RF INTEGRATED CIRCUITS FOR ENERGY AUTONOMOUS SENSOR NODES

by

Nathan E. Roberts

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Electrical Engineering)
in The University of Michigan
2014

Doctoral Committee:

Associate Professor David D. Wentzloff, Chair
Assistant Professor SangHyun Lee
Assistant Professor Mina Rais-Zadeh
Assistant Professor Zhengya Zhang

Acknowledgments

It is impossible to thank everyone appropriately and keep this section shorter than the thesis itself as I am in a world of debt to many people for their help and support. A few shoutouts to my uncles Dean and Ken, aunt LaVonne, cousin Tyler, brother-in-law Duane, and parents-in-law Linda and Michael Chen.

First and foremost, I cannot thank my wife, Michelle (Elle) Roberts enough for being by my side the whole time. Neither of us knew what we were getting into when I agreed to join Prof. Wentzloff's research group, and there have definitely been bumps along the way, but if there was one constant it was Elle's support for me and for our lives together. She kept me grounded and never stopped reminding me of the bigger picture in life. It wasn't always easy and I didn't express my gratitude enough at the time, but I could not have done this without her.

In addition I have to thank my little buddy Jack. While he's wild, crazy, and unpredictable, he also had the unbound love and loyalty that only a dog could have. Hearing Jack bark and run downstairs to the sound of my keys in the door after a long day is something I will always cherish.

As for my family, Doug and Marlene Roberts and Nicole Hastings, I cannot have asked for any more than what they have provided. Not many people would support someone's decision to quit their stable job in the middle of a struggling economy (and sell his boat) in order to pay tons of money to attend a prestigious university on the hopes of finding an advisor before his funds run out. It was a gamble and it paid off. So many times I called my parents excited that a circuit

I designed actually worked when I thought it wouldn't and they would simply tell me they always knew it would. Turns out, even after all my years of education, they still know more than me.

Finally, I would like to send a special note to my future daughter. Someday, if you ever want to read your old man's thesis and come across this page, I hope you know that you are motivating me to finish up this chapter in my life because I cannot wait to start a new one with you…

It has been an incredible journey…and as my family knows well: the joy is in the journey.

TABLE OF CONTENTS

List of Figures

List of Tables

Abstract

The exponential growth in the semiconductor industry has enabled computers to pervade our everyday lives, and as we move forward many of these computers will have form factors much smaller than a typical laptop or smartphone. Sensor nodes will soon be deployed ubiquitously, capable of capturing information of their surrounding environment. The next step is to connect all these different nodes together into an entire interconnected system. This "Internet of Things" (IoT) vision has incredible potential to change our lives commercially, societally, and personally. The backbone of IoT is the wireless sensor node, many of which will operate under very rigorous energy constraints with small batteries or no batteries at all. It has been shown that in sensor nodes, radio communication is one of the biggest bottlenecks to ultra-low power design.

This research explores ways to reduce energy consumption in radios for wireless sensor networks, allowing them to run off harvested energy, while maintaining many qualities that will allow them to function in a real world, multi-user environment. Three different prototypes have been designed demonstrating these techniques. The first is a sensitivity-reduced nanowatt wake-up radio which allows a sensor node to actively listen for packets even when the rest of the node is asleep. CDMA codes and interference rejection reduce the potential for energy-costly false wake-ups. This radio consumes 116nW active power, which is less than the sleep power of a standard sensor node, while achieving a data rate of 12.5kbps and an energy/wake-up of 287.7pJ.

The second prototype is a full transceiver for a body-worn EKG sensor node. Due to size and energy limitations, this transceiver is designed to have low instantaneous power and is able to receive 802.15.6 Wireless Body Area Network compliant packets. It uses asymmetric communication including a wake-up receiver based on the previous design, a 4μW UWB transmitter and a communication receiver that consumes 292μW. The communication receiver has 10 physical channels to avoid interference and demodulates coherent packets which is uncommon for low power radios, but dictated by the 802.15.6 standard.

The third prototype is a long range transceiver capable of >1km communication range in the 433MHz band and able to interface with an existing commercial radio. A digitally assisted baseband demodulator was designed which enables the ability to perform bit-level as well as packet-level duty cycling which increases the radio's energy efficiency. At 2.5mA transmit and 378μW receive power, the transceiver is more than 20X more efficient than its commercial counterpart.

CHAPTER 1

Introduction

1.1. Computing Evolution: From Mainframes to Wireless Sensor Networks

1.1.1. Bell's Law

On October 14, 1984, *Computer* published an article titled, "The Mini and Micro Industries" in which Gordon Bell described a shift in the computer industry as the up and coming "micro-computers" were stealing away market share from the more traditional mainframes and mini computers [1]. Towards the end of his introduction Bell states, "In short, the forecast could be gloomy for mini companies. Just as the mainframe companies were unable to respond to the mini, the mini companies will have difficulty moving to meet the micro challenge because of their large install bases, proprietary standards, and large functional organizations" [1]. 30 years later a similar trend has shaken up the current computing landscape as regular PC vendors like Intel, Microsoft, HP, and others have struggled to respond to the new threat of smartphones and tablets. Players like Qualcomm, NVIDIA, Apple, and Google are poised to take their market share.

Bell's Law states that, "approximately every decade a new computer class forms as a new 'minimal' computer either through using fewer components or use of a small fractional part of the state of the art chips" [2]. He defines a computer class as, "a set of computers in a particular price range defined by: a programming environment such as Linux or Windows to support a

variety of applications; a network; and user interface for communication with other information processing systems and people" [2]. As Moore's Law continues, the 2x increase in transistor density every 24 months [3] enables new classes that have the same functionality as the previous decade, but overtake the older classes because of their lower price, which results in more use and therefore higher volume.

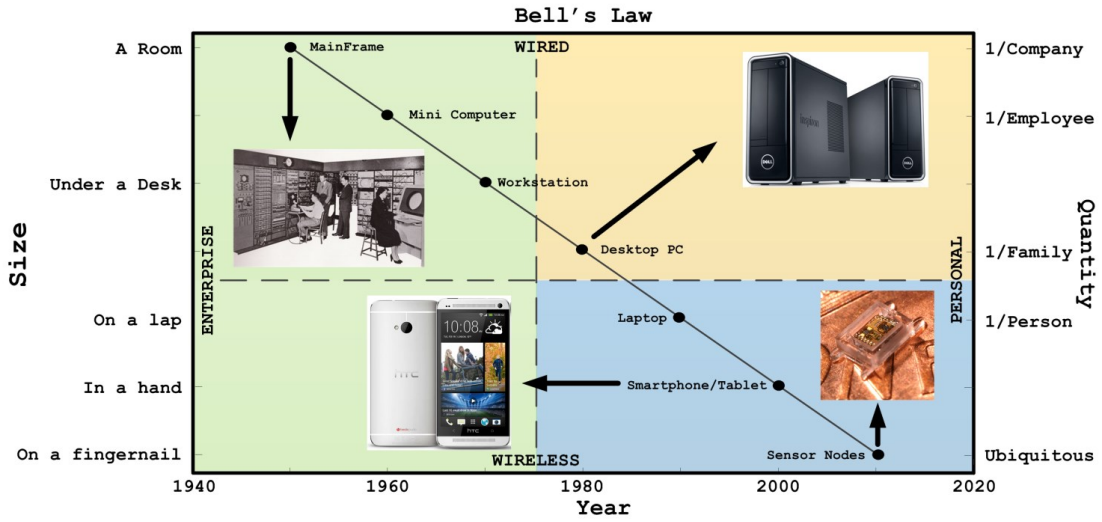

Figure 1: Bell's Law

Figure 1 shows Bell's Law in a graphical context. The y-axis on the left shows an abstraction of volume of each computer class, while the y-axis on the right shows an abstraction of the volume in sales. Mainframe computers in the 1950's were the size of a room and were so expensive they were limited to only one or two per company. Roughly 30 years, and three new computer classes, later the PC emerged, which was the first true computer class that appealed to the everyday consumer. Further evolution has yielded laptops in the 1990's and smartphones in the 2000's. Assuming Bell's Law continues to hold, the next computer class will be smaller and more prevalent than even smartphones are today.

1.2. The Internet of Things and Wireless Sensor Networks

The next computer class will consist of wireless sensor networks and will create the Internet of Things (IoT) era of computing. In 2008, the number of "things" connected to the internet exceeded the number of people living on the earth [4]. By 2020, the number of devices connected to the internet will reach 50 billion according to Cisco's CEO John Chambers [5], which he says is a conservative estimate. These "things" will change the way we interact with the world. For example, the "quantified self" is gaining significant traction in the marketplace with big, and traditionally non-high tech, companies like Nike, Under Armor, and Adidas all releasing wearable devices that monitor metrics like movement and steps and relate them to an energy profile. Mobile healthcare is another example that is receiving a lot of attention where patients can remotely monitor body vitals like heart rate and only communicate with the hospital if an issue is detected. Maybe less obvious are industrial use cases like cattle monitoring. BellaAg uses sensors to monitor the temperature of cattle and transmits the data to a smartphone/tablet. If a cow gets sick its core temperature changes, which can be sensed and alerts the farmer, providing early detection which not only gets the cow the treatment it needs but also helps reduce the risk of infection to the whole herd [6]. The IoT is a revolution that Cisco predicts will be worth $14.4 *trillion* dollars and is viewed as the fourth major evolution of the internet, following email, ecommerce, and social networking [7]. The backbone of the IoT is the sensor node, and their continued development will dictate whether IoT can achieve its lofty

predictions.

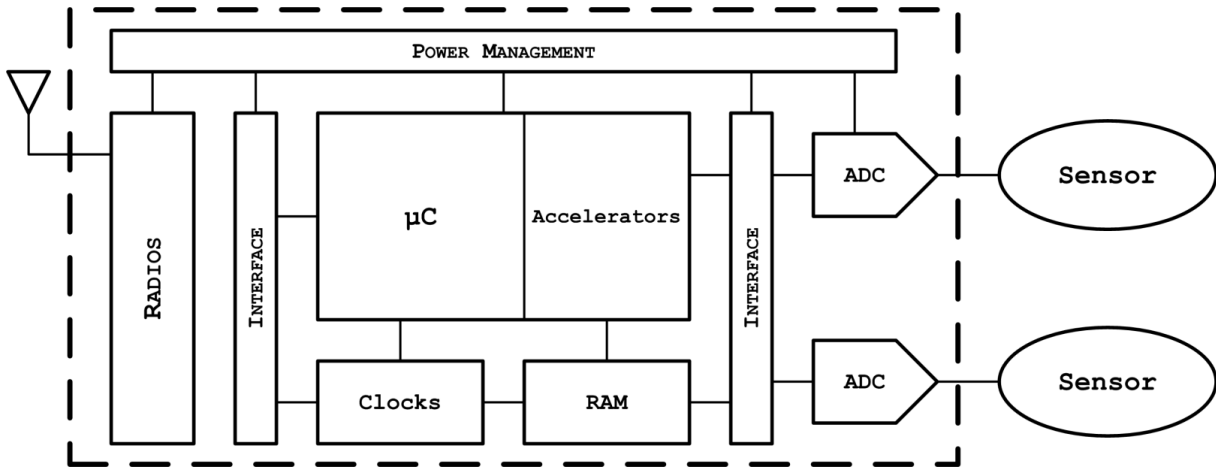

Figure 2: A Generalized Wireless Sensor Node (WSN)

1.2.1. Wireless Sensor Nodes (WSN)

Wireless sensor nodes (WSN) will power the IoT revolution. Taking a cue from Bell's Law, one can expect a sensor node to have the equivalent computing power of early cellphones at a fraction of the volume in size, and that is not too far off for present day WSNs. Figure 2 shows an example of a general WSN architecture. WSNs use one or more sensors to quantify the environment around them and then small energy-efficient microprocessors capable of complex computations process the data. Today, not every sensor has wireless capabilities because of the high power overhead, but increased energy improvement in standards like Bluetooth Low Energy, make wireless communication more prevalent. As broad as the concept of IoT is, the form factors of a WSN will be equally diverse. Contrary to every computer class that has come before, the majority of IoT devices will not conform to a general shape and size. Some may be large enough to contain a screen for interaction, while others will be small enough that it will not be noticeable by the average person. The following section will introduce a few examples of IoT devices available today and what they look like both inside and out.

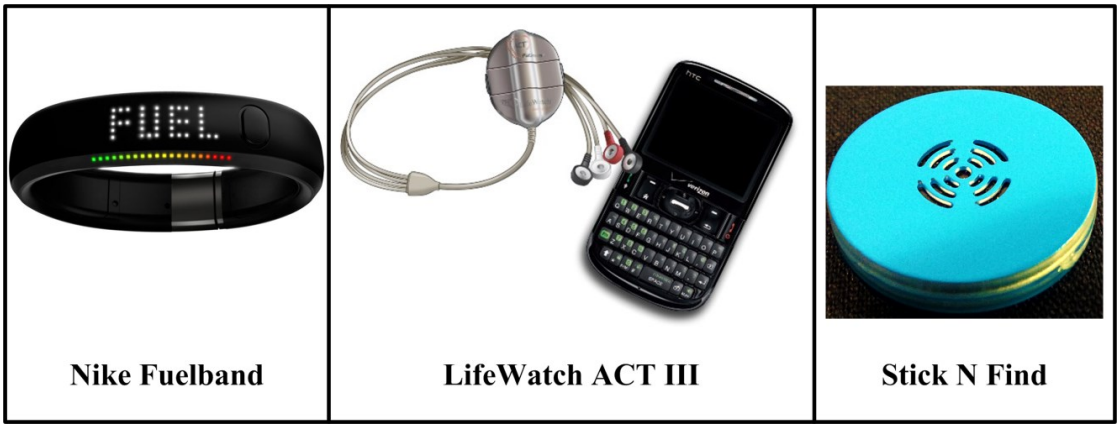

Figure 3: Commercial IoT Devices Today

1.2.2. Example Sensor Nodes

Figure 3 shows three examples of IoT devices that exist today.

Nike FuelBand

The Nike FuelBand is probably the most popular of the wearable fitness trackers today and fits within the profile of a WSN for IoT applications. The FuelBand tracks a person's movement through a 3-axis accelerometer and converts that movement into NikeFuel points so a person can monitor their activity level throughout the day. A teardown [8] of the device revealed that the FuelBand has 32MB of flash memory for internal storage and uses two different microprocessors, a TI MSP430 and an ST Micro ARM-based MCU. A Texas Instruments Bluetooth 4.0 radio and two batteries each with 36mAh capacity are included. Table 1 shows a breakdown of the significant chips in the design as well as their power consumption as reported from their respective datasheets.

Part Number	Description	Active Current consumption
Spansion FL032PIF	32MB Flash Memory	80µA standby mode
TI CC2560	Bluetooth 4.0	34.4mA Active Power
MSP430 F5528	Processor	290µA/MHz
STM32L	ST Micro MCU	214µA/MHz

	Battery	2 X 36mAhr

Table 1: Nike FuelBand's Major Components

The teardown revealed that running the LED bank led to an active current consumption of 23mA with a max current consumption of 60mA, with a 3.6V power supply. When in run mode (accelerometer is active and LEDs are off), the FuelBand consumes 180µA. Under these power consumption conditions, a single charge lasts for an advertised average of 4 days. Details on the radio duty cycling were unavailable, but the Bluetooth 4.0 chipset seen in the teardown has a 34.4mA active power according to its datasheet [9].

Hospital Heart Rate Monitor

The LifeWatch ActIII Platinum heart rate monitor is a state of the art WSN that is actively used in the University of Michigan Cardiovascular Center. Information on this monitor is not as readily available as the Nike FuelBand since it is not a commercial product and therefore no teardown is available, but we do know that it continuously monitors the heart using a 4-lead EKG and wirelessly sends the data via Bluetooth to a provided cellphone, which then sends the data over a cellular network to its parent company for analysis. The sensor, shown in the middle of Figure 3, does EKG processing on-node, but how much and what algorithms are used is unknown at this time. Empirical measurements of the node's lifetime reveal that the average lifetime of the sensor is 124 continuous hours from a single lithium AA battery with a 2600mAh rating, good for an average current consumption of 21mA.

Stick N Find

Stick N Find is a small coin-sized smart "sticker" that can be used to keep track of objects through a phone app and is a good representative of what IoT could look like going forward. Placing these stickers on important objects like a computer or attaching them to key rings allows a person to track those objects and an integrated buzzer and LED helps a user locate their objects

if they go missing. In addition they can be used as a "virtual leash" for a pet. If the pet moves out of range of one's phone then an app will alert them. These devices have a 0.98" diameter and are 0.16" thick, making them pretty unobtrusive. The sticker uses a BLE radio, but does not identify which radio is used. Stick N Find advertises a battery lifetime of one year using a standard CR2016 lithium coin cell battery [10]. A CR2016 from Energizer has a capacity of 90mAh [11], which puts the average power consumption at 10.2μA. With a 3.0V supply the average power is 30.6μW. Assuming the radio is strictly BLE and not Bluetooth 4.0, the Nordic Semiconductor nRF8001 serves as a good example [18]. With an active receive power of 14.6mA and an active transmit power of 12.7mA at 0dBm, then the duty cycle rate will be less than 0.08% for radio communication. Table 2 summarizes these three IoT devices.

Device	Function	Sensors	Battery Capacity	Lifetime
Nike FuelBand	Fitness	3-axis Accelerometer	72mAh	96 hours
ACT III Monitor	EKG	EKG	2600mAh	124 hours
Stick N Find	Tracking	None	90mAh	1 year

Table 2: Commercial IoT Device Summary

These sensor nodes provide a brief glimpse of some of the form factors and applications that IoT devices are integrated into today. Moving forward, while the sensor nodes will change, their basic structure will remain the same: sense the environment, process some of the data on node, and wirelessly communicate with the outside world. The growth of these devices is exponential.

1.2.3. Batteries

A major aspect of the IoT revolution is the sheer volume of WSNs that will be deployed. Referring back to the ActIII heart rate monitor, it only lasted 5 days on a lithium AA battery and averaged 21mA. If Cisco CEO John Chambers is right on his prediction of 50 billion active IoT devices, one must wonder how all of those devices will be powered. Assuming the average IoT

device develops to the point where it can survive for 10 years on a single battery, which would mean we need to change nearly 14 million batteries every day; or 158 batteries every second! Battery technology historically grows at a rate of 10% per year on average [12], which is much slower than Moore's Law, so the burden of IoT falls onto energy efficient circuit design.

