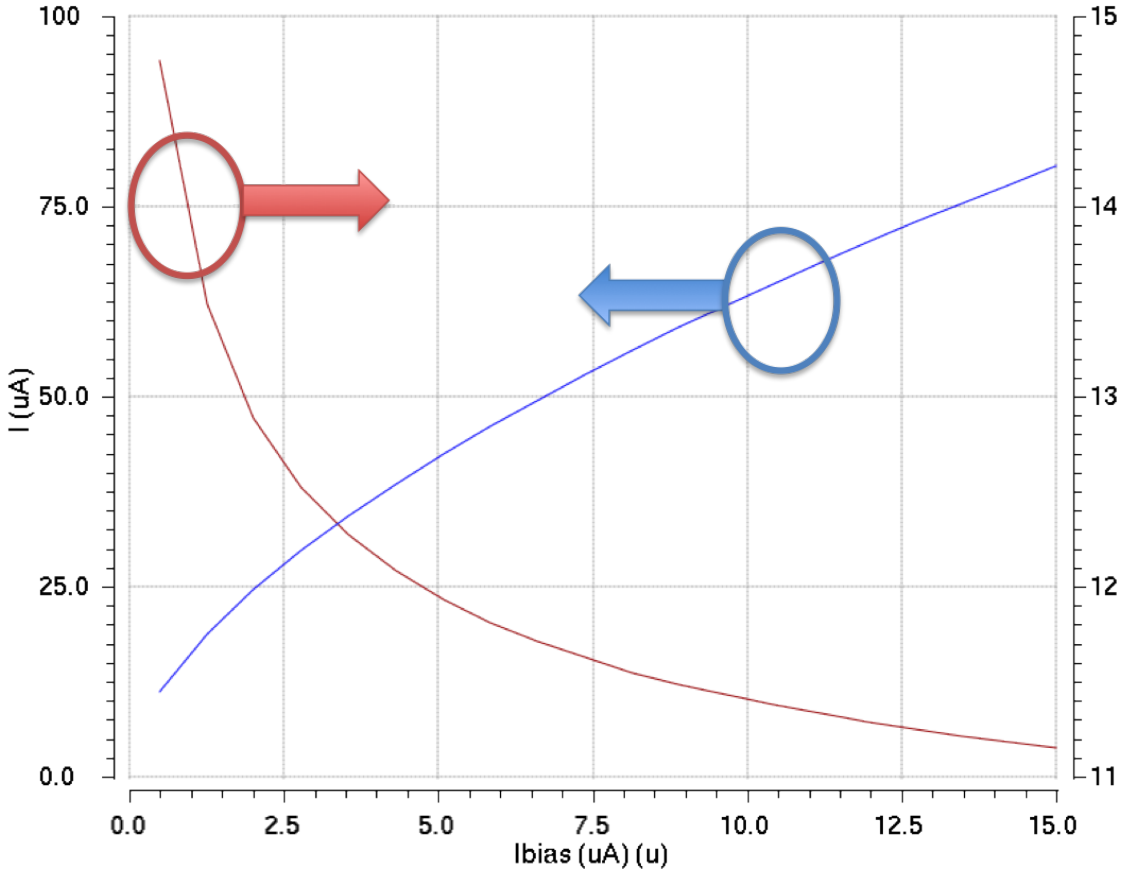

Figure 4.8: LNA performance across bias points

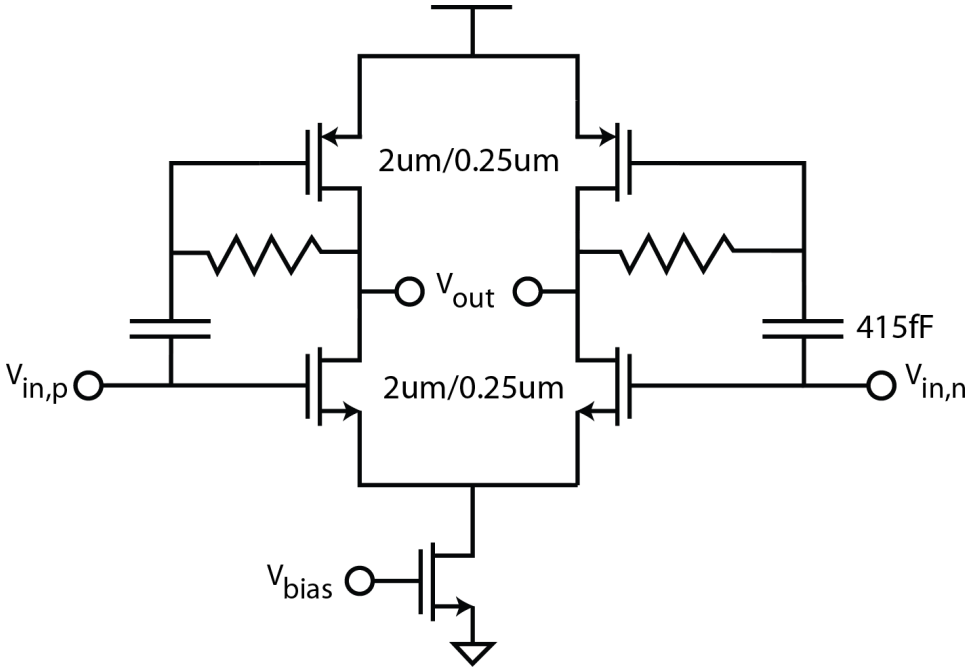

Figure 4.9: IF gain stage

the gain available from a given bias current, allowing the overall current consumption to be minimized. The final amplifier stage consumes only about 2μW and contributes approximately 30dB of gain.

The combined gain of the passive mixer, IF LNA, and additional gain stage is shown in Fig. 4.10, using post-layout extracted circuits for a more accurate simulation. The total gain of the front-end is 62dB, and the bandwidth is 27MHz. The spot noise figure of the entire front end, including the mixer and gain stages, is shown in Fig. 4.11. For this simulation, the matching network and high-Q FBAR was replaced by a simple wideband resistive source and matched load, so this simulation represents the minimum noise figure attainable throughout the IF band. The noise figure in the IF band is 12dB throughout the band in this configuration. Because the front end matching network introduces some additional filtering of the real impedance at the input to the mixer, the average effective noise figure is somewhat lower, at roughly 11dB.

4.2.4 Analog Energy Detection

In Chapter 3, the OOK data was demodulated using an oversampled digital approach, which was designed to reduce the impact of the excess wideband noise from the receiver's topology. While this technique was effective in reducing the impact of this noise, the maximum effective noise reduction possible from the demodulator was limited to around 10dB. In addition, operating the sampler and the digital logic involves introducing switch-