1.2.4. An Academic WSN

Turing to academia, a roughly equivalent EKG sensor node to the previously mentioned Act III heart rate monitor used at the UM Cardiovascular Center has been designed as a joint collaboration between the University of Washington and the University of Virginia [13]. This sensor node measures a similar 4-lead EKG signal, but does so from harvested thermo-electric energy instead of requiring a special AA battery. In order to use harvested energy, the sensor node has an average current of only 19μA. Figure 4 shows a block diagram of the sensor node, which is similar, but more detailed, to that of Figure 2. Several advanced circuit design techniques were employed to reduce the sensor node from an average of 63mW for the ACT III to the 19μW here. Digital design in the subthreshold domain had a major impact as well as sophisticated duty cycling and efficient power management schemes both in the processing and wireless communication. The WSN had a transmitter, but no receiver (the Act III has both transmit and receive capability).

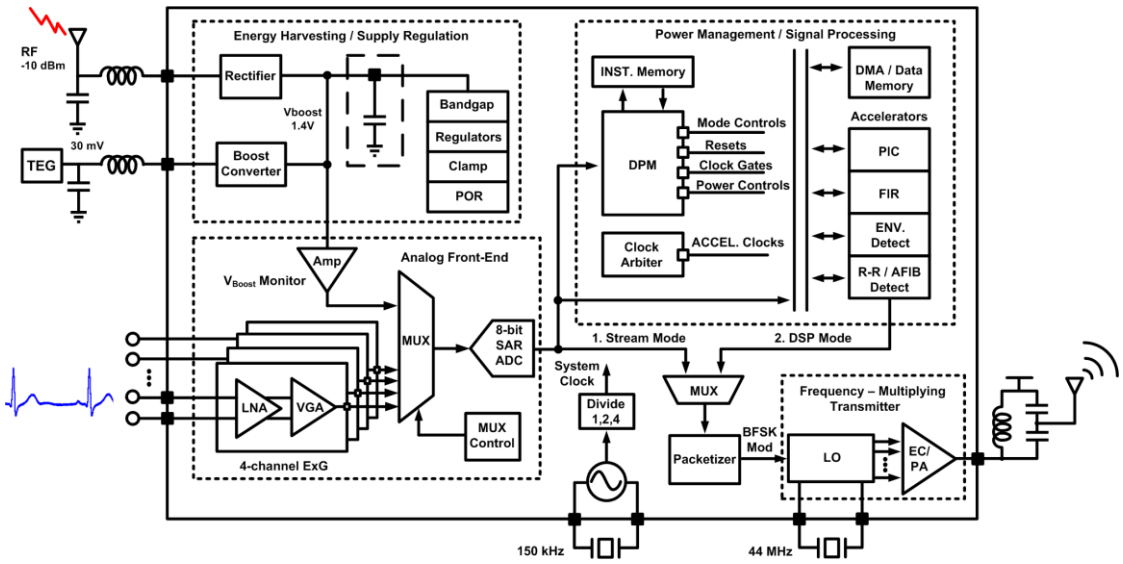

Figure 4: UW/UVA EKG Sensor Node [13]

1.3. Radios in a WSN

Looking at the Nike FuelBand teardown in Table 1, one can see that the radios consume the largest portion of active power by a significant margin. To have a clear reference point going forward of what is state of the art, the academic WSN presented previously is broken down by power and combined with the lowest reported receiver [14], its own transmitter [13], and wake-up receiver [14] to date from other universities and shown in Figure 5. On the left the reported active power consumption of the major blocks in the node is shown. It is obvious that the transmitter and receiver dominate the power budget. Duty cycling is employed to keep these high power circuits off as long as possible with the wake-up receiver staying on for synchronization. To save power, say one sets the duty cycling to 1%, as shown in the figure to the right. Notice that the wake up radio remains active while the node is asleep because it has to be listening in order to alert the node when an outside source wants it to wake up. The transmitter and receiver are no longer the energy bottleneck, but now the energy burden falls

onto the wake up radio. Either way, the radios dominate the energy budget of a state of the art WSN.

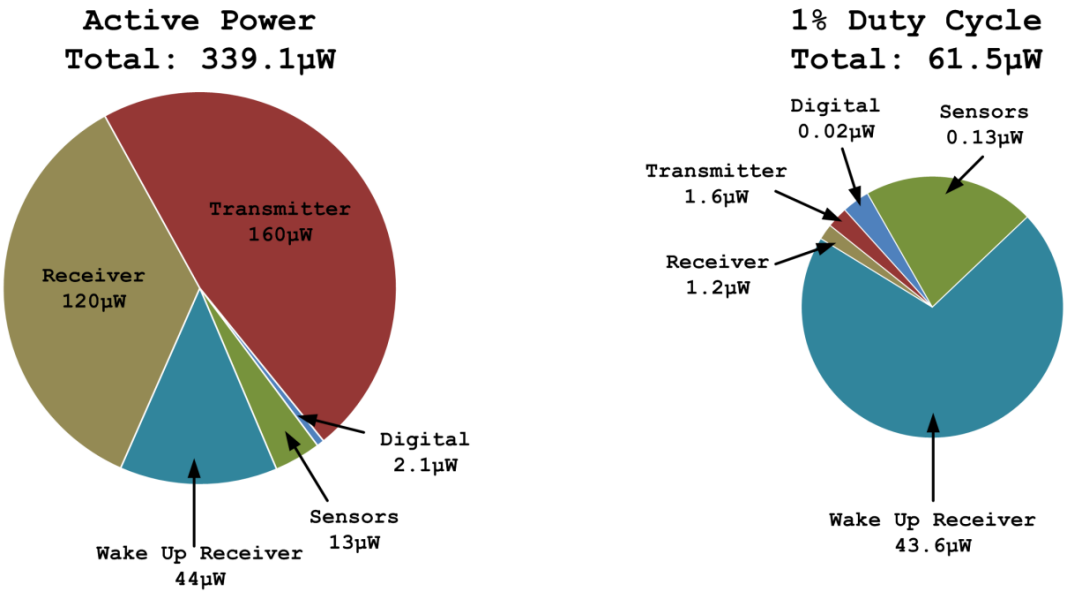

Figure 5: Power Breakdown of Hypothetical State of the Art Sensor Node

There are many reasons why radios consume higher power than other components in a WSN. The main reason is the need to amplify signals at RF frequencies. Beyond that, however, many subtler issues increase the radio's power. Noise performance is an issue for high sensitivity radios and one can make the general statement that noise is improved through higher power circuits. Interference is another concern in radios that lead to techniques like channel and frequency hopping which creates the need to high frequency and high performance PLLs.

Modulation schemes like coherent demodulation can improve performance and/or spectral efficiency at the cost of higher power versus simple modulation schemes like non-coherent energy detection.

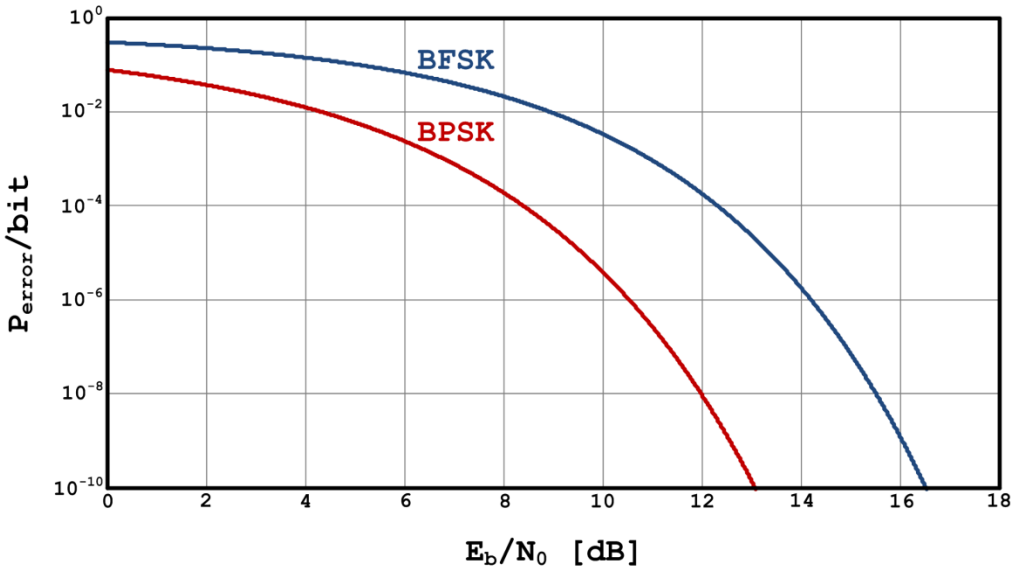

Figure 6: Incoherent (FSK, Energy Detection) vs. Coherent (PSK) Demodulation

Figure 6 shows the BER waterfall curves of standard BPSK (coherent) and BFSK (non-coherent) modulation schemes. BFSK requires a greater energy per bit to noise power ratio, but is simpler and requires less power. Many of these design decisions are not in the hands of RF designers, but instead dictated by industry standards. In the next 2 sections an overview of the most common, IoT applicable, short range standards both commercially and from academics is presented.

1.3.1. Commercial Wireless Communication Standards

Many RF design decisions are dictated by commercial standards that RF designers must follow, even at the expense of additional energy consumption. The following is a brief overview of three popular standards used for IoT devices:

Bluetooth Low Energy

The entire Bluetooth standard is too much to cover in its entirety, so this section is focused on the aspects of Bluetooth that will be most prevalent to IoT applications and to low power

radio design. As of 2013, Bluetooth Low Energy (BLE) is the low energy protocol with the most traction due to Bluetooth Classic's wide acceptance [15].

BLE reduces radio power from Bluetooth Classic by easing the channel and modulation requirements. BLE splits up the 2.4GHz ISM band (2400MHz to 2483.5MHz) into 40 channels each with a 2MHz bandwidth. The modulation format is GFSK with a modulation index of 0.5 and a bit rate of 1Mbps. FSK provides a low power non-coherent modulation while the Gaussian shaping improves spectral efficiency. Power is also reduced through link layer performance that focuses on long duty cycles and synchronous transmission of bursty data, which allows the radio to remain turned off as long as possible.

A BLE radio spends the majority of its time in the low power STANDBY state. In order to know if other radios may be trying to communicate with it, the radio will enter the SCANNING state where it will listen for other transmissions for a few moments before returning to the STANDBY state. If a BLE radio wants to try and contact another radio then it will enter the ADVERTISING STATE. From the ADVERTISING state, the BLE controller can either broadcast information to multiple receivers, listen for another radio trying to communicate, or enter into a 1:1 link with another BLE radio. The interval between consecutive ADVERTISING states can be as short as 20ms and as long as 10.24s, depending on the energy requirements of the application. Communication over SCANNING/ADVERTISING is restricted to 3 of the 40 available channels; the 3 channels are outside the most common Wi-Fi channels in order to avoid interference. If a 1:1 link is requested then both devices enter into the connection state, either through the ADVERTISING state or the INITIATING state. The main takeaway from the BLE link layer is that it saves energy by emphasizing quick connections and long times spent in STANDBY. For example, the Nike FuelBand's TI CC2560 BLE chip consumes an active power

of 34.4mA, but when it is in a connected Master role with a 500ms connection interval the current consumption goes down to 169μA [9].

802.15.4 Wireless Personal Area Networks (WPAN)

Zigbee is the commercial protocol that utilizes 802.15.4 WPAN standard. There are many different frequency bands and modulation schemes associated with 802.15.4, but the focus will be on the 2.4GHz ISM band since that is by far the most common commercially [16]. Across the 2.4GHz ISM band Zigbee is divided into 16 individual channels with 5MHz bandwidth each. The modulation format is OQPSK with half sine pulse shaping and a bit rate of 250kbps.

Zigbee employs Direct Sequence Spread Spectrum (DSSS) which allows for very low duty cycling without the need for close synchronization. For multiple access and collision avoidance, instead of frequency hopping like BLE, Zigbee uses clear channel assessment (CCA) to decide whether or not to communicate in a particular channel. CCA is used to determine the current status of the channel. The simplest form is to use an energy detector (ED) circuit to monitor the channel and determine if the signal energy is above a certain threshold. CCA has a carrier sense mode where it determines if an 802.15.4 compliant signal can be detected. The third common CCA mode does both energy detection and carrier sense. Avoiding noisy channels is a method commonly used to avoid the need for extra power in the hardware.

802.15.6 Wireless Body Area Network (WBAN)

802.15.6 is the newest standard of the three mentioned here and it is specifically targeted towards medical applications and wearable electronics, hence the name Wireless Body Area Networks (WBAN). Similar to 802.15.4, the standard covers multiple frequency ranges and modulation techniques, but again the focus is on the 2.4GHz ISM region. However it should be noted that WBAN does have a specification for a newly released Medical Body Area Network

(MBAN) in the 2360-2400MHz "quiet" band. This will accompany the other common medical band, the Medical Implant and Communication Service (MICS) from 402-405MHz [17].

In the 2.4GHz ISM band, WBAN has 80 channels at 1MHz bandwidth and uses a π/2-DBPSK modulation with square-root raised cosine and an unspecified modulation index. The data rate is 187.5kbps. WBAN contains a Burst Mode which indicates whether the next packet is part of a packet burst. WBAN also contains a CCA mode similar to that of WPAN [17].

Table 3 briefly summarizes the three commercial radio standards that were presented:

STANDARD	CHANNELS	CHANNEL BW	MODULATION	LOW POWER TECHNIQUES
Bluetooth LE	40	2MHz	GFSK	Pseudo-Freq Hopping
WPAN	16	5MHz	O-QPSK	CCA
WBAN	80	1MHz	π/2-DBPSK	CCA

Table 3: Summary of Commercial Low Power Radio Standards

Diverse techniques are used in these standards to reduce power consumption. Increasing the number of channels helps with interference because it gives the radio more chances of finding a quiet channel, but more channels often means smaller channel bandwidth which requires better performance from the hardware…and more power. Modulation also trades off between spectral efficiency and power. Modulation formats like BLE's FSK is a simple modulation, but less spectrally efficient than PSK. However, PSK requires knowledge of phase information, which requires more power. Lastly, techniques like CCA help avoid noisy channels where a radio would have to consume more power to ensure reliable communication, while the pseudo-frequency hopping allows for long off-times while maintaining synchronization.

1.3.2. Low Power Academic Radios

Radios published in academia come in two distinct flavors: those who conform to a commercial standard like Bluetooth or Zigbee and those who ignore any standards. The low power radios that ignore standards have the lowest power since they are able to avoid many complexities involved in being standard-compliant [19]-[54].

Energy Detection

Radios that do not conform to standards use various forms of energy detection techniques. Similar to the hardware used in CCA for both WPAN and WBAN, these radios detect the presence of signal energy to determine if a '0' or a '1' has been sent. Several different flavors of energy detection exist like: on-off keying (OOK) where information is decoded by detecting signal energy above a given threshold in a channel [24], frequency-shift keying (FSK) where signal energy is present in one of two possible frequencies [14], and other techniques like super-regeneration where an oscillator is kick-started by the amount of signal present at the receiver. Based on the speed of the oscillator startup time, the amount of signal energy can be detected and decoded [20].

While these radios can be very low in power, they are subject to several drawbacks. First, even the lowest reported 120μW communication receiver listed in the hypothetical WSN in Section 1.3 dominates the power budget of a sensor node. Second, the energy detection technique, while low power, is susceptible to interference, which makes them un-usable in the real world without some power-increasing techniques like multiple channels or code division multiple access.

Multi-User Access

Most academic radios work well in a lab on a test bench, but if multiple radios are deployed in a multi-user environment, their energy detection architectures would interfere with each other and communication quality would suffer. Techniques like frequency, time, or code division multiple access (FDMA, TDMA, and CDMA, respectively) are needed so that multiple users can communicate at the same time using the energy detection architectures.

1.4. Radio Power Reduction Techniques

Many of these techniques have been mentioned in previous sections, but this next section will summarize some of the common techniques used to reduce the power in radios. All these techniques come with drawbacks of their own, but when power is critical, these are often employed.

1.4.1. Duty Cycling

Duty cycling, the process of turning on and off an active circuit so that it only consumes power when it is needed, is the most common method for reducing power in radios, both in industry and academia.

BLE implements duty cycling into the standard by having packets transmitted followed by long periods of silence before sending/receiving more information. When in a direct link with another BLE radio, up to 4 seconds can pass between packets [18]. When not communicating, the power hungry front end receiver and transmitter are turned off, while the baseband processor keeps track of timing to ensure communication between two links can be easily re-established when the time is right. The energy required to preserve the timing locally is less than the energy to keep the radios in constant synchronization, so this produces an energy win. Low duty cycling in BLE, like talking for about 1ms every 4 seconds can bring the current consumption down from 14.6mA to 4μA in a 4 second connection interval [18].

1.4.2. Simplicity

The use of simple radios can be seen mostly in academic papers, but various forms also in standards like 802.15.4 and 802.15.6 for the use of clear channel assessment (CCA). These simple radios detect signal energy, at its most basic, and/or appropriate packet structure to determine if a certain channel is uncongested enough for efficient communication. The advantage is that the radio knows ahead of time if noise/interference will be a problem instead of wasting energy transmitting packets that might fail to be received. CCA is accomplished using a received signal strength indicator (RSSI) which is a simpler receiver that detects the strength of the signal instead of demodulating any information. If the signal strength is low then the CCA determines that communication in that given channel will be efficient and will allow the main radio to proceed.

Academic radios, as has been discussed, also use simplicity to reduce power. Energy detection schemes are prevalent in low power radios for their circuit simplicity. Wake-up radios are another example of using a very simple radio for communication while keeping the more complex and power-hungry radios off until absolutely necessary. In addition, exotic passive off-chip components such as FBAR resonators [37] or BAW resonators [21] provide impedance matching, filtering, and voltage boosting which eases the performance requirements of the on-chip active circuits. Commercial radios do not often employ these off-chip components due to cost constraints.

1.4.3. Sensitivity

A radio's sensitivity dictates the smallest signal that it can correctly detect and demodulate and it directly impacts the overall power consumption of the radio. A lower sensitivity number (more negative) means the radio can detect smaller signals. Improving sensitivity requires

amplification and good noise performance, most coming at RF or IF frequencies, which always means extra power. If the sensitivity is too low then communication becomes difficult and/or unreliable. Because of this, sensitivity traditionally is a specification and that has stringent requirements and often leads to over-design to ensure communication even in adverse conditions. However, with the advent of IoT devices, it is important to re-evaluate even some of the most traditional specs. With communication ranges shrinking (like on-body communication), and without the need to communicate constantly, like for cellular networks, sensitivity can become a useful knob from which we can tune power.

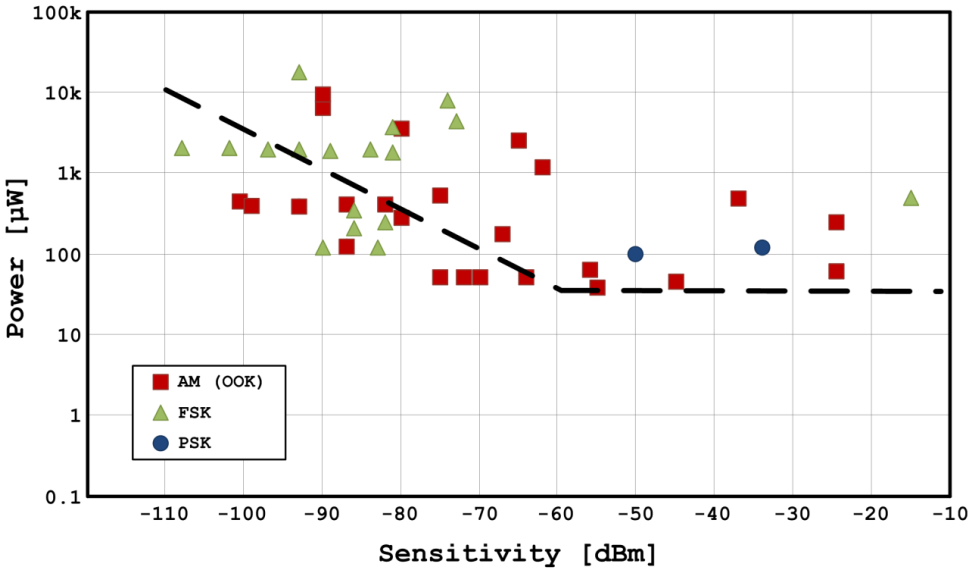

Figure 7: Low Power Radio Survey: Power vs. Sensitivity

Low Power Radio Survey

A survey of published ultra-low power radios was conducted to demonstrate the relationship between sensitivity and power [19]-[54]. Figure 7 shows this survey spanning major conferences and journals from 2006 to 2013, comparing the radios' power versus sensitivity. It should be noted that these are all custom ultra-low power radios that use some form of energy detection architectures and therefore commercial radios like Bluetooth and Zigbee are not shown here.

Also, note that this survey includes radios of different architectures, different frequencies, and different data rates; none of which is separated in the plot. With sensitivity, in dBm, on the x-axis and power, in μW, on the y-axis, two distinct trends can be seen. First, when looking at sensitivity greater than –60dBm (to the right on the x-axis) it can be seen that changing sensitivity does not impact the power of the receiver. This is mainly because at these sensitivity numbers the radio is not noise-limited and therefore does not need to spend extra power amplifying the signal. However, there is a floor around 50μW which suggests there is a minimum power required of the radio regardless of sensitivity. Increasing sensitivity from -60dBm (to the left on the x-axis) there is a linear trend with a slope of roughly -1/2, suggesting that there is a correlation between sensitivity and power. Slope-fitting from the plot, it can be seen empirically that a 20dB change in sensitivity results in a 10x change in power. For short range radios, this knowledge can be used to make key design tradeoffs.

Channel Modeling

Knowledge of the wireless channel between two radios is critical for making a sensitivity versus power design tradeoff and the frequency and amount of variation in the channel for IoT devices is different from more traditional models used in the past. Older and more traditional cellular networks focus on long range communication, and movements made within a small time frame appear essentially static to the channel. However, with certain IoT applications, like body to body communication, movement in a small time frame will have a dramatic influence in the range and position of two links. Understanding this channel and modeling is therefore important yet complicated and also provides an opportunity for lower power design. See Appendix A for more details on a channel modeling study that produces intuitive equations that account for channel periodicity in short range communication.

1.5. Contributions

Wireless sensor nodes may be the next step in Bell's Law, but at the present time they are not ready for the ubiquitous deployment many are predicting. Power consumption for the average node still necessitates the need for a battery, which will prevent WSNs from reaching the lofty multi-billion deployment predictions. Many of the energy constraints have been eased with the continual improvement of digital processing, leaving wireless communication as one of the largest remaining bottlenecks that prevent true energy autonomy. A popular area of research, many publications aim to improve this issue, but do so with significant tradeoffs. In addition to needing further power reduction and the ability to handle multiple users, most published radios are also susceptible to interference (both in-band and out-of-band) as well as jamming and have no way to overcome this due to their simple, power-saving architectures. The contributions I present below are aimed at breaking this energy bottleneck so that wireless communication using harvested energy is possible while maintaining the sophistication needed for the radio to be deployed in a multi-user environment. Power will be reduced using sensitivity reduction and duty cycling, while simple protocols schemes like FDMA and CDMA will be implemented to improve deployment success.

Design Techniques for Nano-watt Radios

As seen from the low power radio survey in Figure 7, the empirical floor for minimum power consumption in ultra-low power receivers is ~50μW, which was shown earlier to still make it the dominant source of energy in a heavily duty cycled WSN. Because WSN sleep mode power consumption is in the nanowatt range, there is a need for radios that consume nanowatt power as well. I will present a 98nW wake up receiver that utilizes the analysis from the dynamic channel modeling study to efficiently reduce sensitivity for short range communication in body area

networks. A novel subthreshold-biased active rectifier front end was designed to achieve the target sensitivity while maintaining the target data rate without the need for a low noise amplifier.

Using CDMA to Overcome Interference for Energy Detection Architectures

A common problem among low power radios that use energy detection architectures for demodulation is the impact of interference on bit error rate (BER). To resolve this many different protocols can be used ranging from very basic, like Aloha, to more standard techniques like TDMA or using codes. I will present a CDMA-based wake up radio that utilizes a similar architecture to the previous one to maintain nanowatt power, but this radio will add a baseband processor capable of decoding up to 4 parallel CDMA codes. The low-power clock generation circuits and subthreshold baseband processor are all included on chip. In the presence of interferers an automatic threshold feedback loop will modify the sensitivity of the receiver and overcome the interference.

Ultra-Low Power Coherent Communication

All reported academic radios use energy detection architectures, which are non-coherent, because they are simple and therefore low power. Coherent demodulation is difficult because it requires knowledge of the phase of the received signal, which requires more power hungry circuits like PLLs. Even so, there are advantages to coherent demodulation like improved SNR and spectral efficiency. I will present a MICS-band receiver that is compatible with 802.15.6 WBAN signals for coherent demodulation with an overall power low enough to sit within other non-coherent academic radios. This radio uses a low power PLL-based processor for demodulation while also using a PLL in the LO to allow up to 10 physical channels in the MICS band.

Transceiver Design for Batteryless WSNs

Two different transceiver designs will be presented; both of which allow for wireless communication on a sensor node that powers from harvested energy. Each transceiver will address in-band and out-of-band interference, jamming, multiple user support, and low power. However, the two transceivers serve very different applications.

The first transceiver integrates with an EKG sensor node used in body area networks and will keep active power low because harvested energy and energy storage will be limited due to size constraints. The second transceiver will be higher power, but achieve long range communication instead of the short range communication most other low power radios focus on. This transceiver will be compatible with a commercial low power radio, but with a 20x to 100x reduction in power.

CHAPTER 2

Nanowatt Receivers

2.1. Introduction

As stated in the Introduction chapter, there are several challenges to overcome before sensor nodes can be deployed ubiquitously and enable IoT at the magnitude that is being predicted. Receiver power and synchronization were discussed as energy bottlenecks and several different tools were offered that can be used to alleviate the energy burden from receivers. This chapter introduces two variations of a wake-up receiver (WRX) that breaks the energy barrier and introduces robust radio design at nanowatt power while maintaining the sophistication needed for ubiquitous deployment. A WRX is a receiver that has much lower power than a standard communication receiver and remains active while the rest of the sensor node enters a low power sleep state. Upon detection of a signal, the WRX wakes the node so it can perform necessary operations, whether that be collecting sensor data or turning on a higher power transceiver for communication.

Section 2.2 will introduce the original version of the nanowatt wake-up receiver (WRX). Section 2.3 will introduce the second version of the wake-up receiver (CDMA WRX) which adds multiple access and robustness to the original through the use of CDMA codes and automatic threshold control logic. Section 2.4 and 2.5 will discuss my individual contributions to the WRX project and provide acknowledgements.

2.2. A Nanowatt Power Wake Up Receiver

Heavy duty cycling in WSNs is necessary for ultra-low power consumption, but results in the nodes spending most of their time in a low power sleep state. Several different methods can be used to take the node out of sleep, including sensor interrupts, timers, or through the radio. While sensor interrupts can be low power, it does not give a user much control if they need to wake up the node to interact with it. Timers can also be low power [55], but result in a periodic wake-up and become inefficient if the node wakes up too often for no reason. A WRX, being active all the time, is a good solution, because it allows a user to wake up the sensor node, but only wakes the node up when it needs to perform some specific function, maximizing the duty cycle efficiency. However, with the sleep power of sensor nodes in the single digit μW range, due to subthreshold leakage in the SRAM and digital logic, even state of the art wake up receivers [21],[23] dominate the overall power budget while listening for packets. In order to be practical, the wake up receiver's active power should be below that of the sensor node's sleep mode.

Looking at the low power radio survey in Figure 7 in Section 1.4.3, the power floor is ~50μW because of the need for RF transconductance which is used for noise performance and to extend communication range. However, because many IoT applications do not require long communication range, RF transconductance is not always necessary and provides an opportunity to significantly reduce power in a wake-up receiver. The next two sections will introduce and motivate the concept of reduced sensitivity RF communication followed by the original nanowatt WRX.

2.2.1. Introduction

The continued push towards smaller WSNs has created the opportunity to use these sensors in novel applications previously unrealizable. One such opportunity is the creation of body area networks (BAN) where a single sensor or multiple sensors unobtrusively monitor the biological functions of a person and relay the information to an aggregator, such as a cellphone, or base station. Such devices will allow doctors to remotely monitor patients and also provide them with the ability to detect potential health-related complications earlier than before [56].

In order to improve the power and energy efficiency of the radios in WSNs even further, the applications themselves need to be exploited. For body area networks, a person's body is always in motion, at least a little bit, and knowing the dynamics of that channel can lead to optimized receiver design [57]. For the design of this receiver, the results from a characterization study of the wireless channel between two nodes on a body were used to reduce the receiver's sensitivity to an optimum point in order to reduce power while maintaining efficient communication. Using the targeted sensitivity and the low power performance of weak inversion devices, the receiver can operate at nanowatt power levels.

2.2.2. Motivation

Channel Path Loss

Short range communication devices, specifically for BAN, differ from more traditional long range communication devices in two significant ways. First, as seen in Figure 8, small movements between two nodes a long distance away do not seem significant over short time periods. For a BAN and other similar short range applications the distance and position of the nodes relative to each other changes significantly in short time periods. Second, the amount of path loss between nodes at long range is much more significant than the path loss between nodes

that are close together. Because of these changes, receivers that are suitable for longer range communication might not be optimal, or may be over-designed, for short range applications. To see this in more detail the wireless channel for a BAN needs to be analyzed.

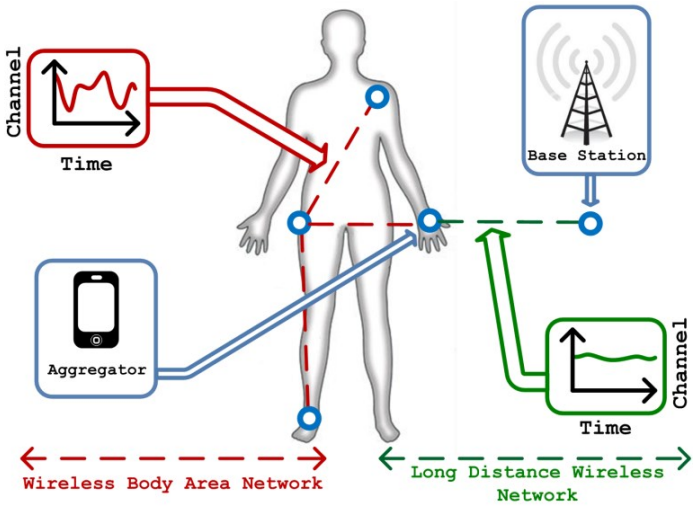

Figure 8: Long Range vs. Short Range Communication

The channel for wireless BAN involves the environment and electromagnetic propagation around the human body, which is complicated due to the body's effect on antenna performance [58]. One reason the channel for a BSN is unique compared to other channel models is the channel's dynamic characteristics due to the motion of a person's body and the proximity of the sensors to the body itself. The distance and angle of antennas mounted to sensors worn on the body will constantly change in relation to one another as a person performs daily activities. Intuitively, the channel should oscillate between strong and weak conditions as the person walks or runs. Furthermore, these oscillations should be bounded by the physical limitations of the body's movement or lack of movement (i.e. the human body is never perfectly still) [57].

Intermittent Communication

For most applications of wireless sensor nodes, streaming data is not necessary and therefore the radios are not required to be in constant contact with an aggregator or another node in the

network. This knowledge can be used to design radios that do not need to communicate in the entire channel. To better understand the dynamics of the channel, custom hardware using discrete components was designed and recorded RSSI between two sensor nodes at a rate of 1.3kS/s, which is then converted to path loss [57]. Figure 9 shows an example of such data. The blue periodic waveform in the background was taken with a user wearing a sensor on their chest and their wrist while running. One can see the periodic movement that would be expected from such a repetitive exercise. Also note the peak-to-peak variation and average path loss in the figure. If it is assumed that the transmitted power and antenna gain are 0dBm and 0dB, respectively, then the sensitivity of the receiver is equivalent to the measured path loss. Using Figure 9 as an example, sensitivity of the radio can be reduced to -40dBm and still be able to communicate almost half the time.

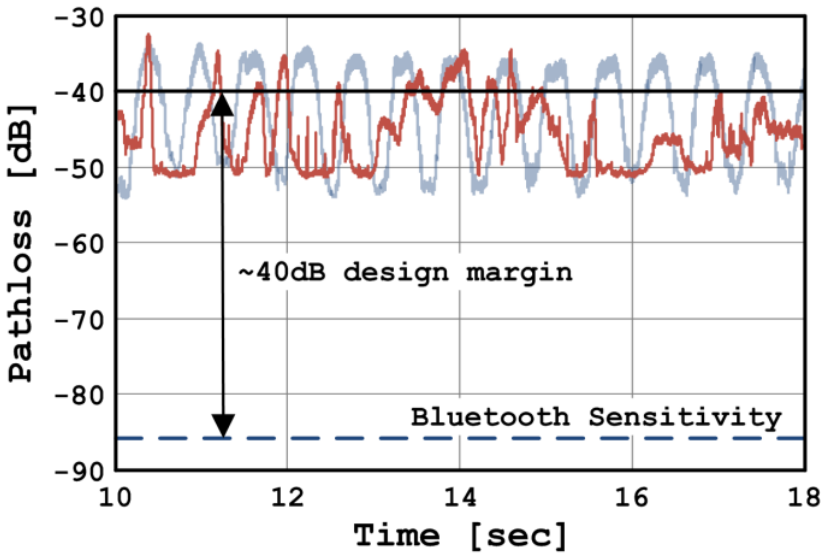

Figure 9: Short Range Communication Channel Variation

Realistic Channel

The background waveform shown in Figure 9 is a controlled example, where channel periodicity is guaranteed. A more realistic channel can be seen in the red foreground waveform

in Figure 9 which shows path loss measured between a sensor on someone's wrist and hip while they played tennis. Also on the plot is the average sensitivity of a Texas Instruments Bluetooth receiver [59]. As can be seen, for this specific application the Bluetooth receiver is significantly over-designed. Reducing the sensitivity of the receiver by 30dB would allow for reduced design constraints in the receiver while still allowing for communication in the entire channel. Allowing for intermittent communication, design constraints can be relaxed even more. According to the channel model, if the sensitivity of the receiver is set to -40dBm, same as for the running example, then there will be many instances in time where communication is possible, and when possible, the receiver can communicate while realizing the power benefits from the sensitivity reduction.

Target Specifications

Having established that reduced design constraints are realizable through sensitivity reduction, and using our understanding of the sensitivity vs. power trade-off from Section 4.3 in the Introduction chapter, we can target an optimal sensitivity that will balance the need for communication along with the need for low power. Looking at the -40dBm sensitivity region chosen through the running and tennis examples, and knowing how dominant the low power receiver is to the overall average power in a sensor node, a receiver with a power consumption that falls below the floor at -40dBm sensitivity is the target of this work.

2.2.3. System Analysis

A block diagram of the ultra-low power receiver is shown in Figure 10. A 915MHz OOK modulated signal is received at the input of the antenna at a data rate of 100kbps. The input signal is then boosted by the off-chip matching network before going on-chip to the gate of the active rectifier. In the presence of an incoming signal the active rectifier performs down-

conversion and pulls its output voltage, V_{active}, low. The other half of the rectifier circuit creates a stable bias voltage, V_{bias}, which is used as a reference voltage for the hysteretic comparator. Based on the voltage difference between V_{active} and V_{bias} and its own hysteresis settings, the comparator will toggle either high or low.

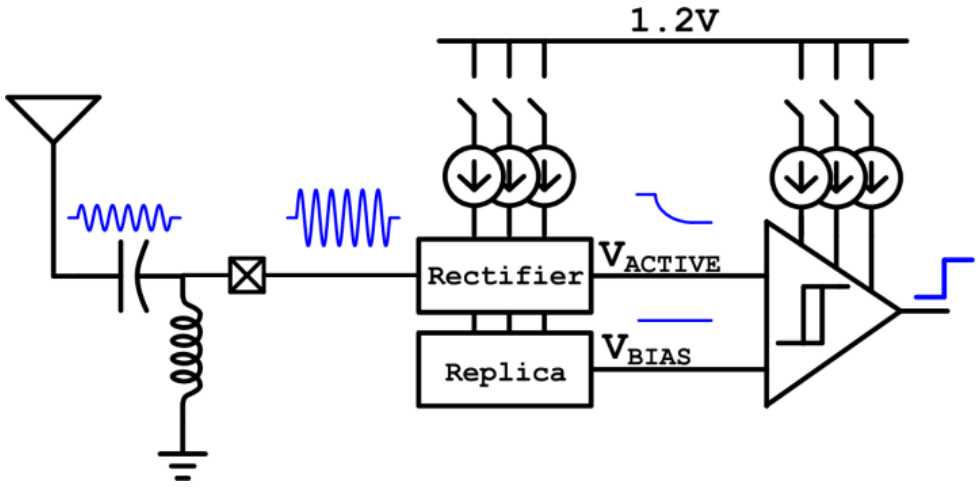

Figure 10: WRX Block Diagram

Off-chip Resonant Tank

For an input sensitivity of -40dBm, the input signal at a 50Ω antenna will be around 2.2mV. The input signal coming into the rectifier is desired to be as large as possible to maximize the rectifier's conversion gain. Since RF transconductance cannot be used due to power constraints a power match is implemented using a series capacitor and shunt inductor resonant tank to boost the received voltage. Before implementing the resonant tank the rectifier's measured input impedance was 0.7-j18.8, so a 1pF capacitor and 5.5nH inductor was used and measured results show the sensitivity of the receiver increased by 12dB.