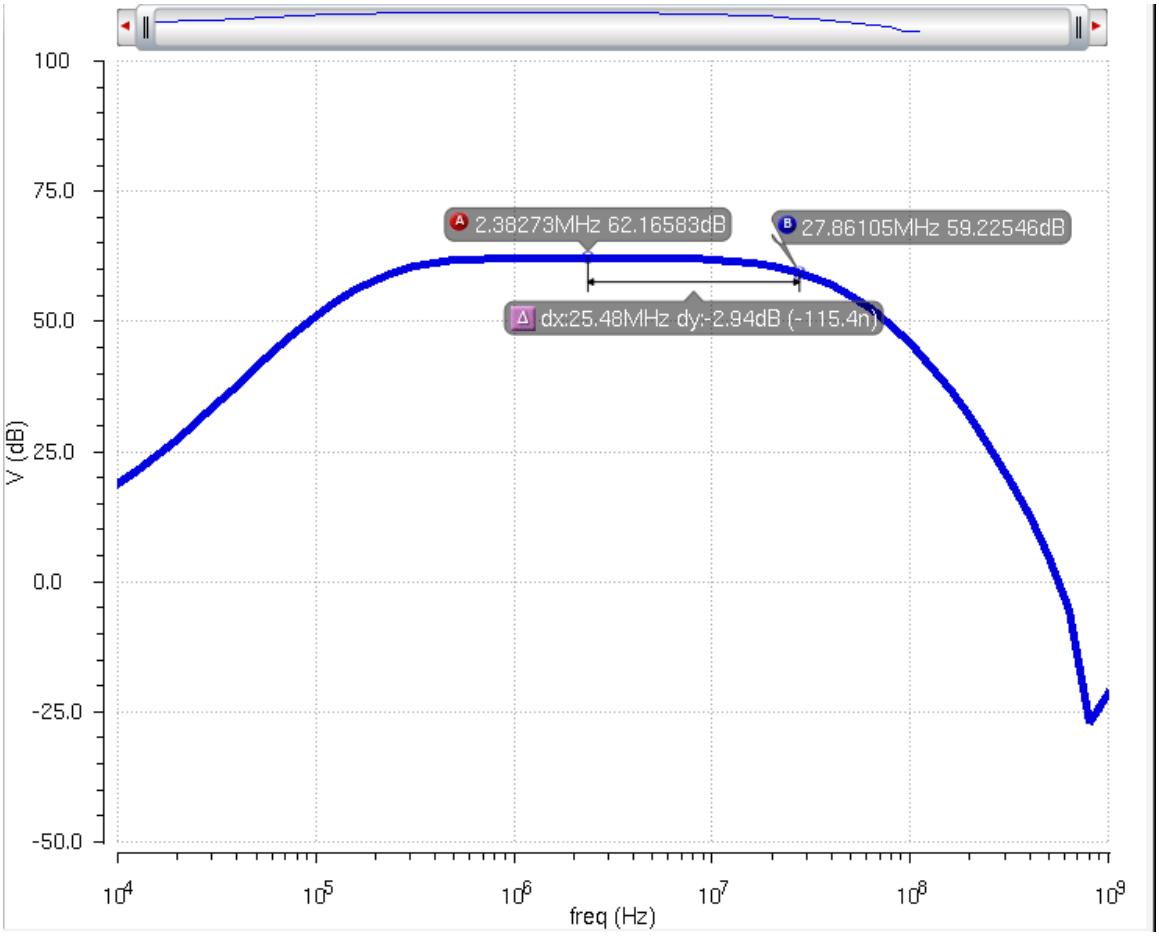

Figure 4.10: Simulated gain and bandwidth of IF amplifiers

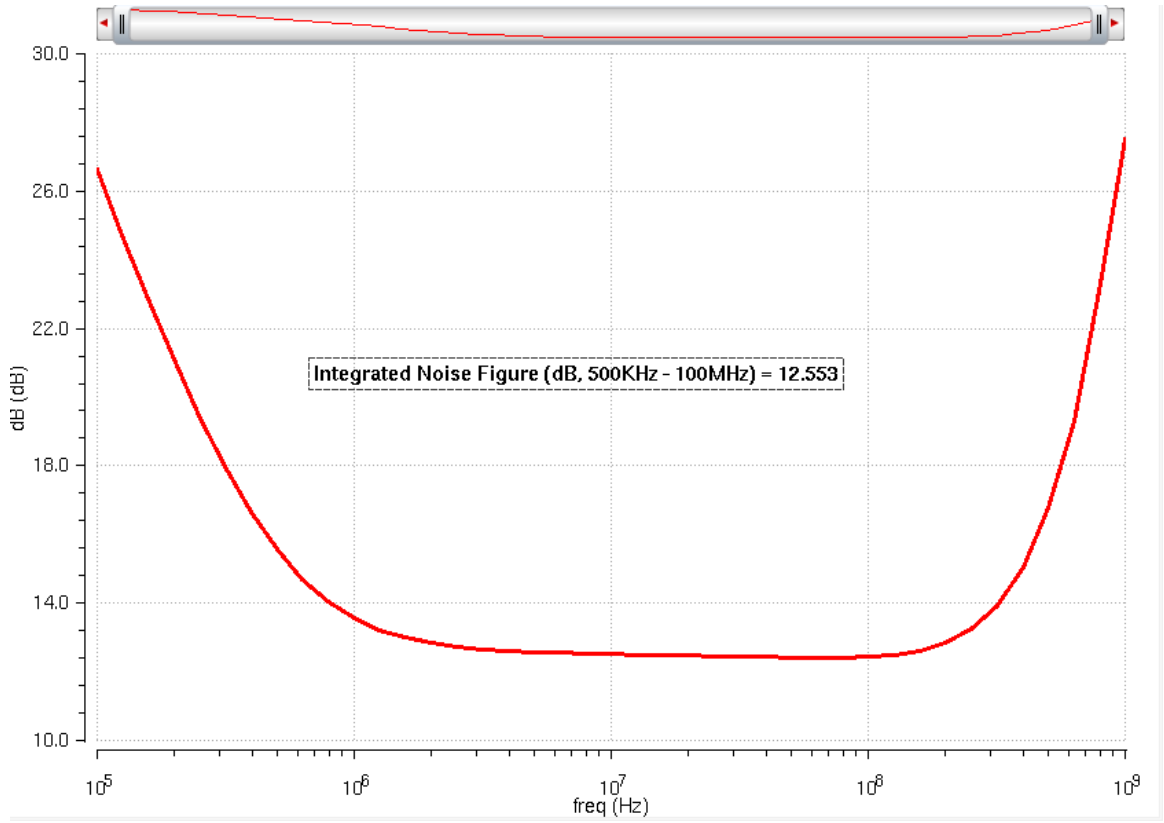

Figure 4.11: Baseband noise figure, assuming a wideband impedance match instead of a narrowband resonator match

ing logic nearby to the sensitive receiver front end and LO generation, which resulted in degradation to the sensitivity of the receiver. This degradation could have been reduced or eliminated by more careful layout and isolation between the blocks, but it is also desirable to remove the possibility of the problem.

In contrast, the demodulation used by [39, 21] relied on simple, analog envelope detectors provided by nonlinear circuits. These envelope detectors do not require a large amount of power, but because of their nonlinear nature, they present some downsides that must be considered. Because both take advantage of the 2nd order nonlinear behavior of MOS transistors, the gain they are able to achieve is quite limited. This means that they require a substantial amount of additional gain following the energy detection, increasing the power consumption of the demodulation as well as potentially introducing additional noise.

Instead of relying on the nonlinearity of transistors, this work explicitly implements a squaring function by using a mixer and driving both the input and LO paths together, as in Fig. 4.12. To minimize the power required, the squarer uses a passive mixer, so it requires no DC current to operate, and is implemented with a fully balanced topology to minimize the IF signal leaking to the output. A diagram of the mixer, without its biasing details, is shown in Fig. 4.13. The gain of the energy detector as a function of the input signal level is shown in Fig. 4.14, which shows that the gain begins to saturate with large input signal levels.

One drawback of using an energy detector is that it is an inherently wideband circuit, which takes all energy from the entire IF band and mixes it together. As discussed previously, this results in a substantial amount of excess noise. In addition, the mixer produces both the desired downconverted signal as well as an upconverted signal at twice the frequency of the IF. This means that a low-pass filter is required to remove the undesired high frequency term. This low-pass filter also acts to remove some of the excess noise generated by the wideband IF, providing an improvement to the sensitivity similarly to the action of the oversampling energy detector of Chapter 3.

An alternative to the use of a low-pass filter following the energy detector is to use a strategy called integrate-and-dump, in which the output is integrated onto a capacitor and periodically discharged or reset. This technique has been commonly used in the field of optical communications systems, and has been shown to implement a matched filter for an OOK signal following an energy detector [9]. This concept was also analyzed by a colleague, Wenting Zhou, for this purpose and was shown to provide a roughly 5-10dB improvement in the sensitivity of the receiver compared to simply low-pass filtering [52].

The circuit implementation for both the low-pass filter operation and the integrate-and-dump operation was shared, as illustrated in Fig. 4.15, using a simple differential-pair amplifier with a capacitor across the output terminals and a switch to reset the capacitor. When operating as a low-pass filter, the switch is simply kept open, yielding a low pass filter with a corner frequency of approximately 20kHz. To operate as an integrator, this switch is periodically closed, shorting the capacitor and resetting the value stored on it. When the switch is opened, the capacitor is gradually charged as integration is performed,

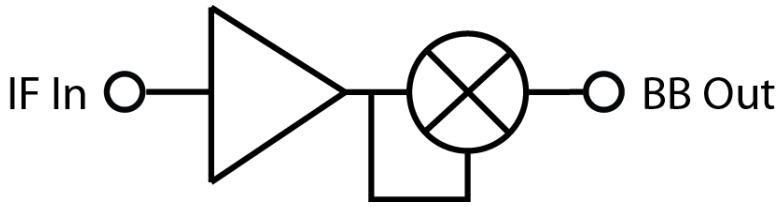

Figure 4.12: Self-mixing energy detector block diagram

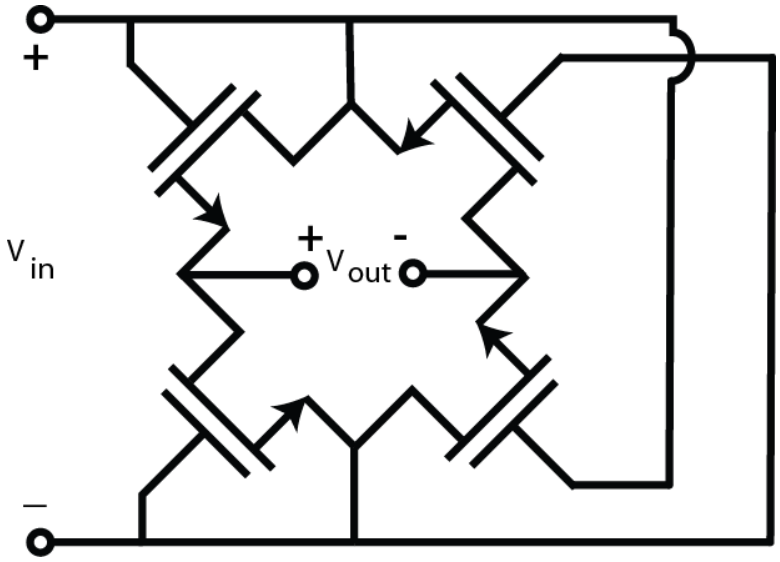

Figure 4.13: Energy detector schematic. Note that biasing details are not shown for simplicity, but all transistors are biased to be roughly at their threshold to maximize output.

and this output can then be thresholded to demodulate the desired signal. The common-mode bias level is set by a simple common-mode feedback circuit using resistors (not shown in the schematic for simplicity) to measure the common-mode voltage level.

The main storage and filtering capacitor for this block was implemented on-chip using an MOM finger capacitor, to minimize the number of external components. This capacitor's size limits the corner frequency achievable, along with the maximum rate at which the integrator can be clocked. For this implementation, the total capacitance included was approximately 25pF, which could be supplemented with off-chip PCB capacitors if a lower frequency corner is desired. The AC response of the filter is shown in Fig. 4.16.

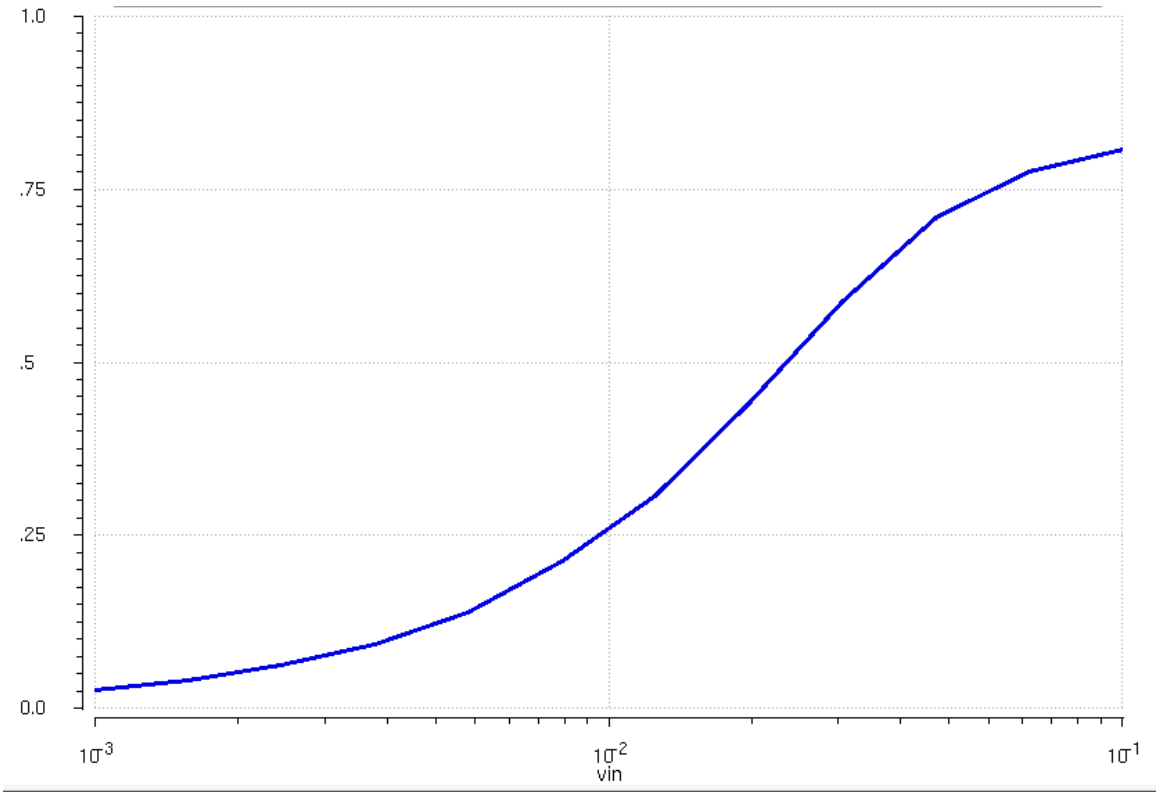

Figure 4.14: Energy detector gain. The X-axis shows the input signal level in volts, and the Y-axis shows the gain in V/V.

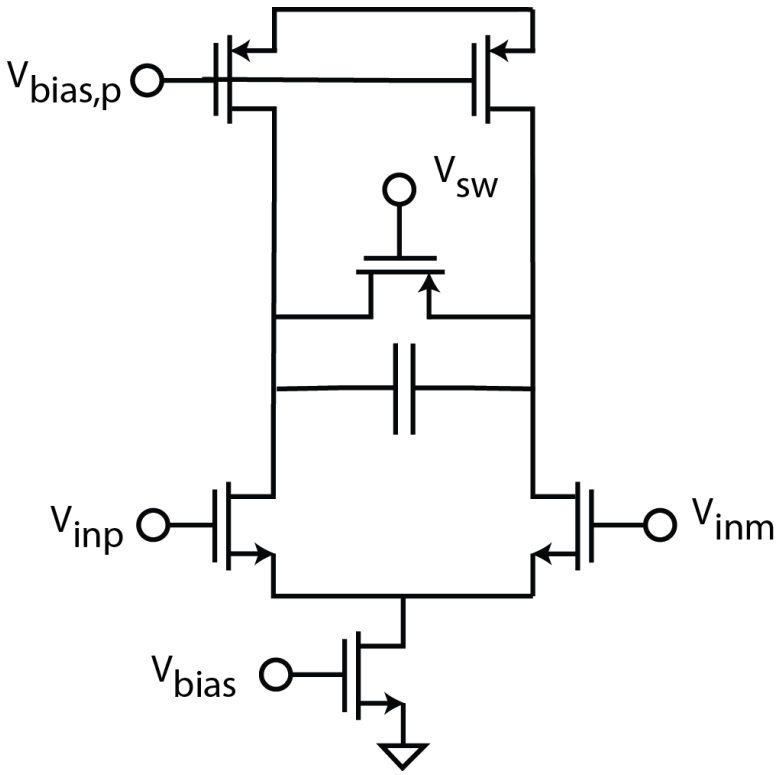

Figure 4.15: Integrator schematic

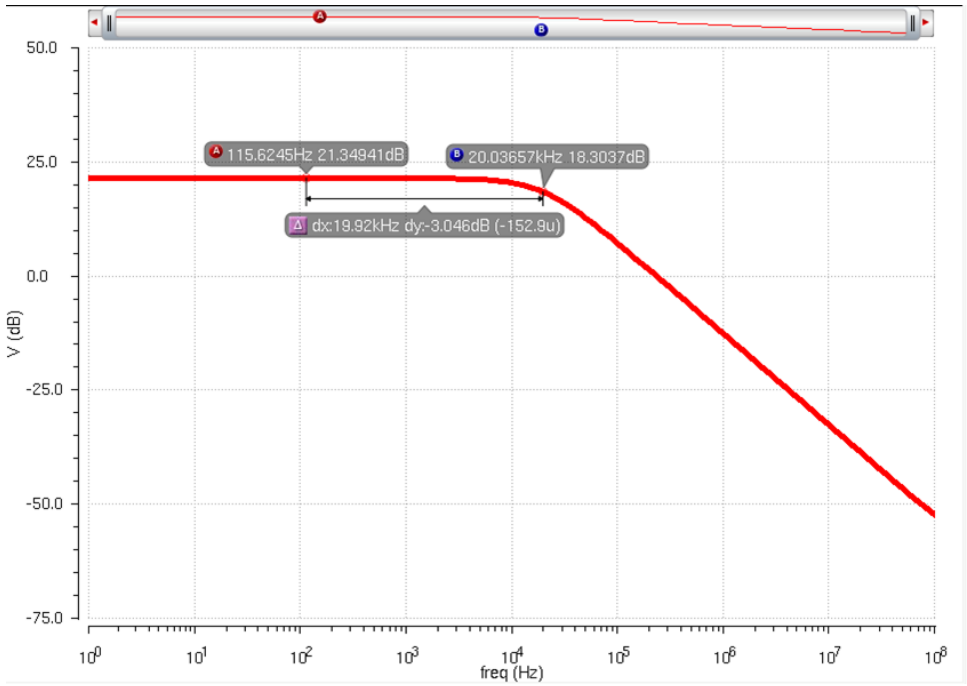

Figure 4.16: AC performance of the low-pass filter and integrator.

4.3 Sensitivity Analysis

Analyzing the expected sensitivity of the wake-up receiver is somewhat complicated by the presence of the non-linear energy detector element. The technique used in Section 3.3 provides a useful starting point for this analysis. In particular, we can again split the calculation into two parts: first, find the noise power present at the point of detection, and then use the SNR_{min} required to demodulate the signal to find the sensitivity.

To find the noise power, we can first start with the calculation from Chapter 3:

$$P_{noise} = -174dBm + 10 \cdot \log(BW) + NF \tag{4.3}$$

but we need to still account for the impact of the energy detector on the output noise levels. From [18], it was shown that the output SNR of a squaring energy detector is 3dB worse than the input SNR, and because this calculation uses input-referred noise power, this simply means adding a term of 3dB to the noise power. At this point, the noise still includes the full wide-band noise caused by the wide IF bandwidth used in the receiver's front-end, so the impact of filtering the excess noise needs to be taken into account. Because of the quadratic relationship to the output, filtering the noise only accounts for a 5dB per decade reduction in noise power, however, rather than the usual 10dB [22]. This means that a term of $-5 \cdot \log\left(\frac{BW_{IF}}{BW_{BB}}\right)$ should be included in the above

calculation. Putting this all together yields

$$P_{noise} = -174dBm + 10 \cdot \log\left(BW_{IF}\right) + NF + 3dB - 5 \cdot \log\left(\frac{BW_{IF}}{BW_{BB}}\right) \quad (4.4)$$

In the implemented design, the IF bandwidth was roughly 26MHz to accommodate the loose tuning for the LO. Because the data rate was not a critical aspect of the design, it was scaled down to less than 10kbps, and the filter following the energy detector has a 15kHz corner frequency. This means that the additional filtering reduced the noise power level by $5 \cdot \log\left(\frac{26MHz}{15kHz}\right) = 16.2dB$, illustrating an advantage compared to the digital energy detection proposed in Chapter 3, which yielded a maximum of roughly a 10dB improvement in SNR.

Given that the average noise figure of the front end was approximately 11dB, we can find the noise power as

$$P_{noise} = -174dBm + 10 \cdot \log\left(26MHz\right) + 11dB + 3dB - 16.2dB = -102dBm \quad (4.5)$$

The SNR_{min} required for demodulation of an OOK signal is 12dB for a 10^{-3} BER, so this means that the expected sensitivity level of this receiver is -90dBm.

4.4 Peripheral Circuits for Additional Functionality

One of the primary drawbacks of using a ring oscillator to generate the LO is that the frequency accuracy, stability, and phase noise of ring oscillators are substantially worse than can be achieved with other types of oscillators. For either simple OOK modulation or detecting the presence of energy in the channel, this does not cause a problem beyond requiring the excess IF bandwidth; however, to use the same architecture for more complex modulation schemes such as phase or frequency modulation, it is helpful to have a more accurate LO.

Although, as discussed in Section 2.4, power consumption of an FBAR oscillator is currently too high to be used as the main oscillator for the receiver, they can be useful for the duty-cycled portion of the receiver proposed in Fig. 4.1, and can then be used to injection lock the ring oscillator, which will cause the accuracy and phase noise performance of the LO to track the high quality FBAR oscillator. Because the ring oscillator is fully differential, the FBAR oscillator will also need to be differential to avoid unbalancing the ring oscillator and disturbing its phase balance. The FBAR oscillator is designed as in Fig. 4.17, using a classical cross-coupled transconductor topology along with an active load with common-mode feedback to set the DC level. The minimum current required to start the oscillator is roughly 150μA, with the nominal operation at roughly 200μA to provide sufficient margin to cover any process variations. This is coupled through a triode MOSFET switch into the ring oscillator core, providing a large injected signal to ensure a good locking range.

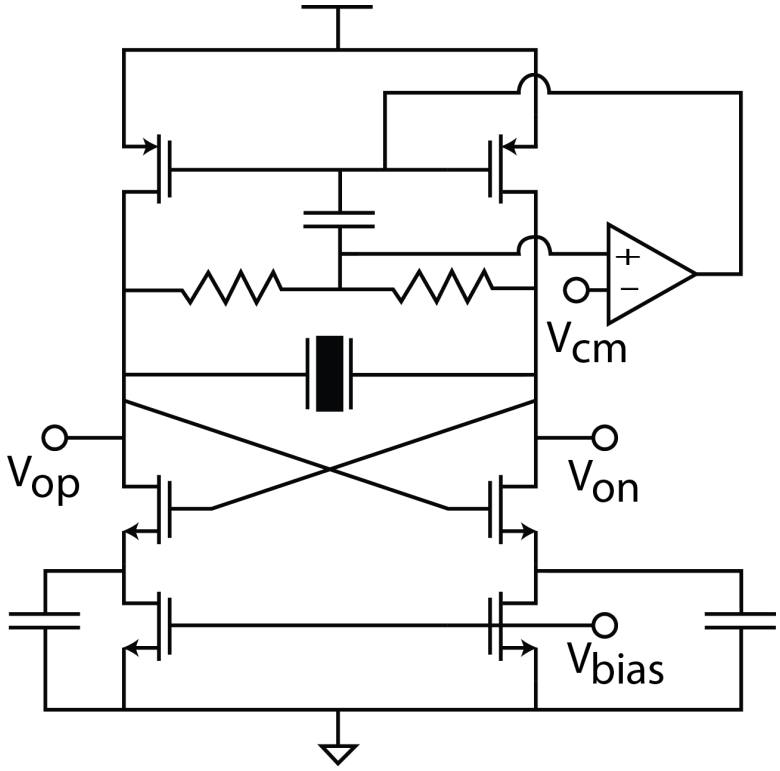

Figure 4.17: FBAR oscillator for injection locking

In order for this injection locking scheme to be useful during a two-phase wakeup operation, the FBAR oscillator must be able to start up and lock the main oscillator quickly. Fig. 4.18 shows the simulated startup time to be less than 10μs, so it would take only 1 bit period at the wakeup radio's nominal 8kbps data rate to start up. In addition, if the startup time needs to be reduced further, increasing the bias current reduces the amount of time required to start, as shown in [13].