Self-biasing Rectifier

The design of the rectifier is critical because it serves as the RF front-end of the receiver. Because of this the rectifier has to have high enough sensitivity to detect the desired signal and

fast enough speed to achieve the data rate specifications, all while doing so at nano-watt power levels.

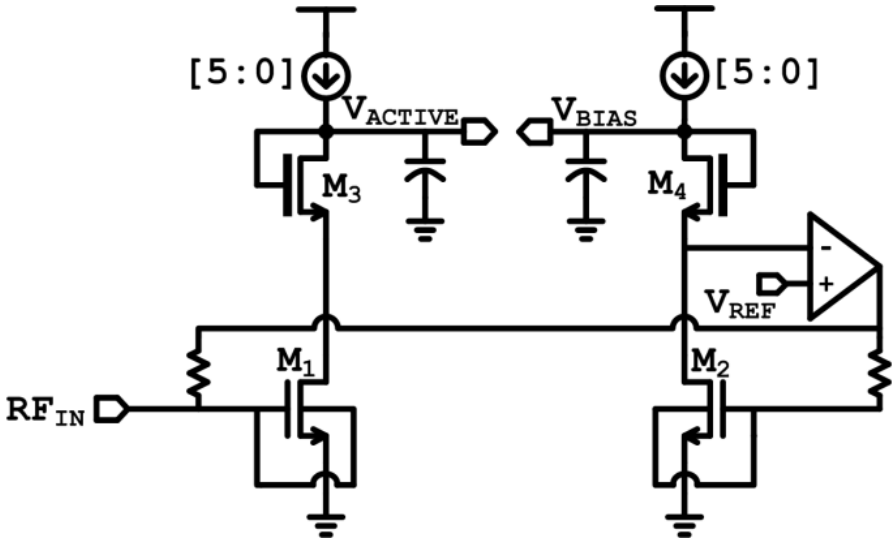

Figure 11: Self-biasing RF Front End

The 23nW rectifier architecture is shown in Figure 11. It begins with a common-source connected device, M_1, in series with a diode-connected device, M_3, with a current load on top. M_3 acts as a level shifter so the output of the active rectifier falls within the dynamic range of the comparator. The series connection of the current load, M_3 and M_1 is then mirrored over to create a replica bias (M_2 and M_4). Without an input signal connected to the replica bias, the output voltage will remain stable. From here the drain voltage of M_2 is sensed and compared with an on-chip voltage reference which is then used in feedback to bias the M_2's gate. This bias voltage is then fed forward to M_1. The current sources used on both the active and bias rectifiers are binary-weighted current DACs with 6 bits of precision in order to control the power consumption over process variation.

The design of the rectifier self-biases M_1 in deep subthreshold. Because the region of operation for a MOSFET in subthreshold is mostly dependent on M_1's drain to source voltage, the off-chip voltage reference is an effective way to ensure proper biasing so that the drain

voltage is very close to the minimum headroom of M_1. In the presence of an incoming signal, M_1 is biased at a point where the signal will cause M_1 to pull its drain voltage down far enough to be detected by the comparator, but not too far, so it can do so very quickly. When the signal disappears the current load will pull M_1's drain voltage back high, resetting the rectifier.

To achieve the sensitivity target M_1 was biased in deep subthreshold in order to maximize its g_m/i_d ratio; maximizing transconductance for the current consumed. M_1 is also connected in a DTMOS configuration by connecting its body to its gate. This will modulate V_{th} as the gate voltage changes, improving conversion gain and therefore sensitivity [60].

The rectifier was also designed for speed with reduced current. Knowing that the rectifier response time is controlled by the change in voltage multiplied by the load capacitance all divided by the current through the rectifier, the circuit can be optimized. Current is minimized as much as possible to reduce power without sacrificing g_m due to the deep subthreshold operation. The product of the load capacitance and the change in voltage is also minimized to keep the response time within spec. The change in voltage has a lower limit set by the comparator offset. Simulations show that a current as low as 5nA can produce the desired speed if mismatch is neglected.

On-chip Voltage Reference

The on-chip voltage reference used to ensure the biasing of the rectifier is shown in Figure 12. The design of the rectifier allows for slack in the voltage reference output voltage so negligible power consumption was prioritized over temperature invariance. To achieve this the CTAT source to gate voltage of M_1 is used and applied to the gate of M_2. This creates a PTAT current through the cascade current mirror and back to the source of M_6 [61]. Typically, M_3 would be diode connected and produce a PTAT current, but tying the

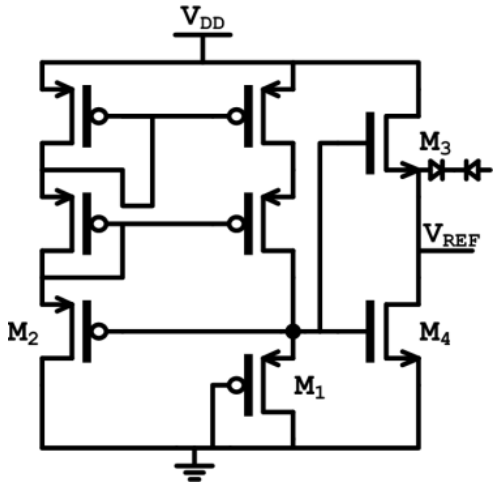

Figure 12: On-chip Voltage Reference

gate of M_3 to VDD gives it more drive strength and M_3 cannot be made sufficiently long to properly reduce its drive strength in relation to the rest of the circuit. Instead, the gate of M_3 is connected to the gate of M_2 and M_4. This change reduces the voltage reference's temperature independence, but also significantly reduces power. Figure 13 shows the measured temperature performance and simulated power of the reference.

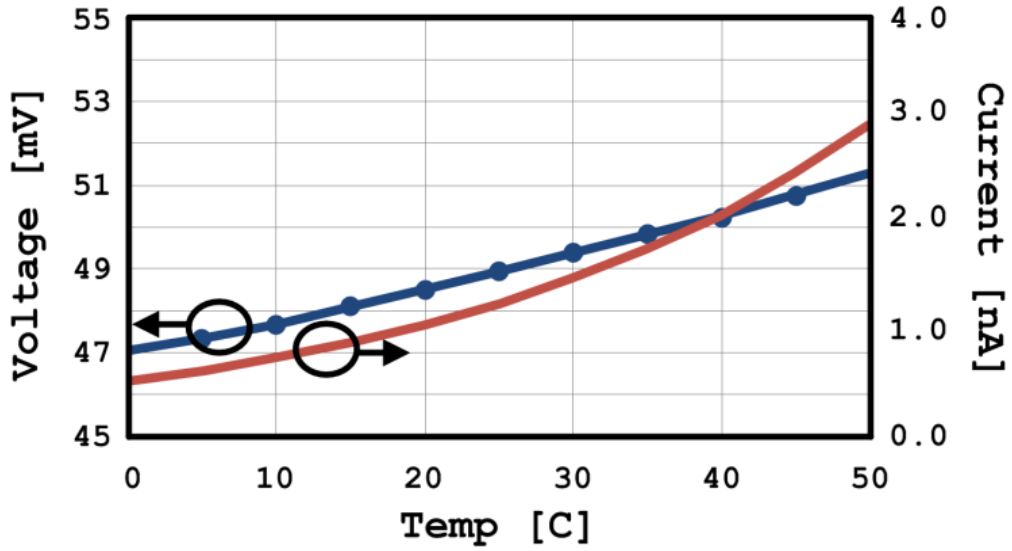

Figure 13: Voltage Reference Temperature Performance (Measured) and Current (Simulated)

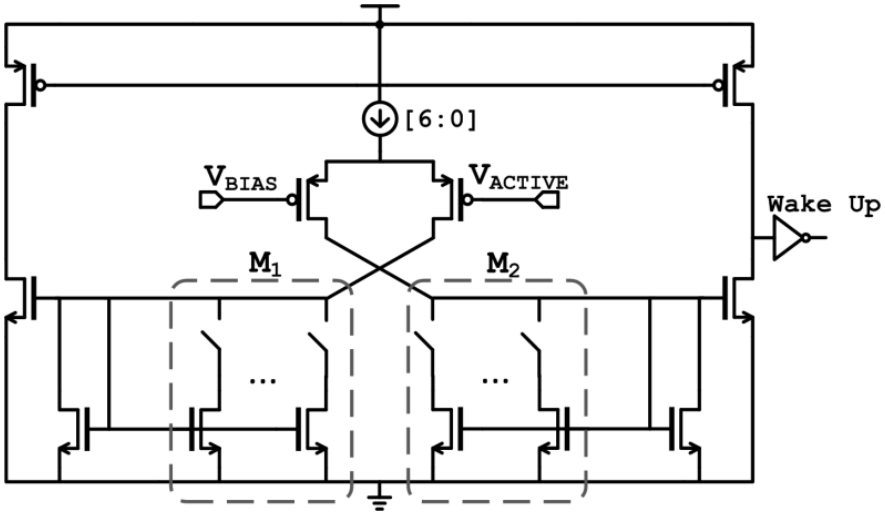

Figure 14: Hysteretic Comparator

Hysteretic Comparator

A 28nW hysteretic comparator [62], with its input devices biased in subthreshold, is used to compare the active and replica outputs of the rectifier. Programmable asymmetric hysteresis is implemented to help offset the impact of mismatch. The hysteresis can be programmed using a 4-bit binary-weighted calibration scheme. As seen in Figure 14, device M_1 is replaced with parallel devices of varying sizes with switch headers which create the desired hysteresis levels.

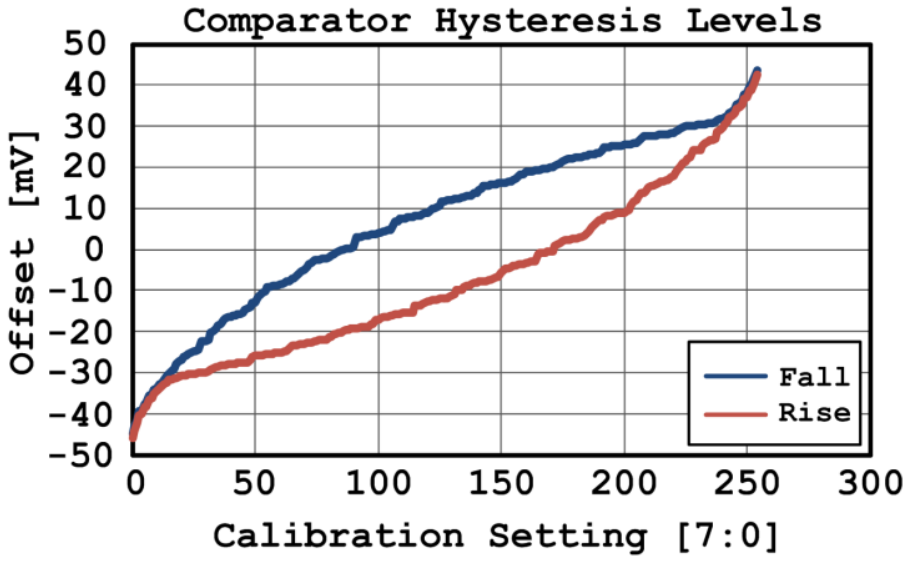

Figure 15: Programmable Hysteresis Settings in Comparator

Binary switching among the four devices on each side changes their effective width, and therefore the input offset. Adding more parallel devices will allow for a greater hysteresis tuning range, but will also increase load capacitance, which will impact switching speed. The comparator design was optimized using four parallel devices to ensure the speed of the comparator meets the 100kbps data rate design spec at very low power levels. Common mode dynamic range is 300mV to 600mV and the input hysteresis ranges from +40mV to -40mV as seen in Figure 15. Power is designed to be stable across process variation using a 7-bit binary-weighted current DAC, similar to the one used in the rectifier. Thick oxide devices are used to reduce power consumption and device leakage.

Mismatch

Mismatch is significant in the subthreshold region and requires careful design to mitigate. The largest impact of mismatch in the receiver comes from the difference between the output of the active and bias rectifiers, V_{active} and V_{bias}. The first priority for mitigating mismatch in the receiver is through the programmable hysteresis in the comparator. The comparator's programmable hysteresis has a range of +/- 40mV with a resolution of around 2mV. Figure 15

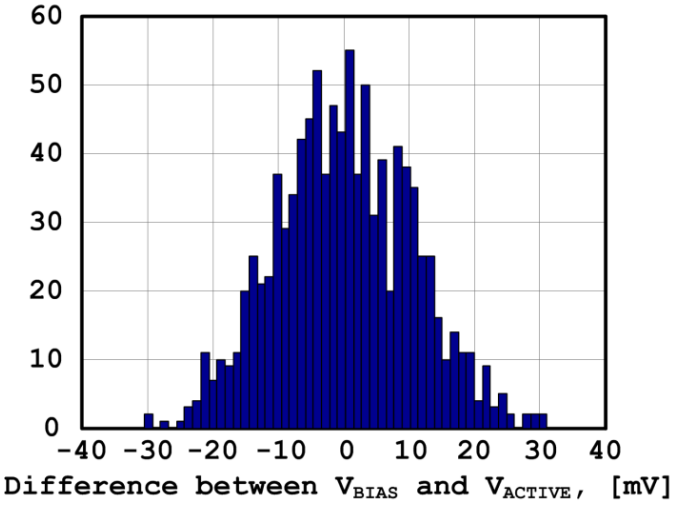

Figure 16: Monte Carlo Simulation when Individually Biasing Active and Replica Rectifier

shows a simulation of the different hysteresis levels achievable through switching the comparator's binary-weighted devices. Monte Carlo simulations shows this range will satisfy mismatch in over 50% of the chips. Further mismatch control is performed by individually varying the power in the active and bias current DACs. While this improves mismatch performance, it does so at the cost of power.

If mismatch is more significant than the range of the comparator and current DACs can handle, then both active and replica rectifiers can be self-biased independent of each other. To achieve this, a second mode of operation for the receiver is possible where two feedback amplifiers are used, individually sensing the drain voltages of M_1 and M_2 and comparing them against the same on-chip voltage reference. Monte Carlo simulations in Figure 16 show that the mismatch in this case stays within the +/-40mV range of the hysteretic comparator. The tradeoff with this approach is a reduction in sensitivity due to the feedback amplifier on the active rectifier side. Although the amplifier has a low bandwidth and is very slow to respond, it will sense the drop in the drain voltage due to the presence of a signal and will slowly alter the gate voltage to compensate.

Sleep Power

The receiver has been designed with a low-power sleep mode to support a duty-cycled wake-up strategy. This places extra emphasis on the importance of energy consumption in the receiver's sleep mode. To improve sleep mode energy, thick-oxide power gating devices were used throughout the design with above minimum lengths. Wake up time from sleep depends heavily on the startup time of the voltage regulator, which has been measured to be 110µs.

2.2.4. Results

The ultra-low power short range communication receiver, fabricated in IBM's 0.13µm CMOS technology, has an active area of 156x190 μm^2 and operates using a single 1.2V power supply. A die photo of the receiver can be seen in Figure 17. The receiver requires two independent IOs, the RF input and received data output. The other necessary IOs like power supplies and scan chain can be shared if the receiver is implemented into a larger system.

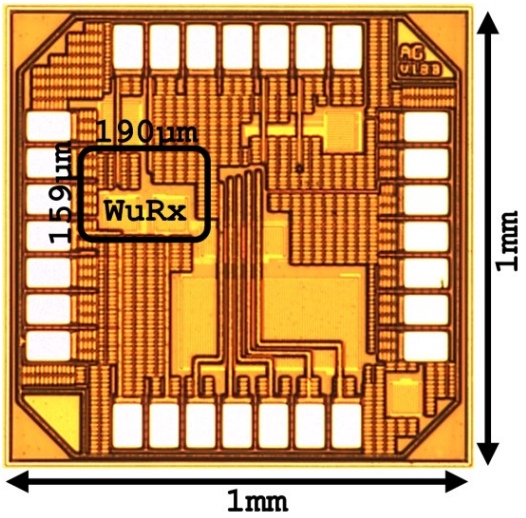

Figure 17: WRX Die Photo

A 915MHz OOK modulated signal with -41dBm sensitivity and a data rate of 100kbps was connected directly to the receiver input and the signal output was monitored on an oscilloscope. Figure 18 shows the transient operation of the receiver running at 98nW total power. Communication using patch antennas has been demonstrated at a distance of 4ft using a transmit power of 0dBm, which is roughly half the theoretical communication distance of 8.5ft based on the Friis equation.

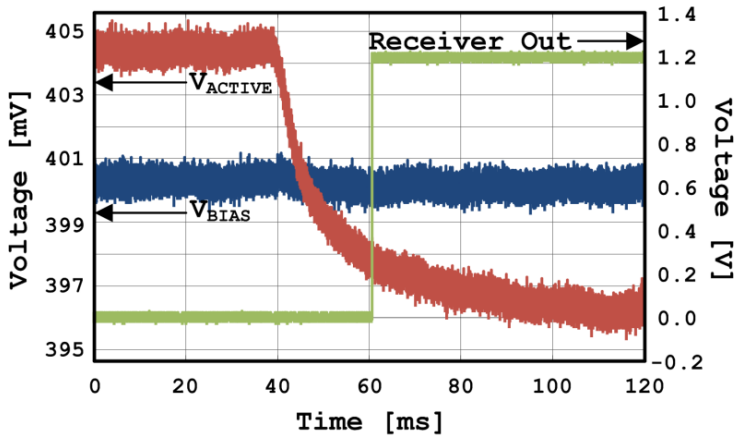

Figure 18: Transient Operation of WRX

Figure 19 shows the measured output voltage of the rectifier as a function of input amplitude. It can see that there is a slight knee in the line around the -40dBm mark, which represents the region of the receiver's peak sensitivity. Following the line to the left, a weakening signal, it can be seen that the rectifier is not able to produce a change in output voltage detectable by the comparator. On the other hand, with higher input, the conversion gain improves and allows for more potential power savings. For example, at -30dBm, the power consumption of the wake-up radio can be reduced to 53nW.