While the OOK demodulator has a limited frequency performance in this architecture, due to the need to remove excess noise, if a faster data rate is desired, alternative demodulators can be used in its place. By switching to either an FSK or a BPSK demodulator, the front-end will support data rates of up to roughly 1 Mbps, which is instead limited by the front-end bandwidth of the filter and matching network. This demodulation can be disabled while the energy detection functionality is operating, and then enabled once energy is detected in the signal band, enabling a more robust set of functionality than a simple energy detection radio alone.

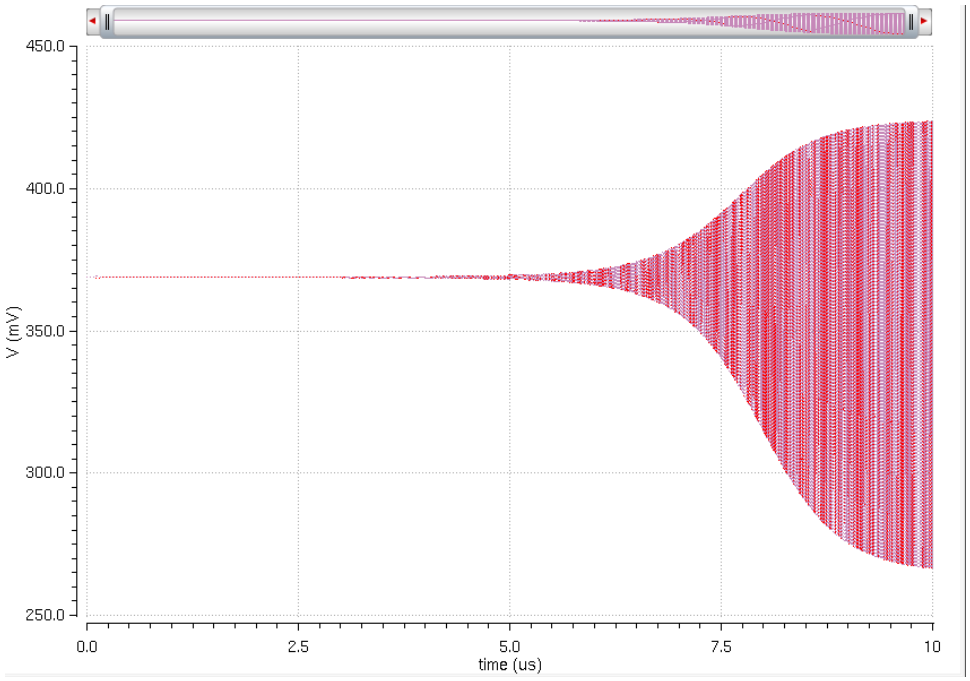

Figure 4.18: FBAR oscillator startup transient

4.5 Measured Results

The wakeup receiver was implemented in a standard digital 65nm CMOS process with no extra process options. A die photograph showing both the CMOS die alongside the two FBAR resonator dies is shown in Fig. 4.19. The complete chip measures 1x1.2mm, and the layout is shown in Fig. 4.20, but the design was pad-limited and the active area occupied only 220x450µm. A significant portion of the total active area went to the main filter and integration capacitor on the output stage; with either denser on-chip capacitors or an external main capacitor, the active area could be reduced significantly. The layout of the main wake-up receiver path, including the matching network, passive mixer, IF amplifiers, energy detector, and baseband filter are shown in Fig. 4.21. For simplicity in fabrication and testing, the FBAR resonator dies were placed to the side of the CMOS die and wirebonded together; however, if area is a premium and it is necessary to limit the space devoted to these structures, the FBAR resonators could have been placed above the CMOS as in [33], permitting a smaller board area.

4.5.1 RF Input Performance and LO Generation

The front-end matching network was designed expecting to use an FBAR resonator at 900MHz with a Q of approximately 1500. The resonators that were actually received and used in the test chip were slightly different, however, due to the available stock

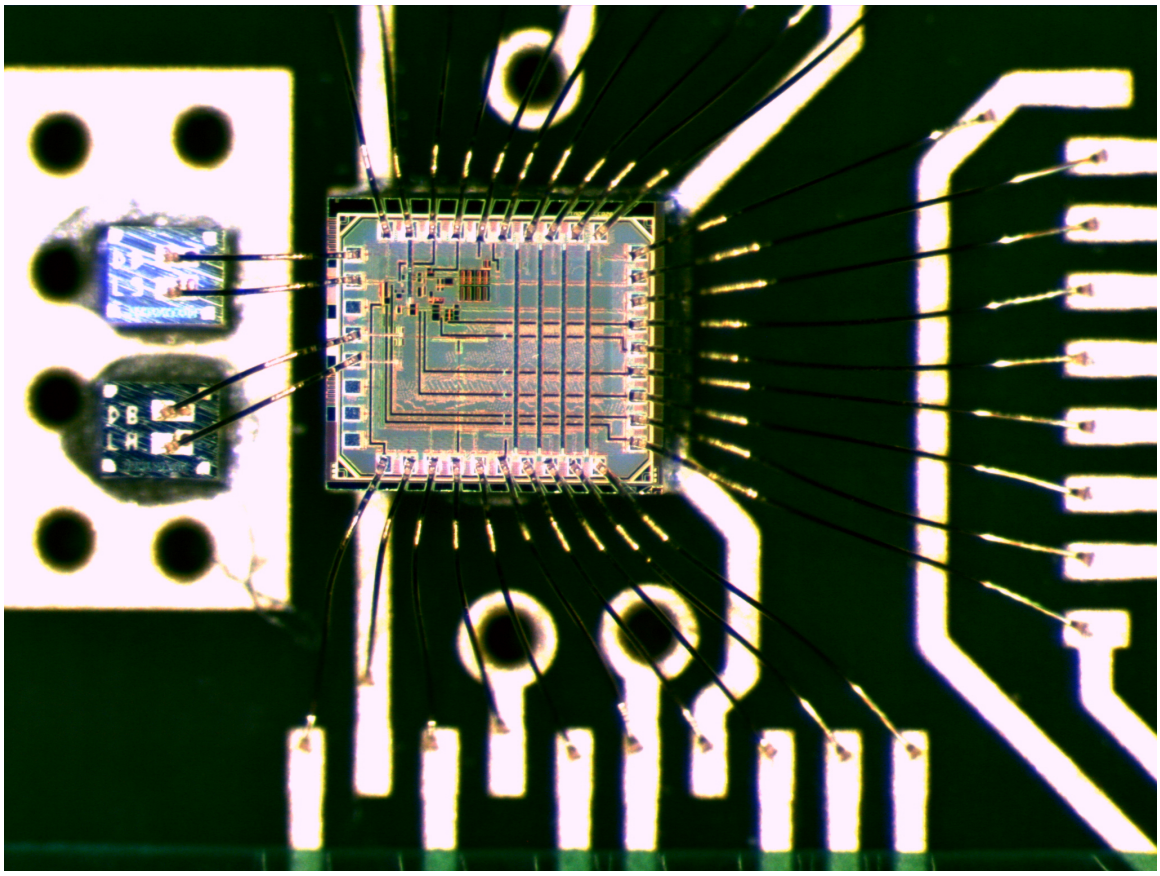

Figure 4.19: Die Photograph of Bonded Chip and FBARs

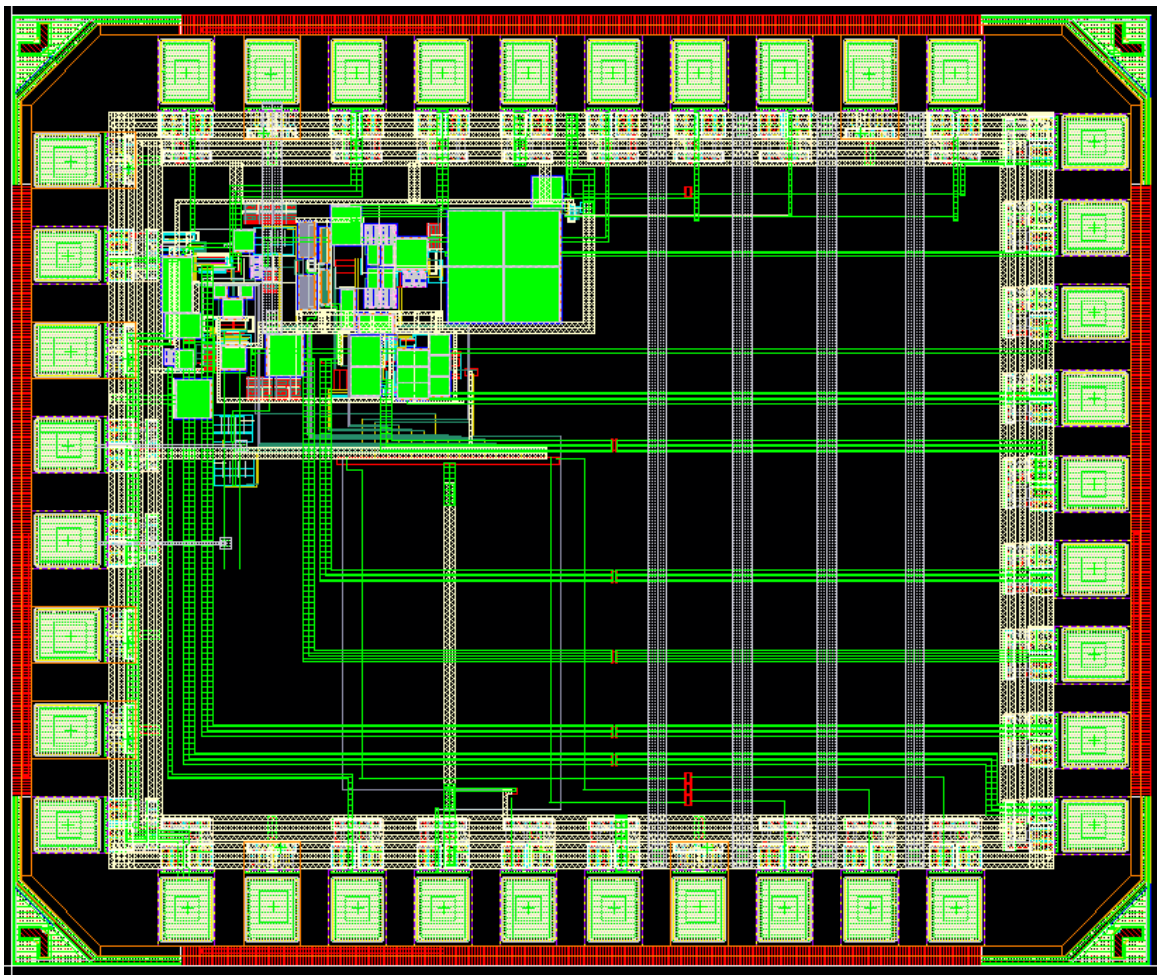

Figure 4.20: Layout view of the complete design

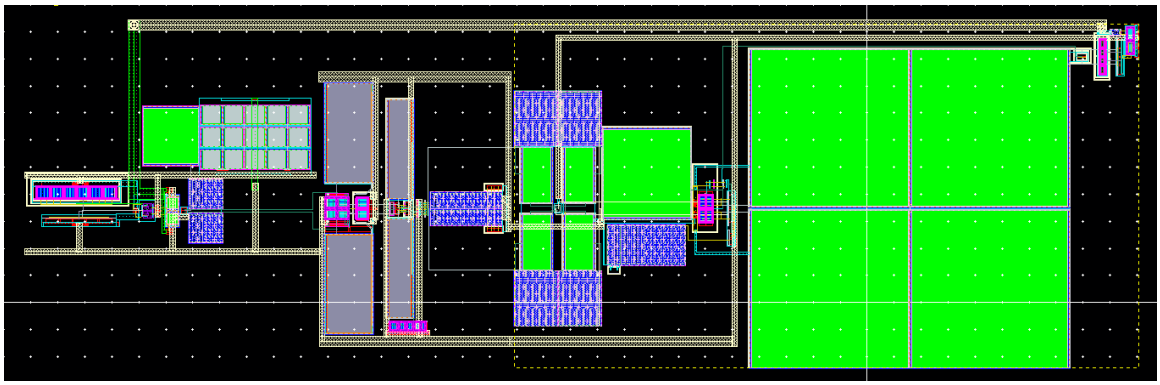

Figure 4.21: Layout view of the core always-on wake-up receiver

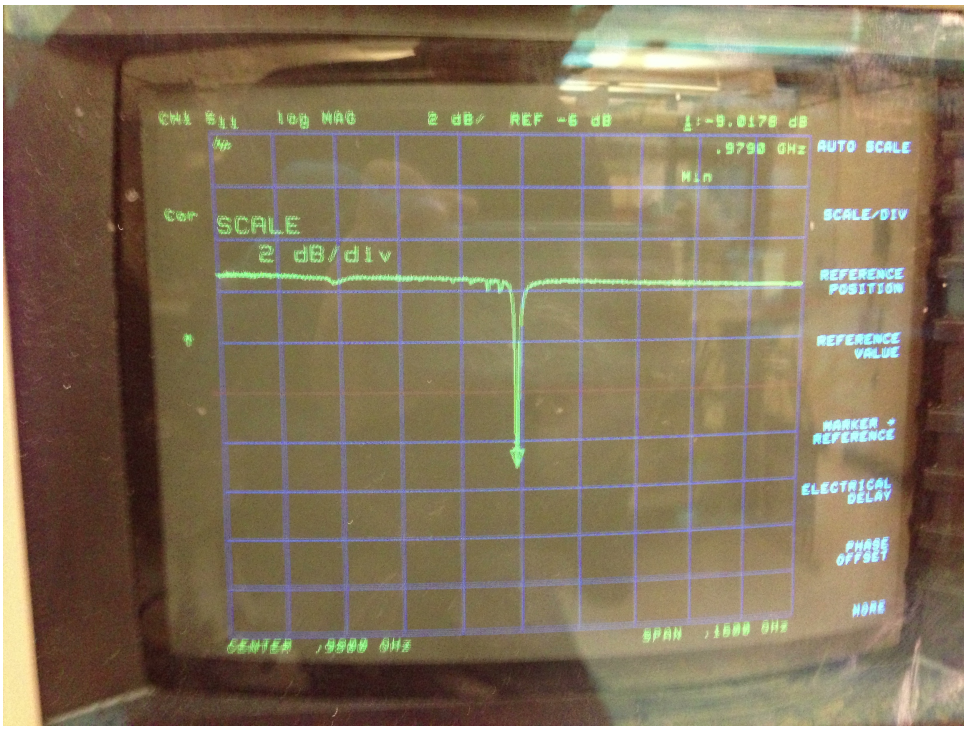

Figure 4.22: Measured S11 of Receiver

from Avago Technologies, with an unloaded parallel resonant frequency of approximately 1GHz and a Q of 2500. This indicates that the resonators actually perform better than was expected at design time; however, this also means that the input matching network is somewhat mis-tuned for the actual resonators. Fig. 4.22 shows the measured input S11, which achieves a minimum of only -9dB at 988MHz after loading by the matching network, rather than the commonly used -10dB to indicate a proper match. This could be corrected in future silicon by re-tuning the matching network, but is still sufficient to test the receiver.

The LO frequency matched well with the expected frequency based on simulations, indicating that the biasing and control circuitry are operating correctly. Fig. 4.23 shows the measured current consumption of the oscillator along with the buffers driving the passive mixer. At the 980MHz nominal tuning point for this resonator frequency, the LO generation consumed 39μA, which matches well with simulated expectations, given the slightly higher operating frequency due to the FBARs. The temperature stability of the oscillator was checked by raising the temperature of the die from room temperature up to 80°C, resulting in the operating frequency increasing from 980MHz to 1,004MHz, which matches well with the simulated result.

The output spectrum of the main ring oscillator is shown in Fig. 4.24. Because the oscillator operates open-loop, the spectrum can be seen to "wander" during operation, where its center frequency moves around a roughly 5-10MHz range. In addition, the

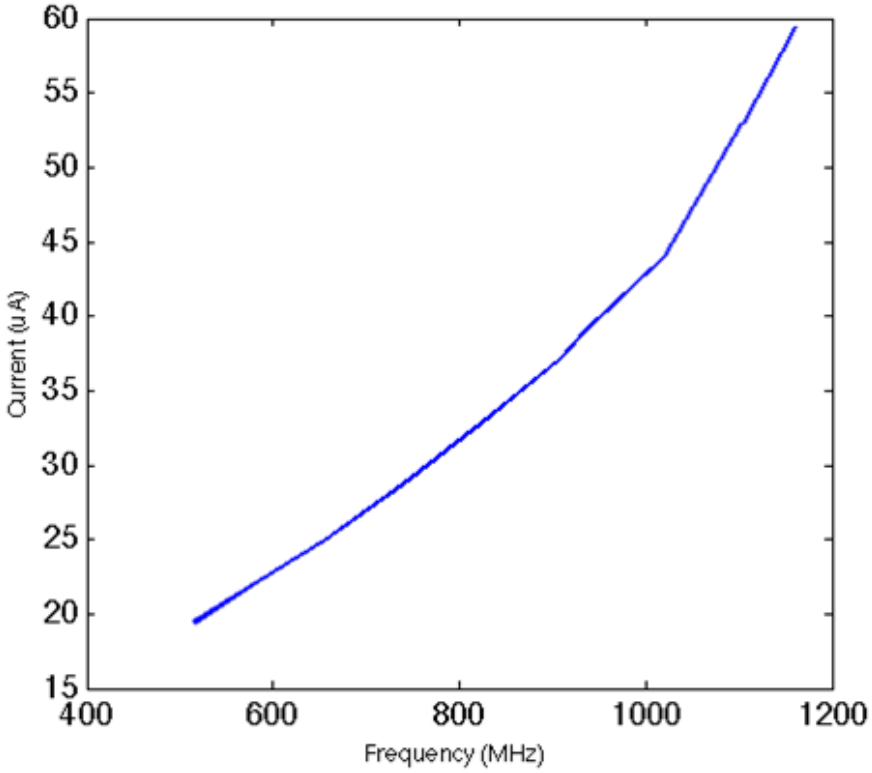

Figure 4.23: Power Consumption of LO

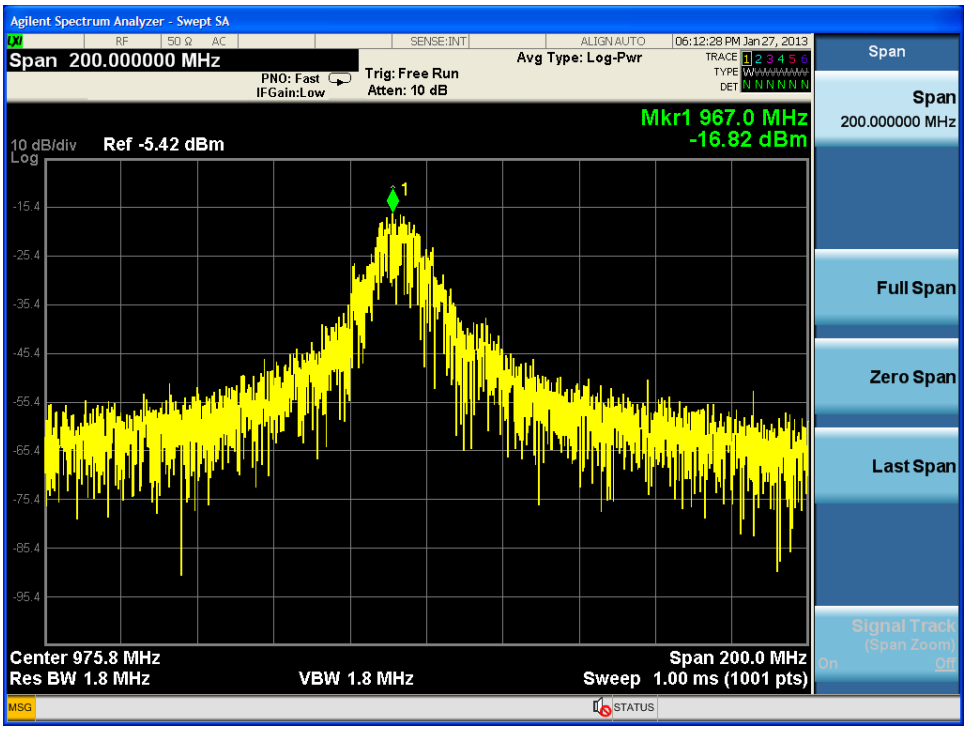

Figure 4.24: Measured ring oscillator output spectrum

very large amount of phase noise present from the ring oscillator can be seen from this measurement as the wide range of energy present from the oscillator output.