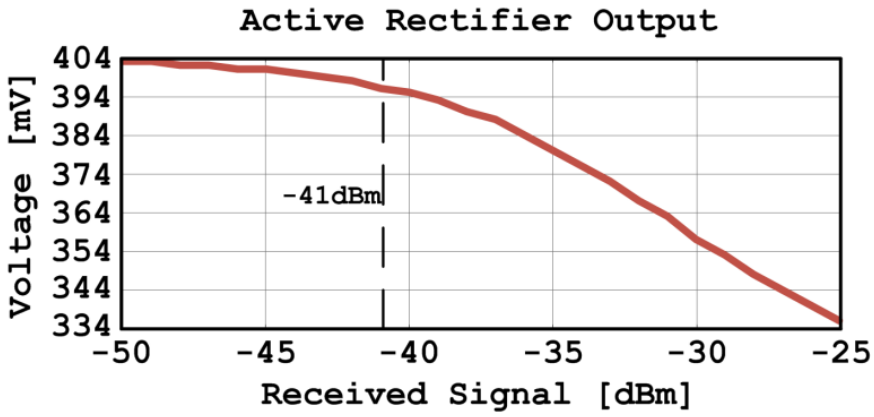

Figure 19: Active Rectifier Output Due to Input Signal Magnitude

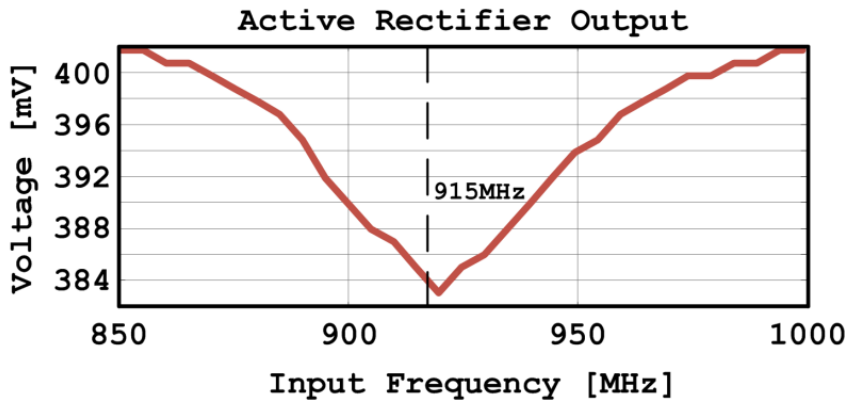

Figure 20: Impact of Resonant Tank

Figure 20 shows the impact of the simple off-chip resonant tank by measuring the change in output voltage on V_{active} as frequency is swept. Peak conversion gain is actually close to 920MHz. In addition the frequency sweep shows very rudimentary filtering which is important to help avoid interferers.

Energy detection architectures in general perform poorly in the presence of interferers due to their inability to distinguish the energy of a desired signal from an interferer, but one of the inherent advantages of this receiver's reduced sensitivity is that an interferer must either be very strong or very close to corrupt the desired signal. For example, a broadcasting 915MHz interferer transmitted with a power of -10dBm would be able to get up to 0.82 meters from the receiver before jamming it, calculated using the receiver's sensitivity and Friis free space path loss equation. At transmit powers of 0dBm, 5dBm, and 10dBm, the distances would increase to 2.6, 4.6, and 8.2 meters, respectively.

Table 4 shows a power breakdown of the receiver under the same conditions as Figure 18. The entire receiver consumed a total of 98nW with a measured sleep power of 11pW. The right side of the table shows the calibration setting resolution as well as the startup time from sleep for the on-chip voltage reference used in the feedback amplifier.

Circuit	Power [nW]		Min	Max
Bias Circuits	20	Hysteresis Calibration	-31mV	+29mV
Rectifier	23	Hysteresis Resolution	2mV	
Replica Bias	23	Current DAC Resolution	1.4nA	
Feedback Amp	4	On-chip Voltage Ref startup time	110µs	
Comparator	28			
TOTAL	**98**	Sleep Power	11pW	

Table 4: WRX Characterization Table

Table 5 shows a comparison with other low power receivers. The two key metrics to focus on between these receivers is the tradeoff between power and sensitivity. Looking at the two middle columns and comparing against this work's receiver in the rightmost column, one can see roughly a 30dB sensitivity difference between them. However, if one compares the same receivers in power, one will notice that the receiver in this work has a 27dB improvement.

	[50] JSSC'07	**[21] JSSC '08**	**[23] ISSC '10**	**This Work**
Power [µW]	500	52	51	0.098
Sensitivity [dBm]	-37	-72	-75	-41
Frequency [MHz]	916.5	2000	915	915
Data Rate [kbps]	1000	100	100	100
Energy/bit [pJ/bit]	500	520	510	0.98
Die Area [mm^2]	1.82	0.1	0.36	0.03
VDD [V]	1.4	0.5	0.5	1.2
Technology	0.18µm	90nm	90nm	0.13µm

Table 5: Comparison with State of the Art

2.2.5. Conclusion

Looking back at the hypothetical sensor node power breakdown (right side of Figure 5), the sensor node's low power wake up radio consumed almost 94% of the average power in the node. Replacing the hypothetical sensor node's WRX with this work's WRX reduces the receiver's power contribution from 94% to 3.3%.

By understanding the dynamics of a BAN node's channel model and making design tradeoffs between sensitivity and power, an ultra-low power receiver for short range communication designed in 0.13μm CMOS was presented. By avoiding the need to generate transconductance at RF, power of the receiver was reduced beyond the surveyed 50μW power floor. Using this methodology a 98nW receiver with -41dBm sensitivity and 100kbps data rate was demonstrated.

2.3. Adding CDMA codes for Multiple Access and Robustness

2.3.1. Introduction

Most published WRXs use simple OOK modulated energy detection architectures to keep power low. In such an architecture, the incoming signal may or may not be amplified at RF and then is directly converted to baseband using an envelope detector or energy rectifier. Bit-slicing can be done with a simple comparator. However, any ambient signal with enough strength at the proper frequency can trigger a false wake-up of these radios, and false wake-ups result in significant amounts of wasted energy on the node. In order to prevent this, a WRX must have enough local processing to differentiate a wake-up event from ambient interference, both in-band and out-of-band, without use of the node's main processor. In addition, if all the nodes in a sensor network rely on the WRX strategy using energy detection architectures, then when the transmitter tries to wake a node, it will either wake all the nodes at the same time or none at all, again producing significant wasted energy. To solve these issues, we present 116nW WRX

complete with crystal reference, interference compensation, and a subthreshold baseband processor capable of demodulating a selectable 31-bit OOK modulated code before toggling a wake-up signal. The front-end operates over a broad frequency range, tunable by an off-chip band-select filter and matching network, and demonstrated in the 403MHz MICS band, as well as the 915MHz and 2.4GHz ISM bands. Additionally, the baseband processor implements automatic threshold feedback to detect the presence of interferers and dynamically changes the receiver's sensitivity, mitigating the false wakeup problem inherent to previous energy-detection WRXs. The WRX radio has a raw OOK chip-rate of 12.5kbps. The radio is 0.35mm^2 and operates using a 1.2V supply for the crystal reference and RF demodulation, and a 0.5V supply for baseband processing in the subthreshold region.

2.3.2. System Architecture

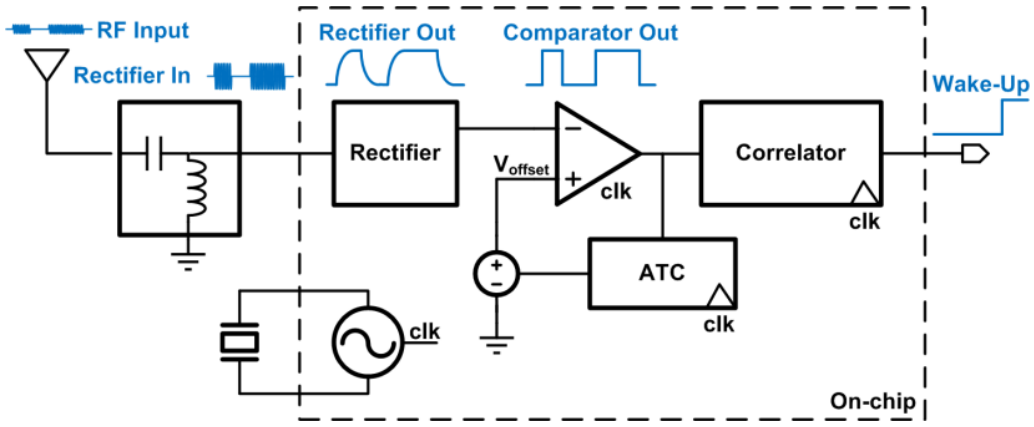

Figure 21: CDMA WRX System Architecture

Figure 21 shows a block diagram of the CDMA WRX. An OOK modulated RF signal passes through a passive input matching network that filters and boosts the signal before going on-chip. A 30-stage rectifier down-converts the RF signal to baseband, which is then sensed by the dynamic comparator clocked at 4X the chip-rate. The offset voltage of the comparator is controlled by the ATC (Automatic Threshold Controller) which compensates for interferers. A

bank of 4 parallel, bit-shifted 31-bit correlators continuously compare the received chip sequence with a programmable wake-up code, and toggles the wake-up signal only when a correlation result exceeds a programmable correlator threshold. The reference clock for the receiver is generated using an off-chip 50kHz crystal with an integrated oscillator. The oscillator and comparator operate from a 1.2V supply while all the digital logic operates in subthreshold at 0.5V. Because of the broadband design in the front end rectifier, the WRX can operate from the 400MHz to 2.4GHz bands, and can be tuned using a band-select filter and matching network off-chip. Sleep power was carefully designed using thick-oxide header devices on all the circuits.

My contribution to the CDMA WRX is the system design analysis as well as the implementation of the new low power crystal oscillator.

2.3.3. Circuit Description

In this section, the major circuit blocks of the WRX are explained in detail. All circuits use a thick-oxide PMOS header to improve sleep power.

Off-chip matching network

For the WRX, a 2 element off-chip matching network was used and provided a passive 5dB voltage boost. The input impedance of the chip was measured on a network analyzer to be 23-j35 Ω at 400MHz so a 12pF series capacitor and a 15.7nH shunt inductor were used. Devices like BAW or FBAR resonators can also be used to tune to the desired frequency of operation.

RF rectifier

Because the sensitivity has been reduced, an LNA is not necessary to amplify the received signal. Instead a zero-power RF rectifier replaces the LNA, saving significant power and allowing communication in the nanowatt range. As seen in Figure 22, the rectifier's structure is

the same as the Dickson Multiplier [63], but the output voltage calculation is different because all transistors operate in subthreshold due to the small RF input [64].

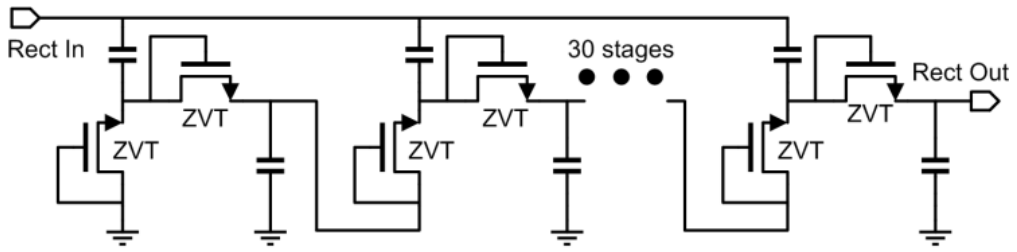

Figure 22: RF Rectifier Used as the Front End of the CDMA WRX

This subthreshold rectifier uses zero-threshold transistors and 30 stages to achieve sufficient RF gain with a fast charging time. The input impedance of the chip is 23-j35 Ω at 400MHz, 12-j13 Ω at 900MHz and 88-j5.8 Ω at 2.4GHz. The Q factor of the input impedance is low, due to a voltage limiter that prevents the rectified voltage from exceeding the breakdown voltage of the FETs, so a broadband matching network could be implemented.

Comparator with ATC

The clocked comparator, shown in Figure 23, applies regenerative feedback clocked by the 50kHz oscillator. Two separate current biases are each controlled by 4-bit binary-weighted

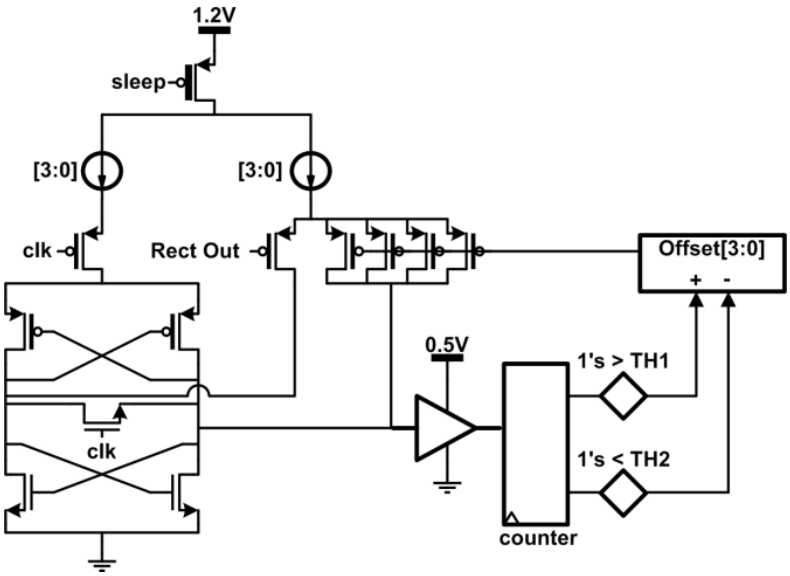

Figure 23: Dynamic Comparator with ATC

current DACs. In addition, the comparator threshold can be programmed to a 4-bit binary-weighted value to tune the sensitivity of the receiver. The threshold of the comparator is controlled in a feedback loop by the ATC which dynamically changes the comparator's offset voltage to overcome interference signals. A diagram showing operation of the ATC can be seen in Figure 24. As RF input signal comes in, the RF rectifier outputs the signal for the comparator. The comparator compares this signal with its threshold. The ATC monitors the samples coming from the comparator output for one 31-bit code period. If the number of 1's is greater than a user defined $value_1$ (indicating the comparator threshold is continuously exceeded by an interferer), then the ATC will increase the comparator threshold to bring the sensitivity of the receiver above that of the interfering signal. When the number of 0's at the output of the comparator reaches a separate, user defined $value_2$ (indicating the interferer is gone and the comparator threshold is never exceeded), the ATC then reduces the threshold to increase the sensitivity of the receiver. Hysteresis is added between these values to eliminate limit-cycles. With this mechanism, the comparator can reject interference signals, and even if the interference signal is modulated by

BPSK, or OOK, the comparator threshold is set to above the maximum level of interference so the comparator will produce the correct output.

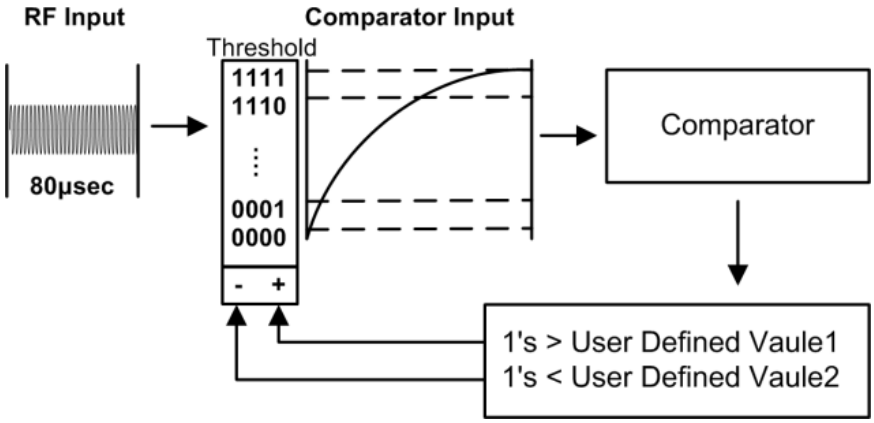

Figure 24: Signal Diagram of ATC

Correlators

A bank of four correlators continuously correlates the 4x-oversampled comparator bit-stream with a programmable, 31-bit CDMA code. This synchronizes to the transmitted code and only issues a wake-up output when the desired code is received. The digital baseband processor was synthesized in subthreshold in order to save power. Eight different Gold codes can be selected using control bits in a scan chain. This will allow for a single transmitter to uniquely wake up 8 possible WRXs. Gold codes are a set of binary sequences whose cross-correlation among the set is bounded into three values [65]. Gold codes are commonly used when implementing CDMA and they are easily implemented with 2 LFSRs (Linear Feedback Shift Register) and an XOR

gate. In this work, 31-bit Gold codes with 3 configuration bits are implemented.

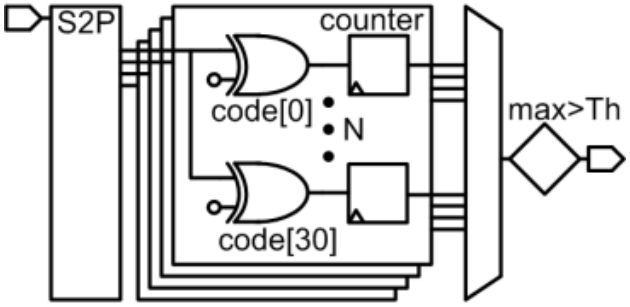

Figure 25: Block Diagram of Correlators

The correlator compares the last two samples in each bit slice. Therefore, each 31-bit code results in a total of 62 comparisons. A programmable correlator threshold allows the user to define a value between 1 and 61 that must be exceeded in order to declare a code received indicating a valid wake-up event. A lower threshold value would mean fewer bits have to match the code, tolerating a higher BER and resulting in better sensitivity, but leads to more false wake-ups. A higher threshold would prevent false wake-ups, but also reduce the sensitivity of the receiver.

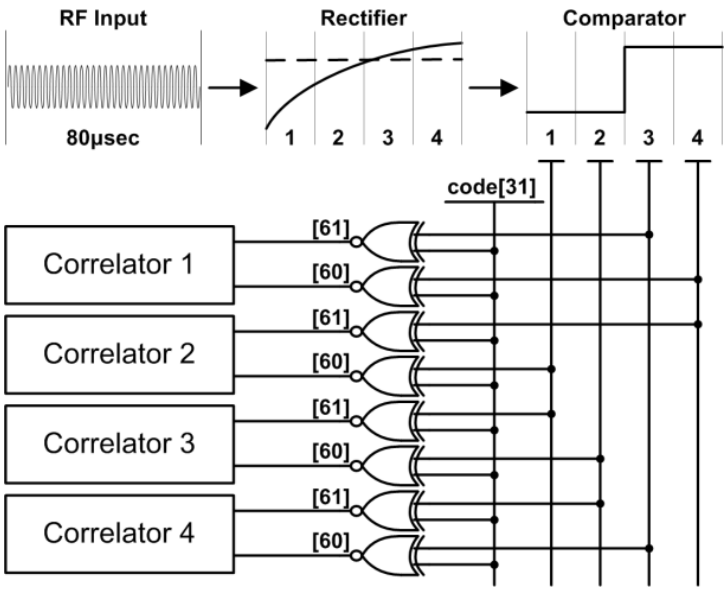

Figure 26: Correlator Signaling

Figure 25 shows a basic block diagram of the correlator. To synchronize the receiver to the transmitted code, 4 correlators operate at the same time and each correlator receives shifted samples of each bit slice since the receiver is 4x-oversampled. In each correlator, all possbile shifts of the 31bit Gold code are simultaneously correlated with the incoming bit stream, so that after a single 31bit sequence, the receiver is guaranteed to synchronize to the wake-up signal. Each parallel correlator will have a different number of correct comparisons based on the code shifts and phase difference between the WRX and the transmitter. If any of the 4 correlators results are greater than the correlator threshold, the wake-up signal will be asserted. An example showing how the correlator handles a single chip slice is shown in Figure 26. This example shows what happens if the WRX receives a '1' as the last chip in the sequence. The OOK signal is gradually rectified during the chip window.