The much improved stability and phase noise performance of the duty-cycled FBAR oscillator can be clearly seen in Fig. 4.25. In the spectrum plot, note that the frequency span is much smaller than the measurement with the ring oscillator alone. In contrast with the ring oscillator, the output frequency of the FBAR oscillator is very stable, as expected due to the very high Q of the resonator. The FBAR oscillator requires a minimum of 450μA to achieve a stable injection lock, which is somewhat higher than expected from simulation. It is believed that this is due to parasitic loading from the routing of the signals being somewhat higher than anticipated and should be able to be improved in future versions of the work.

4.5.2 Sensitivity

The receiver was first characterized without clocking the integrator, such that the output was simply low-pass filtered. The sensitivity of the receiver was measured by using a BER meter to generate a pseudo-random bit sequence (PRBS), which was used to modulate the carrier using an AM-modulation function with a 100% modulation index. On-chip, the receiver includes a differential to single-ended buffer, which is able to drive a large amount of parasitic capacitance such as what might be found in a PCB. To handle

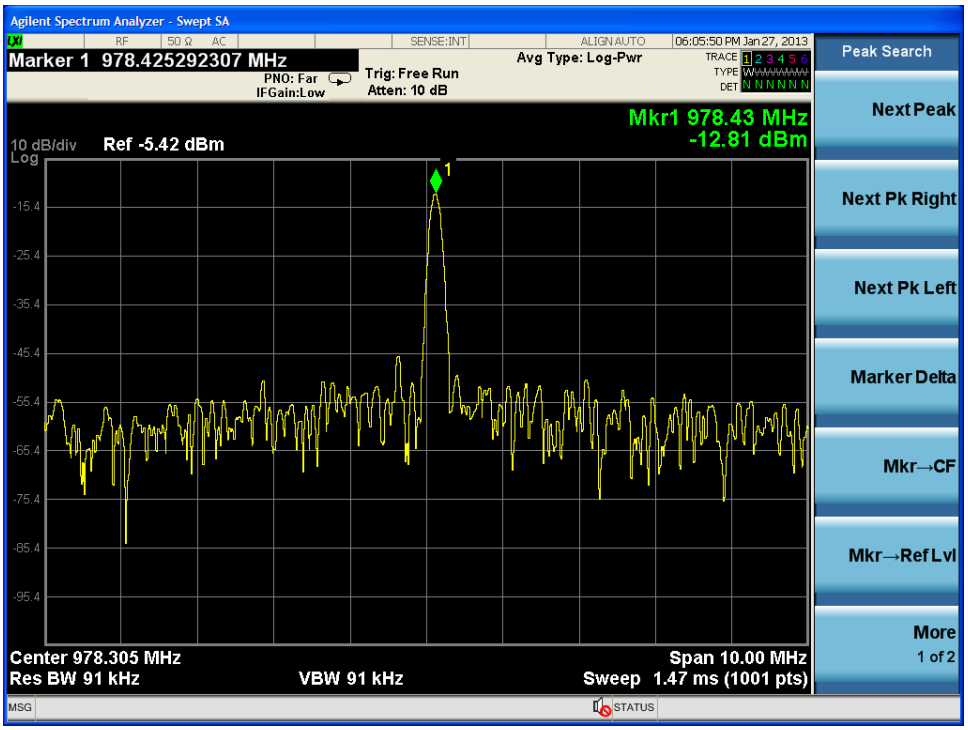

Figure 4.25: FBAR oscillator output spectrum

the conversion between the 0.5V chip power supply voltage and the larger voltage levels required by instrumentation in the lab, first the signal was amplified by 20dB using an off-the-shelf commercial op-amp circuit. The output was sliced by a commercial comparator to provide full-swing logic levels, and fed back into the BER meter.

The measured sensitivity of the receiver at its nominal bias point was -90dBm with an 8kbps OOK signal. The total power consumption at this bias was 37.5μW, with 19.5μW dedicated to the LO and mixer, 15.5μW to the IF LNA, and 2.5μW to the remainder of the analog blocks, consisting of the remaining IF amplifiers, energy detector, and low-pass filter. By varying the LNA bias point, the sensitivity can be improved to -92dBm with a total power consumption of 45μW, while reducing the power consumption to 30μW yields a sensitivity of -83dBm. Fig. 4.26 shows the measured BER curves for each of the bias points of the amplifier.

The bandwidth of the receiver is determined primarily by the FBAR's bandwidth in the front end filter and matching network. The capacitive loading on the FBAR, along with the reduction in Q due to the real resistive match to the antenna impedance, causes the bandwidth to be roughly 2MHz, or a Q of 450, compared with an unloaded Q for the FBAR of 2,500. The Q for the receiver could likely be improved slightly by re-tuning the matching network with less total capacitance. The measured bandwidth of the receiver is shown in Fig. 4.27, illustrating that the bandwidth is determined mainly by the impedance curve of the resonator.

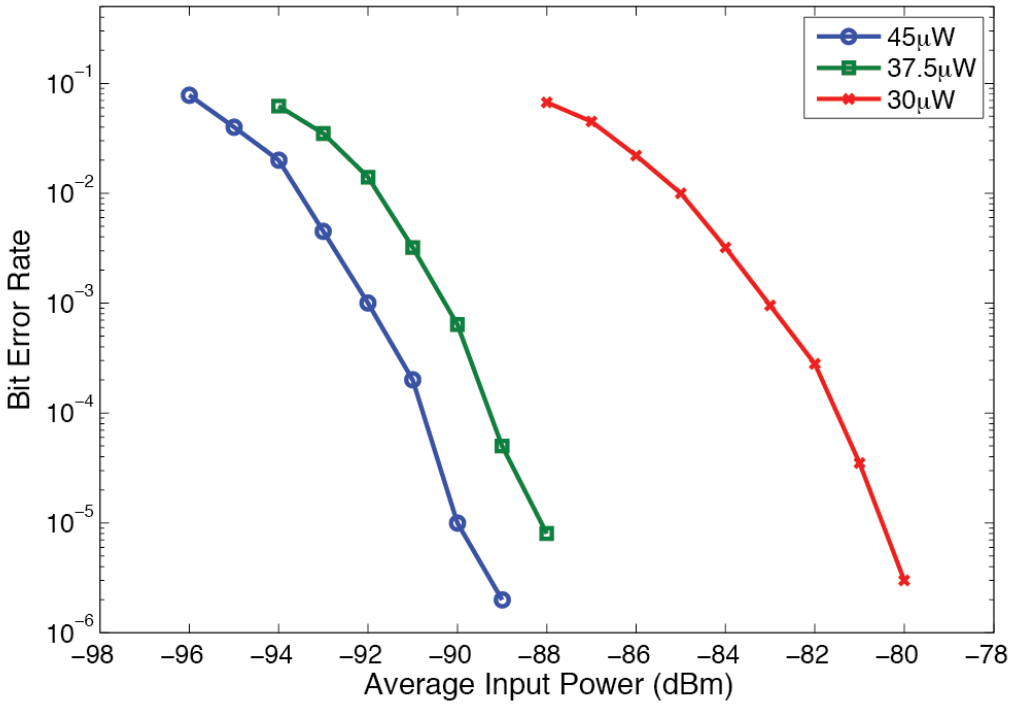

Figure 4.26: Bit Error Rate versus Input Power

Unfortunately, the measured results when clocking the integrator did not yield the expected improvement in sensitivity. In this implementation, the large integration capacitor is placed across the differential outputs of the integration amplifier. This means that the large capacitance is seen in the differential mode, but is not seen in the common mode. The output buffer was implemented as a simple differential pair with a single-ended output, which leads to a relatively low amount of common mode rejection due to the very low supply levels and the resulting small amount of headroom for the tail current transistor. This issue means that the integrator and buffer passes interference or additional noise that reaches the integrator to the output. In addition, by the nature of the integration operation, the output signal levels that are expected to differentiate either the presence or lack of a carrier are fairly small, on the order of 10mV, making it easy to disrupt the detection due to the lack of proper filtering. It is believed that this issue could be corrected in a future implementation of the concept by instead placing the capacitance at the output between the terminals and either ground or the chip supply, which would act to remove common mode signals as well as differential mode signals. According to simulation results and calculations, the use of the integration is expected to yield approximately another 5dB of sensitivity performance. Note that even without this improvement, the receiver as measured still delivers very competitive performance relative to the state-of-the-art.

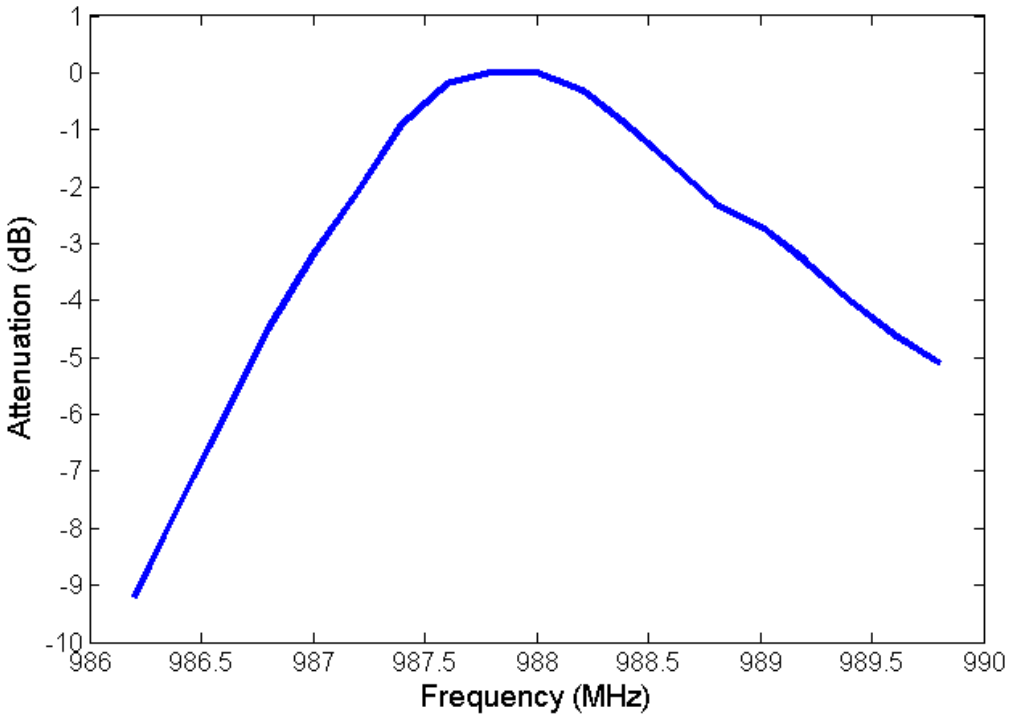

Figure 4.27: Measured Bandwidth

4.5.3 Comparison with Prior Art

The receiver described in this chapter focused primarily on improving the sensitivity compared to prior art, while still maintaining an ultra-low power level of below 50μW. A comparison with previously published receivers with similar goals is shown in Table 4.1. This receiver achieves both the lowest power consumption as well as the best sensitivity for the narrowband receivers, and offers a much lower active power than the UWB receiver in [25].

In addition, compared to the works in [22, 21, 38], this architecture provides the significant advantage of being functionally equivalent to a low-IF receiver front end, because the energy detection occurs at the end of the baseband chain. This property means that, if desired, a future designer could use the same front-end topology providing the very low power consumption, but follow it with a more typical ADC and perform demodulation in the digital domain, which would open the receiver to a range of different possible modulation schemes. This is a key advantage to this type of architecture which, combined with the strong sensitivity performance, suggests a path to integrate ultra-low power wakeup radios alongside current existing standards, allowing their use in ways that were not previously possible.

	[15]	[25]	[38]	[39]	[21]	[17]	This work
Technology	130nm	90nm	90nm	90nm	90nm	65nm	65nm
Architecture	Low-IF	Direct Conv.	Tuned RF	Uncertain IF	Tuned RF / Low-IF	Low-IF	Uncertain-IF
Active Power	330μW	250μW Avg.	65μW	52μW	123μW	415μW	37.5μW
Data Rate	N/A	100kbps	100kbps	100kbps	10kbps	500kbps	8kbps
Sensitivity	N/A	-99dBm	-56dBm	-72dBm	-86dBm	-82dBm	-90dBm
Frequency	2.4GHz	3-5GHz	1.9GHz	1.9GHz	915MHz	2.4GHz	988MHz
Modulation	2-FSK	Pulsed UWB	OOK	OOK	OOK	PPM-IR	OOK

Table 4.1: Comparison table with previously published works.

Chapter 5

Conclusions and Future Work

5.1 Research Summary

The number of devices that can be connected together wirelessly is rapidly increasing, enabling applications and ideas that were previously not possible. In order to continue this push towards making wireless more ubiquitous, radios must become ever cheaper, smaller, and lower in power consumption. By achieving useful and reliable radios with low enough power consumptions, it becomes possible to envision radios that can run indefinitely using only energy harvested from the ambient environment, truly making ubiquitous wireless possible.

In Chapter 2, this work discussed the types of architectures and technologies that are used to design low power radios. It also explored the metrics that are available for quantifying the performance of radios and looked at the types of trade-offs that are required for low-power radios.

This was followed by two receiver prototypes targeting applications that need ultra-low power designs with different requirements. In Chapter 3, the focus was on implementing a complete platform for wireless sensing, including all the basic components that a full system would require such as power management, radios, and a digital baseband to control the system. In this work, the receiver needed to be able to coexist with the other components and acted as the main radio for a sensor. This radio consumed a total of 109μW to receive 200kbps of data with a sensitivity of -66dBm, even in the presence of relatively large amounts of ripple from an on-chip switching power supply. Chapter 4 described a wake-up receiver, designed to be an extremely low power, always-on counterpart for a more complex radio. The primary goal of this project was to push the boundaries of sensitivity. This design resulted in a receiver that achieved -90dBm sensitivity for an 8kbps data stream while consuming only 37.5μW of power.

5.2 Future Work

The pressure to continue pushing power consumption down and improving sensitivity for receivers is likely to continue into the future. While the active power consumption required by these receivers is competitive against others that have been previously designed, they still require more power than many energy harvesting techniques are capable of providing. In addition, further improvements in sensitivity will allow for either longer transmission distances for sensors or for the transmitted power to be reduced.

Figure 5.1 shows a breakdown of the power consumption of the wake-up receiver presented in this work. Because the LO and mixer power is driving digital gates and switches, it can be expected that this power will be reduced with future CMOS scaling, as the device sizes continues to shrink. Because this currently accounts for a majority of the total power in the receiver, this promises to yield improvements in reducing the overall power with further technology improvements; however, the analog portions will be difficult to reduce further because the amount of power consumed is determined by the

amount of current needed to meet the given noise figure targets. Although the short-term trend in technologies is for the threshold voltages to not decrease (and in fact, as seen in Table 2.2, may increase), making it difficult to utilize the other means of reducing power by reducing the supply volage, upcoming devices such as FinFETs and silicon-on-insulator technologies may yield improvements later.

MEMS resonator technology is still a very active area of research. For the designs featured here, the FBAR resonator is the limitation on the bandwidth and interference handling that can be achieved, but if future devices are able to significantly increase the Q, the front-end filter performance can be improved. In addition, the FBAR devices represent an added cost, because they require special processing techniques to apply the piezoelectric material; if it becomes possible to fabricate high-Q resonators in a fully CMOS-compatible manner directly along with the active circuitry, the use of MEMS resonators would become still more attractive. In addition, work is underway on techniques to use passive mixer structures to remove the need for the high-Q resonators [6, 15].

From an architectural standpoint, one of the limitations of these designs remains that they are fixed to a single channel, limiting their ability to cope with congested frequency bands such as the 2.4GHz ISM band. Figure 5.2 shows one possibility for a way to both work around this limitation as well as take advantage of the small sizes offered by recent MEMS components. Rather than using a single channel, using multiple resonators, each connected to its own separate receive path, would enable both additional channels for interference mitigation as well as opening the possibility of more complex operation, such as encoding transmissions across multiple channels for either improved sensitivity performance or increased throughput. In addition, many of the components can be made reconfigurable to dynamically control the performance and power consumption of the receiver, helping to further optimize it for its conditions.

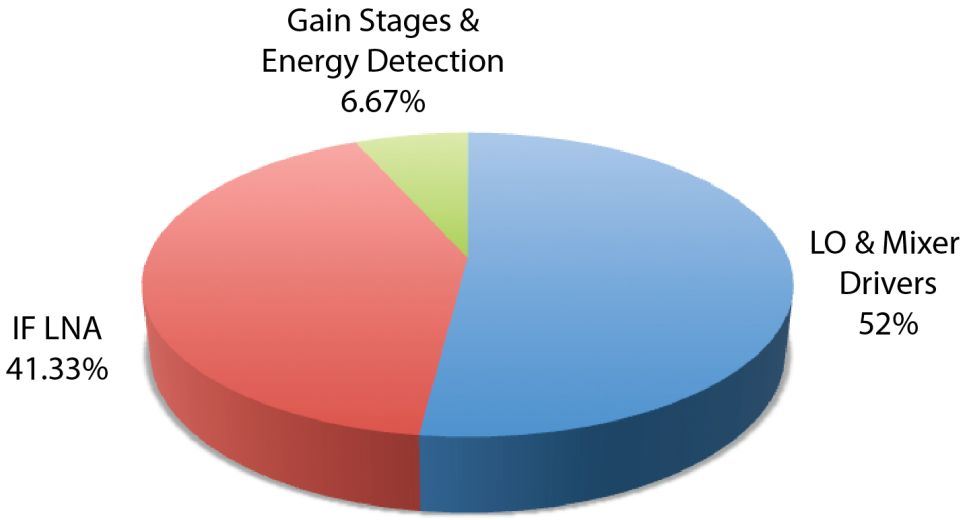

Figure 5.1: Power Consumption Breakdown for Wake-Up Radio from Chapter 4

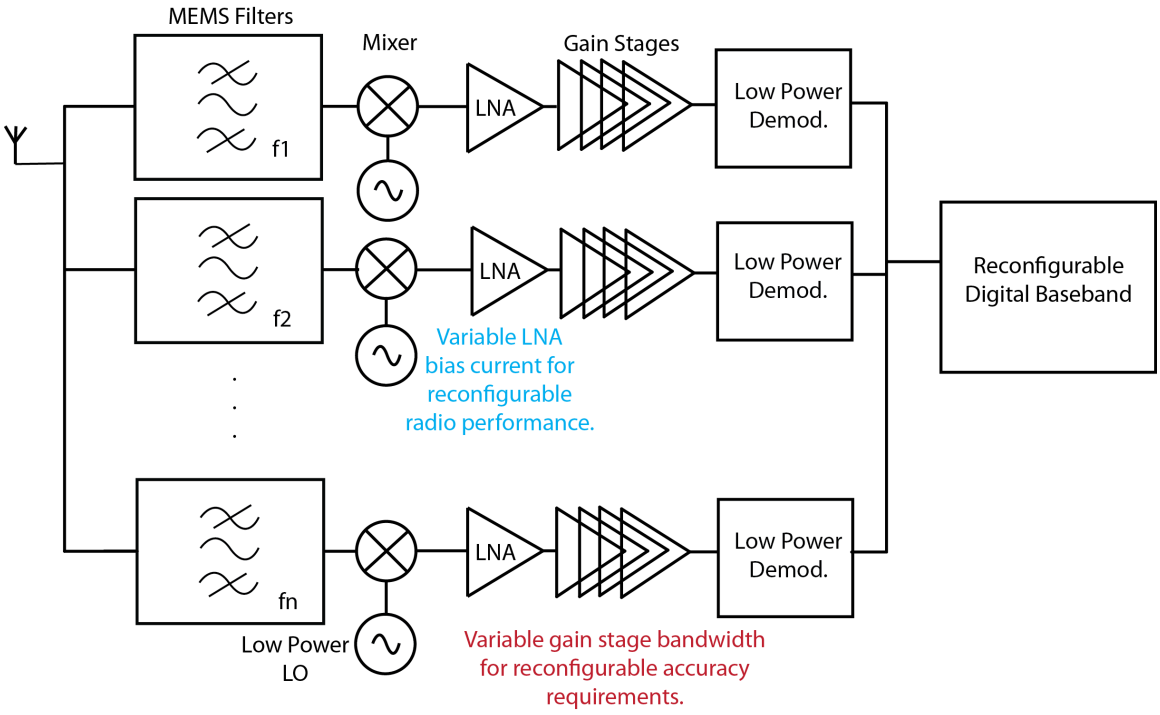

Figure 5.2: Multi-channel MEMS-based Receiver

Bibliography

[1] The International Technology Roadmap for Semiconductors (ITRS), System Drivers, 2009, http://www.itrs.net/.

[2] ISO/IEC 18000-7. Information technology - radio frequency identification for item management - Part 7: parameters for active air interface communications at 433 MHz, 2008.

[3] M.A. Abdelmoneum, M.U. Demirci, and C.T.-C. Nguyen. Stemless wine-glass-mode disk micromechanical resonators. In *The Sixteenth Annual International Conference on Micro Electro Mechanical Systems, 2003. MEMS-03 Kyoto. IEEE*, pages 698–701. IEEE.

[4] AA Abidi. Direct-conversion radio transceivers for digital communications. *Solid-State Circuits, IEEE Journal of*, 30(12):1399–1410, 1995.