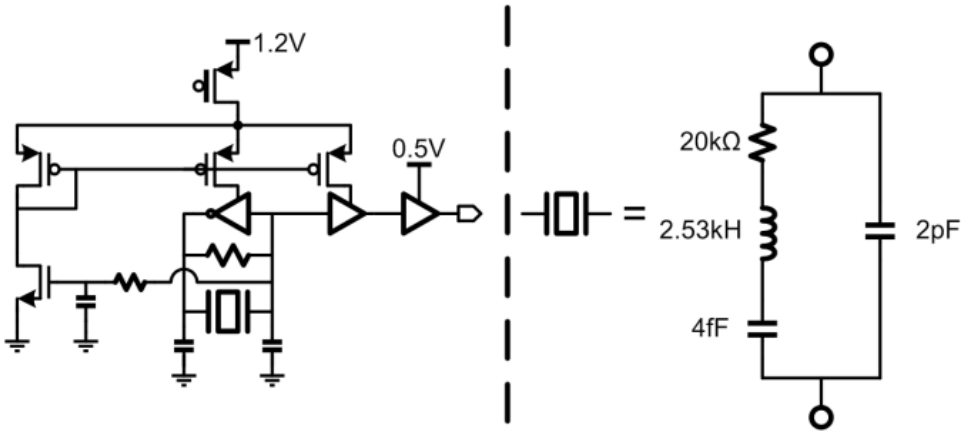

Figure 27: Block Diagram of 50kHz Crystal Oscillator

Crystal Oscillator

A 50kHz crystal oscillator serves as the reference clock of the radio [66]. As seen in Figure 27, an off-chip crystal is used, and the oscillator's primary amplifier is an inverter with resistive feedback. Figure 27 also shows the intrinsic values of the 50kHz crystal. Using these values the critical g_m can be calculated using Equation 1, which is the transconductance of the amplifier that must be produced to achieve sustained oscillations. C and C_0 represent the motional and shunt capacitance of the crystal, ω and Q are the resonant frequency and Q factor of the crystal, and C1 and C2 are the load capacitance in the circuit. If the primary amplifier is biased in the near threshold region where the g_m/i_D ratio is around 10, then the current consumption to reach this critical gm value is around 20nA.

$$\boldsymbol{Critical\ g_m = \frac{\omega}{QC} * \frac{(C_1C_2 + C_2C_0 + C_0C_1)^2}{C_1C_2}} \quad [66] \qquad (1)$$

Initially, the transconductance of the primary amplifier is much greater than the critical g_m of the crystal, which is needed to quickly increase the oscillation amplitude. However, as the oscillation amplitude increases, the DC level of the oscillation also drops and this common-mode signal is used in feedback to starve the primary amplifier until it settles with sustained oscillations. Measured results show the total power consumption is 30nW when sustaining

oscillations using a 1.2V supply. The oscillation is then buffered to provide the reference clock for the WRX. Measured results show the oscillator has an RMS jitter of 6ns.

Figure 28 shows the transient of the startup time for the oscillator. The top of the figure shows the current consumption during startup and shows that the current consumption peaks around 95nA before dropping as oscillations begin to increase. The current consumption finally settles to 25nA after a second has passed. While this start-up time is relatively long, the WRX can be left on continuously in a typical sensor node due to its ultra-low power consumption.

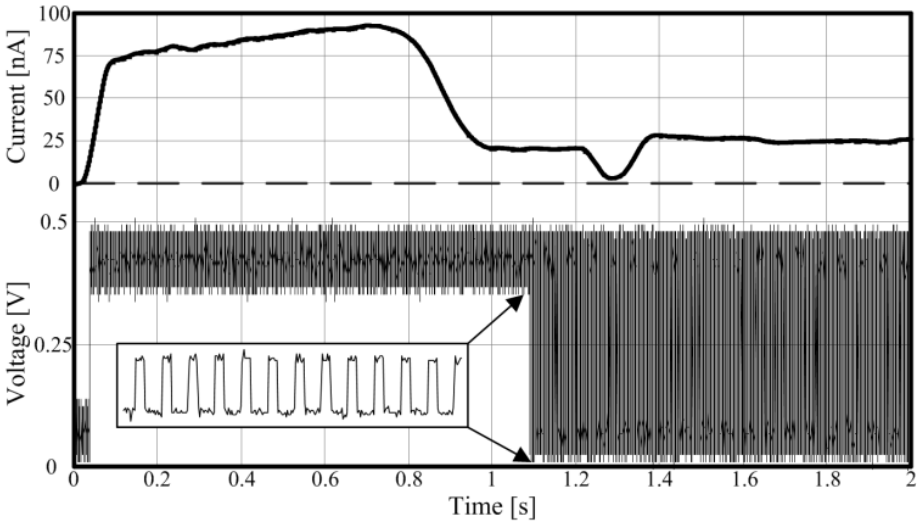

Figure 28: Transient Startup Time of Crystal Oscillator

Sleep Power

Sleep power in the WRX was carefully considered during the design process to support a duty cycled wakeup strategy. To improve sleep mode energy, thick oxide power gating devices were used throughout the design. Wake-up time is dominated by the slow startup time of the crystal oscillator, which takes about 1.1 seconds to oscillate. Putting the WRX into a full sleep mode, where every circuit is power-gated, results in a sleep power of 18pW. Clock gating the crystal oscillator, and putting all the circuits except the crystal oscillator into sleep mode, results in a sleep power of 30nW.

2.3.4. Measurement Results

The WRX was fabricated in IBM's 130nm CMOS process and has an active area of 0.35mm^2 without pads. It uses 2 separate voltage supplies; a 1.2V supply for the crystal oscillator and demodulating comparator and a 0.5V supply for all baseband processing. A die micrograph can be seen in Figure 29.

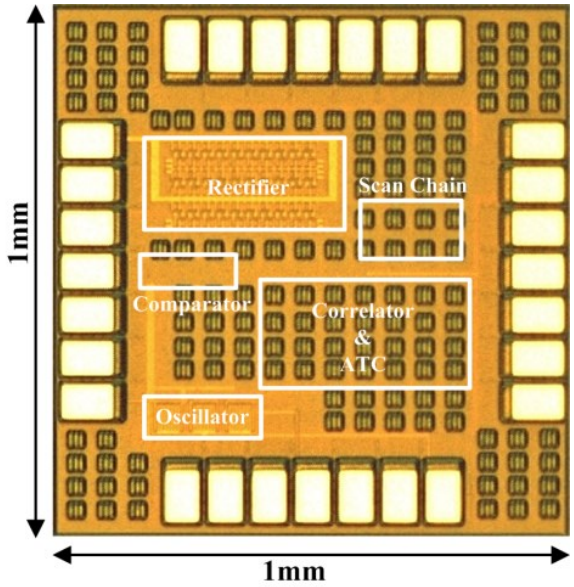

Figure 29: WRX Die Photo

Transient operation of the WRX receiving a 31-bit code is shown in Figure 30. Because of the bit-shifted parallel correlators, the WRX is able to automatically synchronize to the incoming bit stream upon receiving the first wake-up code. The top two traces show the RF input signal and the RF rectifier converting the signal to baseband. The third trace shows the output of the comparator being clocked at 4X the data-rate by the oscillator and the final trace is the wake-up signal being toggled by the correlator. The WRX is capable of CDMA by selecting different codes used by the correlator block. If an interfering signal is strong enough to exceed the

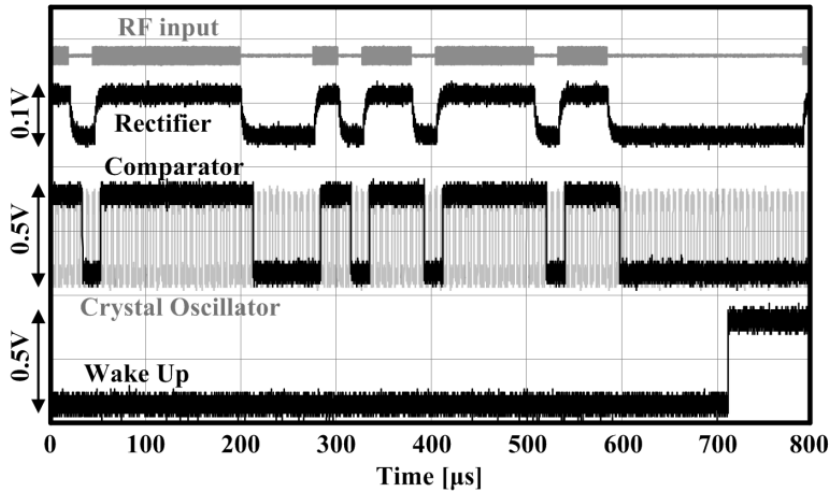

Figure 30: Transient Response of WRX

comparator threshold, then the ATC increases the comparator's threshold until it is above the interfering signal. A transient of this operation can be seen in Figure 31. The top signal is the received RF signal, which is jammed by a 2.4GHz tone at 8ms. With the interferer present, the receiver cannot initially demodulate the code. After 15ms, the ATC has raised the threshold of the comparator above that of the interferer, and the WRX regains synchronization.

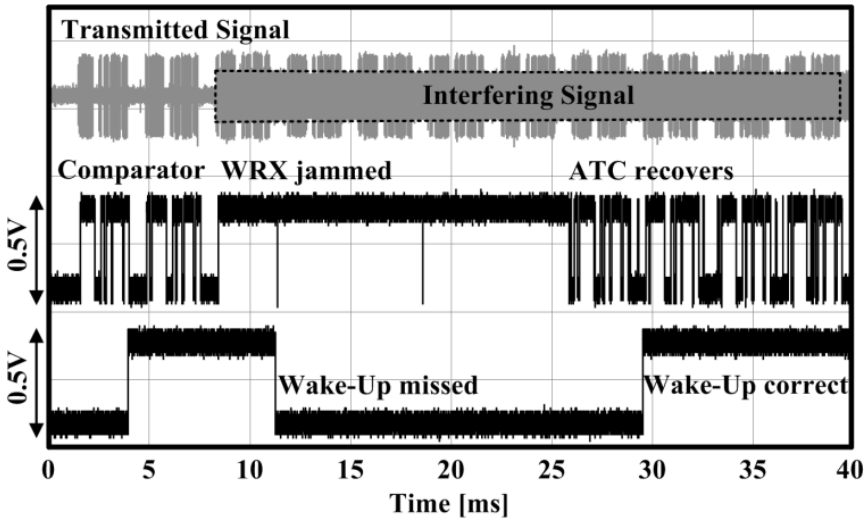

Figure 31: Transient Response in the Presence of an Interferer

The WRX has an active power of 116nW with a sleep power of 18pW. It has a raw OOK chip-rate of 12.5kbps and a symbol rate of 50kbps and sensitivities of -45.5dBm, -43.4dBm, and

-43.2dBm at 403MHz, 915MHz, and 2.4GHz, respectively. The top of Figure 32 shows the chip error rate (BER; chip = 1 bit of data) curves for the 403MHz, 915MHz, and 2.4GHz bands. The bottom of Figure 32 shows the BER as the correlator threshold is varied. The measurements were taken using a -40dBm received signal in the 2.4GHz band. The figure also shows the impact this threshold has on false wake-ups. From these two data sets, the correlator threshold can be set to maximize sensitivity while minimizing the possibility of a false wake-up.

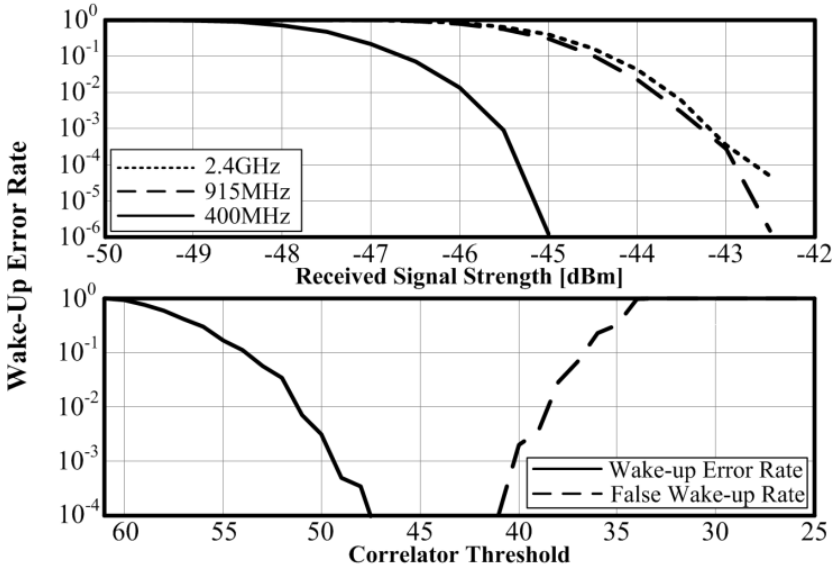

Figure 32: BER vs. Received Signal Strength and Correlator Threshold

To demonstrate the ability of the WRX to be selective in its wakeup code, therefore allowing different codes to wake-up different WRXs, Figure 33 shows an experimental setup where the input signal was connected to two different WRXs, each with a different wakeup code programmed. The top of Figure 33 shows that code 1 and code 2 were transmitted back to back and the resulting output from the WRXs correctly toggles when the proper code is sent, but ignores the incorrect code.

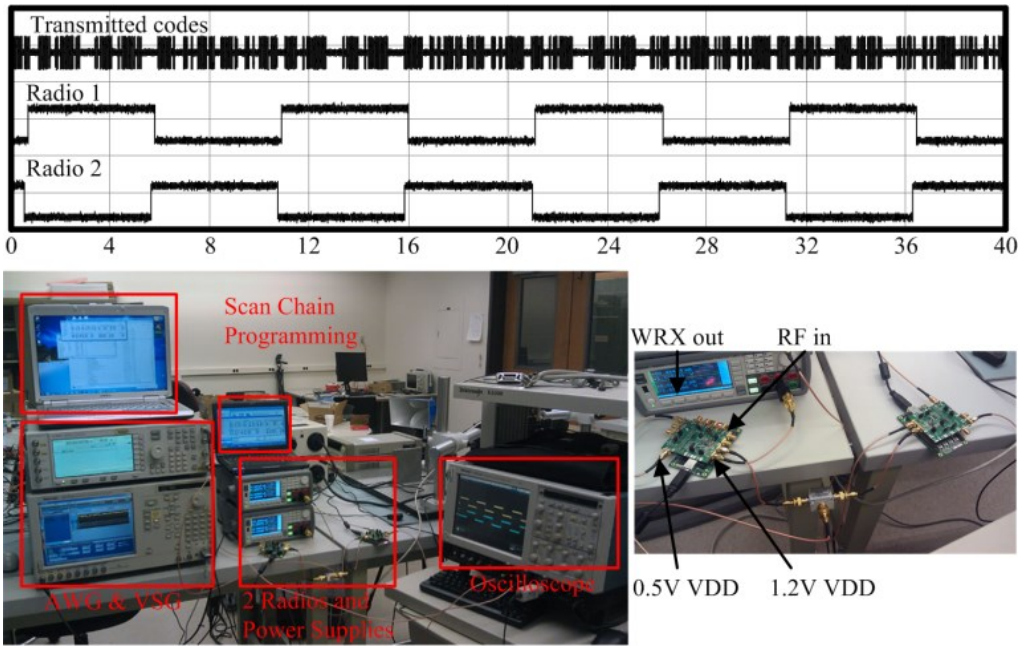

Figure 33: 2 WRX Measurement Setup and Result

Table 6 shows a power breakdown on the WRX on the left. Digital logic consisting of the 4 parallel correlators and a clock divider all synthesized in subthreshold consumes the bulk of the 116nW power. Second is the crystal oscillator consuming 38.4nW. The demodulating comparator consumes 8.4nW and the RF rectifier front end consumes no power, for a total of 116.3nW.

The right side of Table 6 shows the receiver's specs. It requires 9.28pJ/bit and with a 31-bit wake-up code, the total energy to wake-up the node would be 287.7pJ. The maximum signal level tolerable is -15dBm and the maximum interferer level is -20dBm.

Power Breakdown [nW]		Receiver Specs	
RF Rectifier	0	Energy/bit	9.28pJ
Comparator	8.4	Energy/wake-up	287.7pJ
Digital Logic	69.5	Max Signal Level	-15dBm
Crystal Oscillator	38.4	Max Interferer Level	-20dBm
TOTAL	**116.3**	Code Length	31
Sleep [pW]	20	# Pre-defined Codes	8

Table 6: Power Breakdown and Receiver Specs

Table 7 shows a comparison versus the state of the art WRXs previously mentioned. It is apparent the design tradeoffs that are made to both reduce power and interference. In [21] the sensitivity is nearly 30dB better than this work, but the power is nearly 30dB greater. [25] blocks out-of-band interferers, but not in-band interferers and also has higher power. [67] uses the same sensitivity reduction technique to lower power, but does not have any way to compensate interferers.

	This Work			[21]	[25]	[67]
Power [μW]	0.116			52	45	0.098
Sleep [pW]	20			N/A	N/A	11
Frequency [MHz]	403	915	2400	2000	5800	915
Data rate [kbps]	12.5/31			100	14	100
Sensitivity [dBm]	-45	-43	-41	-72	-45	-41
SIR [dB]	3.3	1.7	1.7	N/A	N/A	N/A
Die Area [mm^2]	0.35			0.1	N/A	0.03
VDD [V]	1.2 | 0.5			0.5	3.0-3.6	1.2
Process [nm]	130			90	130	130

Table 7: CDMA WRX vs State of the Art

2.3.5. Conclusion

This section introduced a 116nW WRX that uses automatic threshold compensation to provide interference rejection from both in-band and out-of-band interferers. In addition, with 8 selectable wake-up codes, the WRX can be used in a network with similar WRXs and be able to uniquely wake-up. With reduced sensitivity specifications, the use of a zero-power RF energy harvester was used as the RF front end of the receiver and subthreshold design was implemented to keep entire radio in the nanowatt power region. With power that is less than a typical sensor node's sleep power, the WRX is not the energy dominant circuit when the node is asleep and can provide false wake-up rejection, making it a very suitable synchronization technique for sensor nodes.

A comparison between the WRX and BLE's scanning methodology for synchronizing sensor node radios shows how advantageous the WRX can be. The shortest duration a BLE receiver can be in the scanning state (active) is 2.5ms and the BLE receiver discussed in the beginning of this chapter, with sensitivity shown in Figure 9, consumes 47.4mW of power when active. That equals 118.5μJ of energy for a minimum scan time of 2.5ms. The WRX, for the same amount of energy, can actively listen for 17 minutes!

2.4. Acknowledgements

I would like to thank Seunghyun Oh for his contributions with the CDMA WRX. Seunghyun helped with the system architecture and was responsible for the majority of the IC design, using his previously designed rectifier for the front-end and designing the digital baseband correlators.

CHAPTER 3

Wireless Communication in a Battery-less SoC

Currently, one of the most appealing applications of IoT is in the medical healthcare space. The influx of wearable health monitors being released demonstrates the appeal for both healthcare professionals and consumers alike. Current health monitors such as the Nike Fuelband, Jawbone UP, or FitBit Flex focus on activity and movement which is measured through the use of an accelerometer. Newer health monitors are beginning to incorporate heart rate monitoring in addition to movement. Like most of the IoT devices, energy consumption is a primary concern. As discussed in the intro section, the Nike Fuelband consumes 23mA of current by simply running the band's LEDs while the Lifewatch ACTIII heart rate monitor has a measured average power consumption of 63mW. Power levels that high require the use of a battery and still have lifetimes that only span a few days.

Due to the low data rate required to sense movement and heart rate, ultra-low power processing in the subthreshold domain is possible. The main power bottlenecks therefore lie in the analog front end and the wireless communication. Wireless communication, as has been discussed, can be so significant that often early wearables (Jawbone UP) did not include it as an option. Looking back at the Nike Fuelband teardown in the Introduction section, its wireless communication chip consumed 123.8mW of active power according to its datasheet.

To further enable IoT and demonstrate the ability of hardware capable of running without batteries while maintaining the ability to communicate wirelessly, a 4-lead EKG monitoring WSN (EKG WSN) that runs off energy harvested from a thermoelectric generator (TEG) was designed in collaboration with the University of Virginia (UVA). Targeted to be a body worn device, size is a critical design element which limits the amount of potential energy harvesting and storage, effectively eliminating the possibility of using commercial short range, low power wireless standards due to their significant energy overhead. However, consideration needs to be given to the effectiveness of any custom radio as the EKG WSN will be deployed in a realistic multi-user environment. Many features that are present in commercial standards, but overlooked in academia due to their increase in power, need to be considered such as multiple access and robustness to interference.

This chapter will discuss the design of a transceiver for wireless communication in a small form factor sensor node that runs off harvested energy. The next section will introduce a general overview of the EKG WSN and the transceiver will be introduced and detailed in the sections after.

3.1. A EKG Sensor Node SoC Powered by Thermoelectric Energy

The EKG WSN is used for biological signal monitoring and has hardware accelerators that aid it in heart rate extraction and atrial fibrillation detection. It runs off harvested energy from a TEG and a power management unit boosts the low TEG voltage to a supply voltage of 0.5V for subthreshold processing and memory and 1.2V for analog functions like the analog front end (AFE) and wireless communication.

The EKG WSN uses two different processing units, the primary unit being an *openMSP430* core which contains a 16-bit RISC architecture and a UART debugging interface. The secondary

processor is called the Lightweight Controller Unit (LCU) which is also a 16-bit RISC architecture and controls data flow when the larger *openMSP430* is off. Both processors use a programmable clock capable of providing clock rates from 187.5kHz to 500kHz. Programming for the EKG WSN is controlled through a synchronous SPI interface.

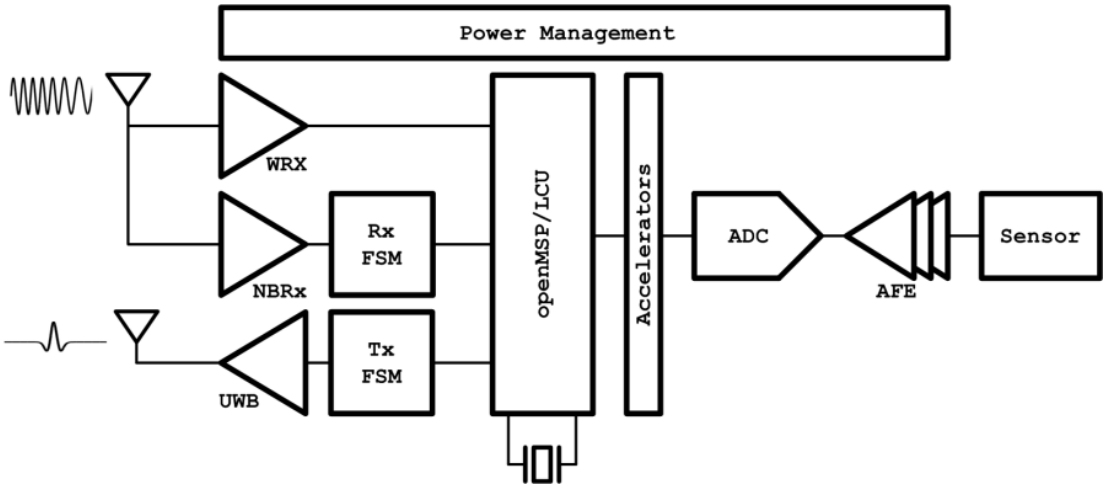

Figure 34: EKG WSN Simplified Block Diagram

Peripherals include a basic timer, a hardware multiplier supporting 16-bit operations, direct memory access (DMA), and 2 hardware accelerators: a FIR and RR-Afib detection. With these accelerators the node is capable of taking a 4-lead EKG and detecting abnormalities in the heart like an arrhythmia.

The EKG WSN consumes an average of 19μA of current which allows it to be continuously powered by the TEG. Figure 34 shows a block diagram of the EKG SoC and Figure 35 shows the chip layout.

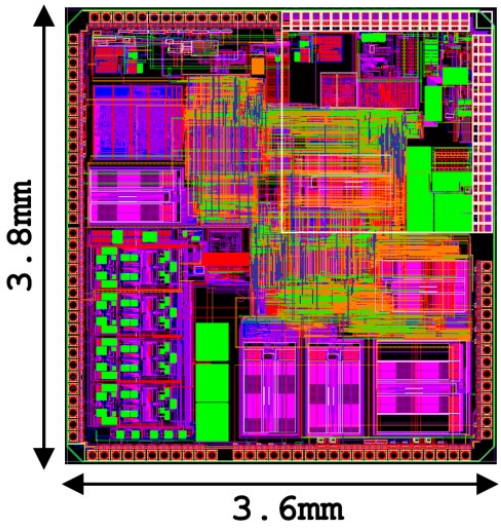

Figure 35: Die Layout of EKG WSN

3.2. The WBAN Wireless Transceiver

The WBAN transceiver was designed to be compatible with the recently adopted 802.15.6 WBAN standard. The purpose of the WBAN standard is to provide "an international standard for a short range, low power, and highly reliable wireless communication for use in close proximity to, or inside, a human body." [17] Current personal area networks (i.e. 802.15.4) do not meet the medical regulations for some application environments, nor do they support the combination of reliability, quality of service, low power, data rate, and non-interference required to support the scope of BAN applications [17], which is what motivated 802.15.6.

Introduction

The design of a WBAN transceiver to be used to enable wireless communication in an EKG sensor node that harvests thermoelectric energy requires a mix of standard-compatible and proprietary systems. Complying with the entire WBAN standard, outlined in the Introduction chapter, requires too much power to enable wireless communication using harvested energy without heavy duty cycling. However, it should be noted that battery-less operation is not in the

scope of the 802.15.6 standard. The designed transceiver contains a proprietary CDMA WRX, a proprietary UWB transmitter that operates in an 802.15.6 compliant frequency band, and an 802.15.6 compatible narrowband receiver.

System Architecture

The SoC transceiver uses an asymmetric communication link where the WBAN receive frequency is in the 402-405MHz MICS band, and the UWB transmit frequency is 4GHz. The lower frequency in the receive band allows for lower power design, but does so at the cost of a larger antenna, which is an issue for body worn sensor nodes.

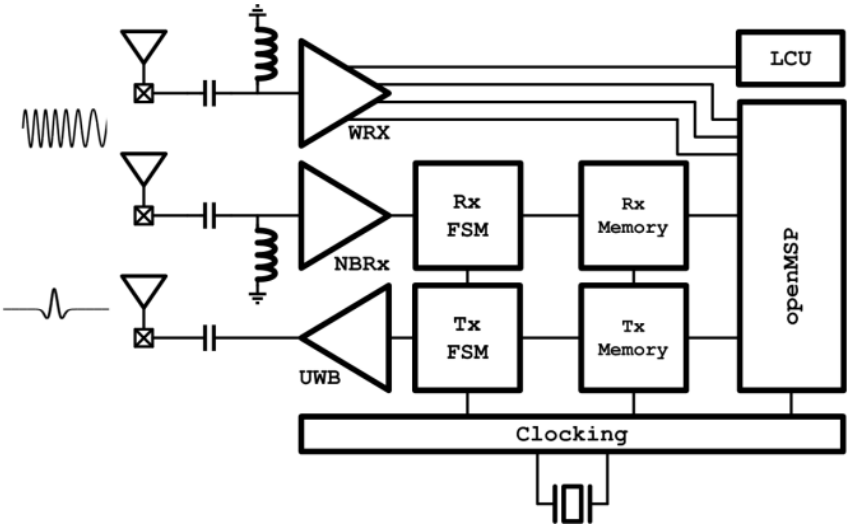

Figure 36: WBAN Transceiver

A block diagram of the entire transceiver can be found in Figure 36. It consists of a CDMA WRX, a narrowband receiver, and a UWB transmitter. In addition, receive and transmit FSMs and memory buffers were designed to provide a fully functional interface to the sensor node. The CDMA WRX has four separate interrupt outputs, one hardcoded to each the openMSP and LCU, with two others being programmable.

I was responsible for the entire system design and architecture of the transceiver and making sure it interfaced with the UVA processor and 802.15.6 WBAN specs. I also designed the

narrowband communication receiver as well as the low power crystal oscillator that was used as the clock reference for the entire SoC.

3.2.1. Wake Up Receiver

The WRX on in the EKG WSN is the same CDMA WRX from the previous chapter, except the baseband processor now correlates four parallel 15-bit codes (Kasami) instead of a single 31-bit code. This subtle, but significant change makes the CDMA WRX much more appealing in the SoC. There are two different ways the codes can be utilized.

The first method is to use each wake-up code as an individual interrupt so the user can use different codes to perform different actions. For example, $code_0$ could wake up the node in its default operating condition. $Code_1$ might tell the SoC to transmit its data buffer to an aggregator. The second method the CDMA WRX can be used for is to become an ultra-low power communication receiver. Since the CDMA WRX is looking for four different codes, each code can represent two data bits. This allows for very low data-rate applications to run strictly through the CDMA WRX. The updated block diagram for the new CDMA WRX can be seen in Figure 37.

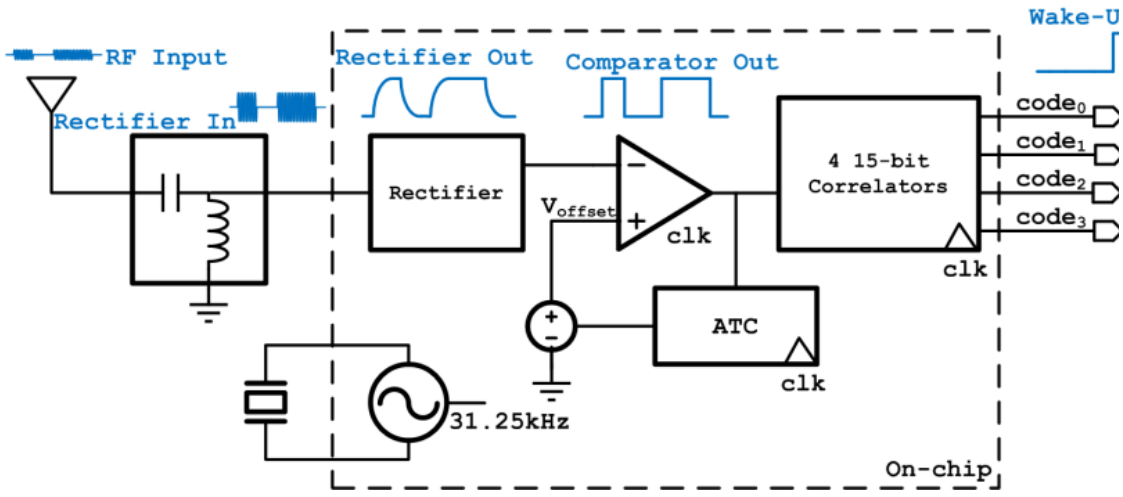

Figure 37: CDMA WRX for WBAN

4 Parallel 15-bit Codes

The disadvantage of a smaller code to correlate is an increase in false wakeups, but the advantage the shorter code provides is that the WRX can now look for four codes at the same time without a significant increase in power. The previous CDMA WRX made four correlations for every bit which resulted in 4*31 = 124 correlations per code. Now the CDMA WRX does 4*4*15 = 240 correlations per code.

Measurement Results

Full measurement results are not included since the CDMA WRX is so similar to the previous CDMA WRX. Table 8 summarizes the performance and power characteristics of the new design.

Power Breakdown [nW]		Receiver Specs	
RF Rectifier	0	Energy/bit	10.75pJ
Comparator	12	Energy/wake-up	161.25pJ
Digital Logic	70	Max Signal Level	-15dBm*
Crystal Oscillator	30	Max Interferer Level	-20dBm*
TOTAL	**112**	Code Length	15
Sleep [pW]	20*	# Pre-defined Codes	Programmable

Table 8: CDMA WRX Measurement Results
***measurements taken from previous CDMA WRX chip…not a new measurement.**

Practical Demonstration

To demonstrate the CDMA WRX's merits, a practical demonstration was developed and presented at the Consumer Electronics Show in January, 2014. The CDMA WRX breakout chip was built into a discrete prototype board that ran off solar energy. Once the solar cells charged up to 3V a custom discrete power management unit (PMU) turned on the WRX and a PIC16LF1459 microcontroller. The microcontroller immediately programmed the CDMA WRX's scan chain and then went to sleep. When a code was received the WRX would interrupt the PIC and it would toggle a 7-segment LCD screen for 2 seconds before going back to sleep.

A push button interrupt would enable multiple different interrupt functions based on the number of button presses within a 2 second window. When the harvested energy fell below 1.3V the PMU would turn off.

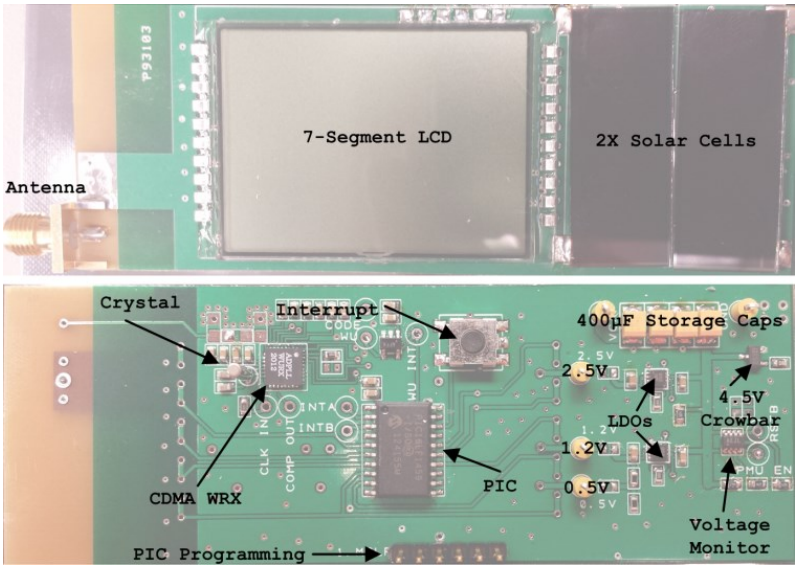

Figure 38: CDMA WRX Demo

Figure 38 shows the CDMA WRX demo highlighting the major components in the design. The total power consumption was 6µW including the CDMA WRX, PIC, PMU, and supporting circuitry. 5µW comes from the PMU alone, making the wireless communication piece a fraction of the overall power.

3.2.2. UWB Transmitter

The UWB transmitter is an OOK modulated transmitter with a data rate of 187.5kbps, chosen to match a data rate option in the 802.15.6 WBAN standard. The center frequency is 4GHz with a 500MHz bandwidth. It consists of a pulse-width generation circuit which uses two separate signal paths with different variable delays to create a short pulse which is then used to enable/disable a ring oscillator, creating the UWB pulses. A Class AB power amplifier then buffers the signal onto the antenna. The UWB transmitter block diagram can be seen in Figure

39. I was responsible for the system design decisions of the UWB transmitter as well as the design of the transmitter's FSM.

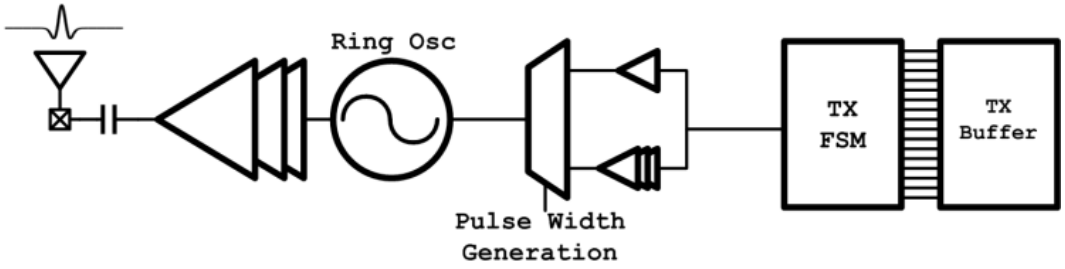

Figure 39: UWB Transmitter

Tx FSM

The Tx FSM has a 2kB memory array that stores the raw data to be transmitted and is written by the processor. Raw data includes any synchronization and preamble headers the receiving base station receiver might need. The Tx FSM is responsible for clocking the data out of the memory array and serializing it before sending it to the UWB transmitter. The Tx FSM also has the option to append a timestamp to the serial data before it is sent to the transmitter.

UWB Ring Oscillator

The 4GHz ring oscillator consists of two parallel 3-stage CMOS-based ring oscillators with the first stage consisting of a NAND gate to allow for the pulse signal to enable and disable the ring's oscillation, as seen in Figure 40. Binary weighted current sources on both NMOS and PMOS control the speed and power of the oscillations.