[5] J Ammer and J Rabacy. The Energy-per-Useful-Bit Metric for Evaluating and Optimizing Sensor Network Physical Layers. *Sensor and Ad Hoc Communications and Networks, 2006. SECON '06. 2006 3rd Annual IEEE Communications Society on*, 2:695–700, 2006.

[6] C Andrews and A C Molnar. A Passive Mixer-First Receiver With Digitally Controlled and Widely Tunable RF Interface. *Solid-State Circuits, IEEE Journal of*, 45(12):2696–2708, 2010.

[7] E.H. Armstrong. Some Recent Developments of Regenerative Circuits. *Proceedings of the IRE*, 10(4):244–260, August 1922.

[8] EH Armstrong. The super-heterodyne-its origin, development, and some recent improvements. *Proceedings of the Institute of Radio Engineers*, pages 539–552, 1924.

[9] E Bodtker, G Jacobsen, and B Jensen. Design and test of novel integrate and dump filter (I&D) for optical Gbit/s system applications. *Electronics Letters*, 27(24):2286–2288, 1991.

[10] J.E. Butler, T. Feygelson, and C.T.-C. Nguyen. 1.51-GHz nanocrystalline diamond micromechanical disk resonator with material-mismatched isolating support. In *17th IEEE International Conference on Micro Electro Mechanical Systems. Maastricht MEMS 2004 Technical Digest*, pages 641–644. IEEE.

[11] YH Chee, M Koplow, M Mark, et al. PicoCube: a 1 cm 3 sensor node powered by harvested energy. *Proceedings of the 45th ...*, pages 114–119, 2008.

[12] Y.H. Chee, A.M. Niknejad, and J. Rabaey. A sub-100 μW 1.9-GHz CMOS oscillator using FBAR resonator. *Radio Frequency integrated Circuits (RFIC) Symposium, 2005. Digest of Papers. 2005 IEEE*, pages 123–126, 2005.

[13] Y.H. Chee, A.M. Niknejad, and J. Rabaey. A 46% Efficient 0.8dBm Transmitter for Wireless Sensor Networks. *VLSI Circuits, 2006. Digest of Technical Papers. 2006 Symposium on*, pages 43–44, 2006.

[14] G Chen, M Fojtik, Daeyeon Kim, et al. Millimeter-scale nearly perpetual sensor system with stacked battery and solar cells. *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2010 IEEE International*, pages 288–289, 2010.

[15] B.W. Cook, A. Berny, A. Molnar, S. Lanzisera, and K.S.J. Pister. Low-power 2.4-GHz transceiver with passive RX front-end and 400-mV supply. *Solid-State Circuits, IEEE Journal of*, 41(12):2757–2766, 2006.

[16] Jan Crols and MSJ Steyaert. A single-chip 900 MHz CMOS receiver front-end with a high performance low-IF topology. *Solid-State Circuits, IEEE Journal of*, 30(12):1483–1492, 1995.

[17] Salvatore Drago, Domine M. W Leenaerts, Fabio Sebastiano, et al. A 2.4GHz 830pJ/bit duty-cycled wake-up receiver with -82dBm sensitivity for crystal-less wireless sensor nodes. In *2010 IEEE International Solid-State Circuits Conference - (ISSCC)*, pages 224–225. IEEE, February 2010.

[18] Simone Gambini, Nathan Pletcher, and Jan M Rabaey. Sensitivity Analysis for AM Detectors. Technical Report UCB/EECS-2008-31, EECS Department, University of California, Berkeley, April 2008.

[19] N Gershenfeld, R Krikorian, and D Cohen. The Internet of things. *Scientific American*, 291(4):76–81, 2004.

[20] P R Gray, P J Hurst, R G Meyer, and S H Lewis. *Analysis and design of analog integrated circuits*. John Wiley & Sons, 2008.

[21] Xiongchuan Huang, P Harpe, G Dolmans, and H de Groot. A 915MHz ultra-low power wake-up receiver with scalable performance and power consumption. *ESSCIRC (ESSCIRC), 2011 Proceedings of the*, pages 543–546, 2011.

[22] Xiongchuan Huang and S Rampu. A 2.4 GHz/915MHz 51uW wake-up receiver with offset and noise suppression. *Solid-State Circuits ...*, pages 222–223, 2010.

[23] K.M. Lakin, K.T. McCarron, and R.E. Rose. Solidly mounted resonators and filters. In *1995 IEEE Ultrasonics Symposium. Proceedings. An International Symposium*, volume 2, pages 905–908. IEEE.

[24] H-P Le, M Seeman, S Sanders, et al. A 32nm fully integrated reconfigurable switched-capacitor DC-DC converter delivering 0.55W/mm2 at 81% efficiency. *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2010 IEEE International*, pages 210–211, 2010.

[25] FS Lee and AP Chandrakasan. A 2.5 nJ/b 0.65 V 3-to-5GHz subbanded UWB receiver in 90nm CMOS. *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International*, pages 116–118, 2007.

[26] T H Lee. *The design of CMOS radio-frequency integrated circuits*. Cambridge university press, 2003.

[27] E Y A Lin, J M Rabaey, and A Wolisz. Power-efficient rendez-vous schemes for dense wireless sensor networks. In *Communications, 2004 IEEE International Conference on*, volume 7, pages 3769–3776. IEEE, 2004.

[28] Y W Lin, S Lee, S S Li, et al. Series-resonant VHF micromechanical resonator reference oscillators. *Solid-State Circuits, IEEE Journal of*, 39(12):2477–2491, 2004.

[29] D. Markovic, V. Stojanovic, B. Nikolic, M.A. Horowitz, and R.W. Brodersen. Methods for true energy-performance optimization. *IEEE Journal of Solid-State Circuits*, 39(8):1282–1293, August 2004.

[30] D. Markovic, C.C. Wang, L.P. Alarcon, and J.M. Rabaey. Ultralow-Power Design in Near-Threshold Region. *Proceedings of the IEEE*, 98(2):237–252, February 2010.

[31] S. Marksteiner, J. Kaitila, G.G. Fattinger, and R. Aigner. Optimization of acoustic mirrors for solidly mounted BAW resonators. In *IEEE Ultrasonics Symposium, 2005.*, volume 1, pages 329–332. IEEE.

[32] G.E. Moore. No exponential is forever: but "Forever" can be delayed! [semiconductor industry]. In *2003 IEEE International Solid-State Circuits Conference, 2003. Digest of Technical Papers. ISSCC.*, volume 1, pages 20–23. IEEE.

[33] Andrew Nelson, Julie Hu, Jyrki Kaitila, Richard Ruby, and Brian Otis. A 22μW, 2.0 GHz FBAR oscillator. In *Radio Frequency Integrated Circuits Symposium (RFIC), 2011 IEEE*, pages 1–4. IEEE, 2011.

[34] C. T.-C. Nguyen. MEMS technology for timing and frequency control. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 54(2):251–270, 2007.

[35] B. Otis, Y.H. Chee, and J. Rabaey. A 400μW-RX, 1.6mW-TX superregenerative transceiver for wireless sensor networks. In *ISSCC. 2005 IEEE International Digest of Technical Papers. Solid-State Circuits Conference, 2005.*, pages 396–398. IEEE.

[36] B P Otis and J M Rabaey. A 300-μW 1.9-GHz CMOS oscillator utilizing micromachined resonators. *Solid-State Circuits, IEEE Journal of*, 38(7):1271–1274, 2003.

[37] B.P. Otis, Y.H. Chee, R. Lu, N.M. Pletcher, and J.M. Rabaey. An ultra-low power MEMS-based two-channel transceiver for wireless sensor networks. *VLSI Circuits, 2004. Digest of Technical Papers. 2004 Symposium on*, pages 20–23, 2004.

[38] N Pletcher, S Gambini, and J Rabaey. A 65 μW, 1.9 GHz RF to digital baseband wakeup receiver for wireless sensor nodes. *Custom Integrated Circuits Conference, 2007. CICC '07. IEEE*, pages 539–542, 2007.

[39] NM Pletcher, S. Gambini, and J M Rabaey. A 52 W wake-up receiver with 72 dBm sensitivity using an uncertain-IF architecture. *IEEE J. Solid-State ...*, 44(1):269–280, 2009.

[40] J Rabaey, J Ammer, J L da Silva Jr, and D Patel. PicoRadio: Ad-hoc wireless networking of ubiquitous low-energy sensor/monitor nodes. In *VLSI, 2000. Proceedings. IEEE Computer Society Workshop on*, pages 9–12. IEEE, 2000.

[41] J M Rabaey. The swarm at the edge of the cloud-A new perspective on wireless. In *VLSI Circuits (VLSIC), 2011 Symposium on*, pages 6–8. IEEE, 2011.

[42] J M Rabaey, A P Chandrakasan, and B Nikolic. *Digital integrated circuits*, volume 996. Prentice-Hall, 1996.

[43] Jesse Richmond. *Digital Energy Detection for OOK Demodulation in Ultra-Low Power Radios*. Master's project report, University of California, Berkeley, 2010.

[44] Jesse Richmond, Mervin John, Louis Alarcon, et al. Active RFID: Perpetual wireless communications platform for sensors. In *2012 Proceedings of the ESSCIRC (ESSCIRC)*, pages 434–437. IEEE, September 2012.

[45] Jesse Richmond and Jan Rabaey. Digital energy detection for OOK demodulation in ultra-low power radios. In *2011 IEEE International Symposium of Circuits and Systems (ISCAS)*, pages 1808–1811. IEEE, May 2011.

[46] R. Ruby, P. Bradley, D. Clark, et al. Acoustic FBAR for filters, duplexers and front end modules. *Microwave Symposium Digest, 2004 IEEE MTT-S International*, 2:931–934 Vol.2, 2004.

[47] R. Ruby, P. Bradley, Y. Oshmyansky, and A. Chien. Thin film bulk wave acoustic resonators (FBAR) for wireless applications. *Ultrasonics Symposium, 2001 IEEE*, 1:813–821 vol.1, 2001.

[48] Michael D. Seeman, Seth R. Sanders, and Jan M. Rabaey. An Ultra-Low-Power Power Management IC for Wireless Sensor Nodes. In *2007 IEEE Custom Integrated Circuits Conference*, pages 567–570. IEEE, 2007.

[49] S D Senturia. *Microsystem design*, volume 3. Kluwer academic publishers Boston, MA, 2001.

[50] Y Tsividis. Operation and modelling of the MOS transistor, 1998. 1998.

[51] S. Zhou and M.-C.F. Chang. A CMOS passive mixer with low flicker noise for low-power direct-conversion receiver. *IEEE Journal of Solid-State Circuits*, 40(5):1084–1093, May 2005.

[52] Wenting Zhou, J Richmond, and Jan Rabaey. High Performance Wake-Up Receiver, 2012.