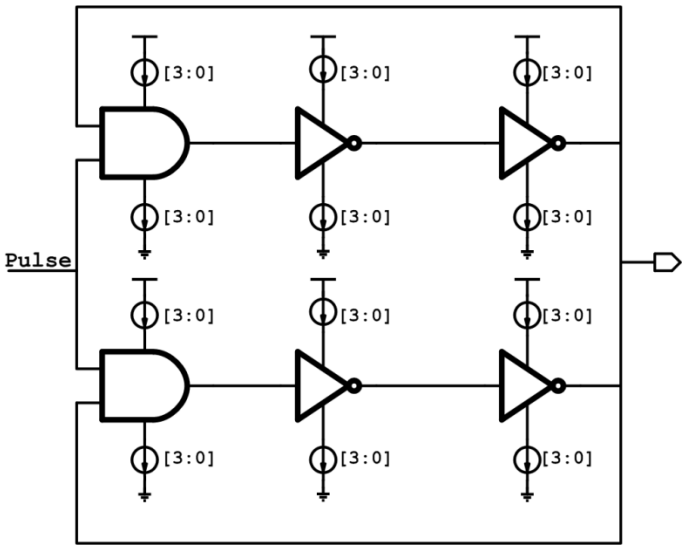

Figure 40: UWB Tx Ring Oscillator

Measurement Results

A breakout chip of the UWB transmitter was tested and Table 9 shows the primary performance metrics. The center frequency was measured at 3.996GHz with a 492MHz bandwidth. Power consumption at the 187.5kbps data rate is 4.18μW. Output power has been measured to be -28.9dBm. Figure 41 shows a transient response of a single pulse.

Spec	Value	Unit
Power	4.18	μW
Data Rate	187.5	kbps
Center Frequency	3.996	GHz
Bandwidth	492	MHz
Output power	-28.9	dBm

Table 9: UWB Tx Measurement Results

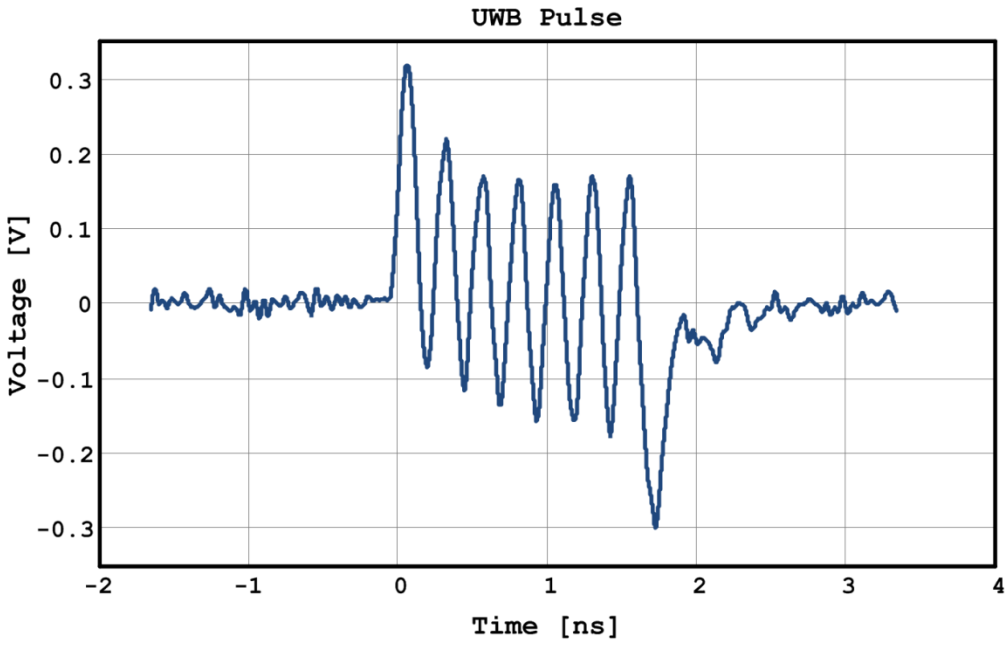

Figure 41: UWB Pulse

3.2.3. WBAN-compatible Narrowband Receiver (NBRx)

None of the ultra-low power radios presented in the survey (Figure 7) in the Introduction are standard compliant because of the energy overhead. Standard compliant radios like Bluetooth and Zigbee consume much higher power and therefore are not shown in the survey. Two factors that standard compliant receivers address that often are not addressed in ultra-low power radios is frequency selectivity and multiple access. Without these, two or more similar transmitters and receivers would jam one another if all were trying to communicate simultaneously. Furthermore, interference from other incumbent radios in the same band can saturate the energy detectors in these radios. In order to be effective in real-world environments low power radios need to have some of the capabilities that more sophisticated and higher power receivers have like the ability to select multiple channels within a band, which allows for multiple users. And conforming to a standard would allow for easier adoption for manufacturers.

Because of the limitations in current ultra-low power receivers, the narrowband communication receiver (NBRx) on the EKG WSN is a low power, coherent, heterodyne receiver that is compatible with the 802.15.6 WBAN standard in the MICS band and capable of channel selection at RF. In order to be compatible with WBAN the receiver can select between 10 300kHz channels between 402-405MHz. In addition the RF front end coherently down-converts $\pi/2$ DBPSK signals at a symbol rate of 187.5kbps, as governed by the standard [17].

System Architecture

Figure 42 shows a block diagram of the receiver. The antenna receives the $\pi/2$ DBPSK signal and passes through a COTS MICS band SAW filter and an external matching network. Once on-chip, the signal is amplified by the LNA and down converted to a 10.7MHz IF frequency by a passive switching mixer and PLL-based LO, which is used for channel selection. After down conversion the signal passes through an external 10.7MHz ceramic channel select filter with a 300kHz bandwidth. After filtering the signal goes back on-chip and is squared up and serves as the reference frequency for the demodulator. After locking, the IF PLL charge pump and LF will modulate up or down based on the direction of the phase shift of the received $\pi/2$ DBPSK signal. An identical second charge pump receives the same PFD output as the IF PLL charge pump, but is used to charge and discharge a storage capacitor. At the beginning of a bit slice interval the storage capacitor is reset to VDD/2 and then charges or discharges throughout the bit slice interval based on the direction of the phase change. At the end of the bit interval the storage capacitor charge is compared against VDD/2. A 1 is produced if the capacitor value is greater than VDD/2 and a 0 is produced if the capacitor value is less than VDD/2. Bit slicing and synchronization is accomplished with a demodulation PLL designed in subthreshold and the Rx FSM.

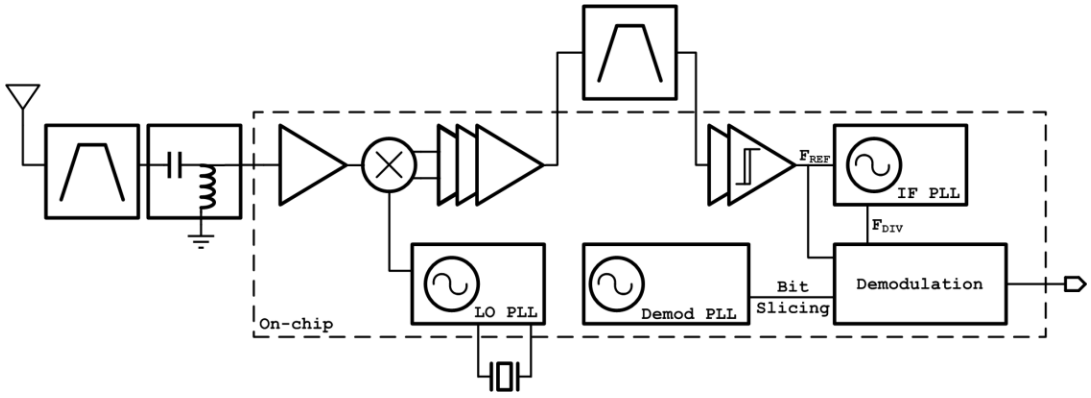

Figure 42: System Block Diagram

Off-chip Matching Network

For the receiver, a 2-element off-chip matching network was used to provide matching for the low noise amplifier. The input impedance of the chip was measured on a network analyzer to be 1.2-j60.2Ω at 403MHz so a 900fF series capacitor and a 47nH shunt inductor were used and provides a 9dB voltage boost [64]. Devices like BAW or FBAR resonators with very large Q can also be used to provide a passive voltage boost off-chip.

LNA and Mixer

As seen in Figure 43, the LNA architecture is a single-ended cascode amplifier. The single-ended architecture was chosen to save power. A 4-bit binary-weighted current DAC allows the LNA to trade-off between power and RF gain. An off-chip inductor and a 4-bit binary-weighted capacitor bank on the LNA's output allows the LNA to be optimally tuned to the MICs band over process variation.

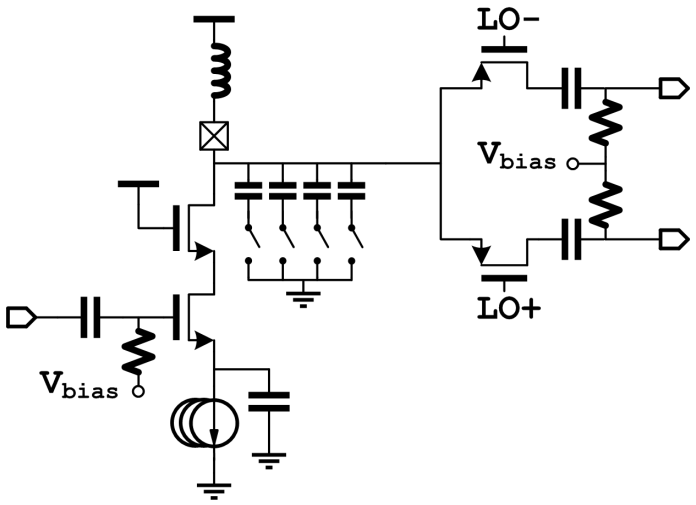

Figure 43: LNA and Mixer

The mixer is a single balanced passive switching mixer. The output of the LNA connects to the drain of 2 PMOS devices and the differential square-wave LO signal switches the gates.

PLL-based LO

Figure 44 shows the PLL-based LO which allows the receiver to tune the channel frequency with 62.5kHz resolution. The PLL is analog with a charge pump and type-2 external loop filter which locks to a 31.25kHz reference signal from a crystal oscillator. The 31.25kHz reference frequency is fed directly from the SoC's crystal oscillator. Such a low reference frequency will impact the lock time of the PLL, but also allows for much lower power consumption. Because the CDMA WRX is responsible for synchronization, the narrowband receiver will not need to duty cycle as much, reducing the need for a quick lock time. A standard tri-state PFD is used and a charge pump with active feedback is used to control symmetry between the UP/DN pulses as well as minimize charge sharing on the loop filter [69]. The VCO comprises of an LC oscillator with an off-chip inductor. A course-tuning capacitor bank sets the VCO to the appropriate frequency range while a fine-tuning varactor is controlled by the PLL to lock to the

desired frequency. The LC VCO is then squared using a differential amplifier followed by a self-biased broadband amplifier.

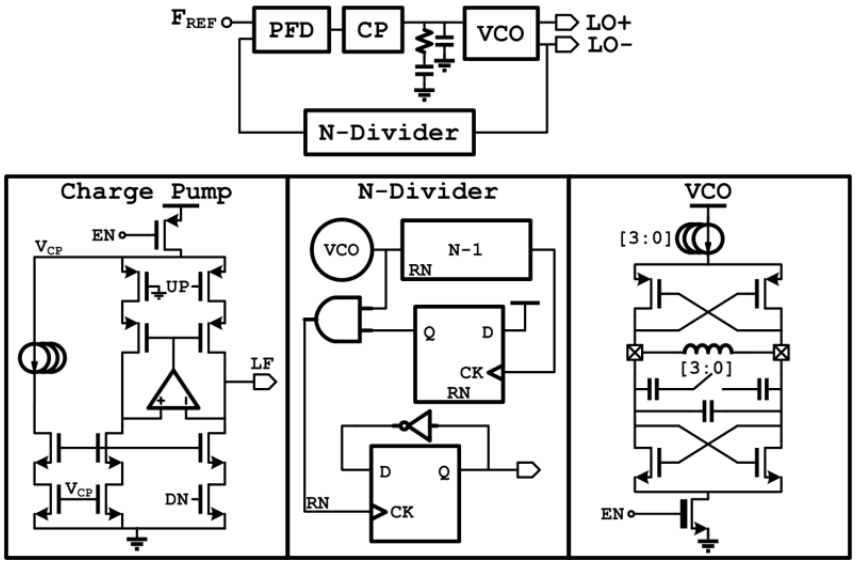

Figure 44: LO PLL

Because the divide value is so large (~13,000), good jitter performance in the divider is critical. The middle box in Figure 44 shows the architecture of the programmable divider. A divide value of N/2-1 is programmed into the divider and when N/2-1 is reached the final output stage is enabled. The final N/2 clock and the subsequent divider output change is then controlled directly from the VCO's clock in order to remove any accumulated jitter from the N/2-1 divides.

IF Gain Stages

The IF gain stages consist of cascaded differential amplifiers followed by a differential to single-ended converter and a drive buffer. Figure 45 shows the different topologies. The first stage is a differential amplifier with resistive load and a capacitor between the source nodes of the input devices, which eliminates accumulated offset voltages so CMFB is not needed [35]. The second stage is a typical resistor loaded differential amplifier. A differential to single-ended amplifier and drive buffer finish the IF stages at which point the signal passes to the external

channel select filter. To account for process variation the differential amplifiers have current tuning DACs that can modify the current through the amplifier within 10% of spec. In addition the amplifier is DC biased using an on chip voltage reference with switchable resistor ladder to ensure the common mode remains within range over process variation.

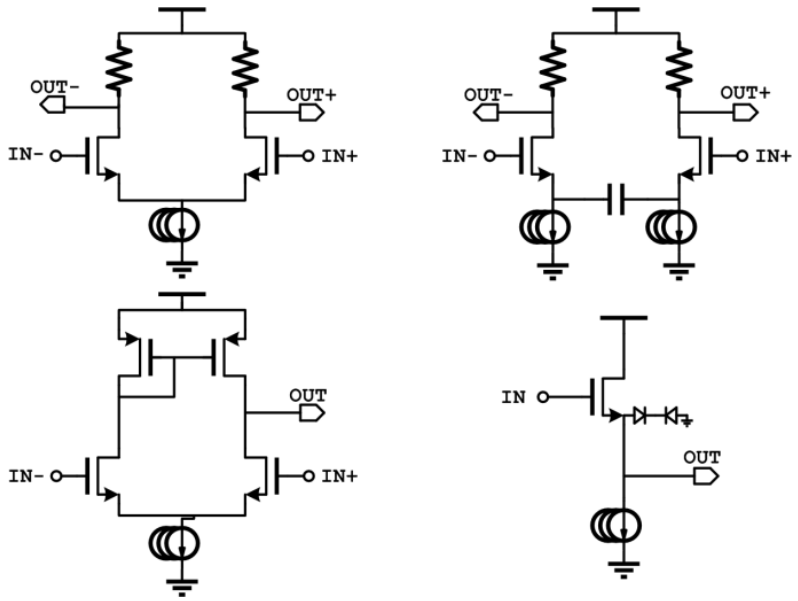

Figure 45: IF Gain Stages

IF PLL and Demodulation

The IF PLL is a charge-pump based PLL that uses the 10.7MHz received input as a reference frequency. The VCO is a ring oscillator and has a divide value of N=1. The critical output of the PLL is actually not the VCO, but the loop filter because the transient response of the loop filter indicates the direction of phase change of the received signal.

The demodulation stages consist of another charge pump identical to the one used in the IF PLL and the storage capacitor. The voltage on the capacitor, which is reset to V_{REF} at the beginning of the bit interval, and a stable reference of VDD/2 (V_{REF}) are then compared using a dynamic comparator that is clocked just before the end of the bit slice interval and produces a 1/0 output. Figure 46 shows the demodulation block diagram in more detail.

Demodulation PLL

Bit slicing and synchronization is accomplished using a demodulation PLL which outputs a 187.5kHz square wave, and the Rx FSM. Each stage of the demodulation PLL's ring oscillator VCO is brought into a MUX. An NBRx-specific synchronization header is sent at the beginning of every packet which the Rx FSM uses to synchronize the receiver. The 101010… sequence allows the Rx FSM to tune the MUX selection and calibrate the phase of the Demodulation PLL to be synchronized with the received IF signal. Three phases from the VCO are brought out, two are NANDed together to create the Bit Slice pulse that resets the capacitor and the third phase clocks the comparator just before the bitslice pulse is asserted and demodulates the data based on the voltage on the capacitor in relation to V_{REF}.

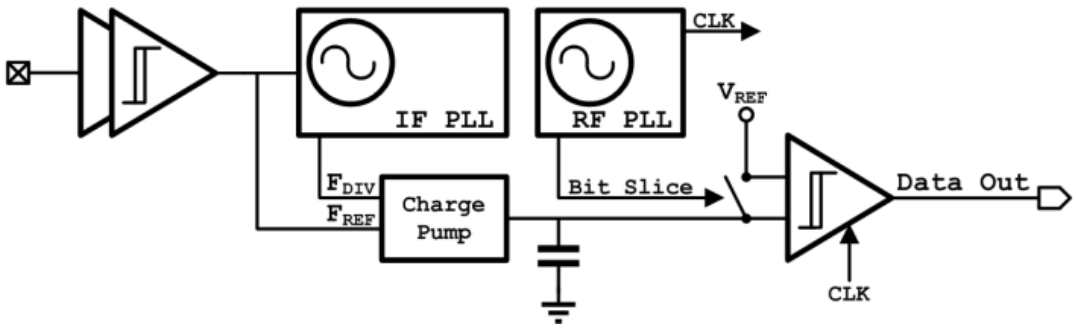

Figure 46: PSK Demodulation

Receive FSM

The Rx FSM processes the incoming bitstream from the narrowband receiver and writes it into a receive memory buffer where it can be accessed by the main processor. At the beginning of a packet the NBRx synchronization header is received and the FSM controls the demodulation PLL's 17-bit MUX select value to synchronize with the data. Following the sync header is the 802.15.6 PLC preamble. According to 802.15.6, two different preambles can be sent, but this FSM only looks for one specific preamble [17]. After the preamble the PLCP header is sent

which is typically used to convey PHY layer information. However, since the NBRx does not have all the configuration options available, this information acts as another preamble. The data information that follows the PLCP header is then stored in its raw form in a 2kB memory array. Any post-processing or MAC level functionality is performed by the microcontroller.

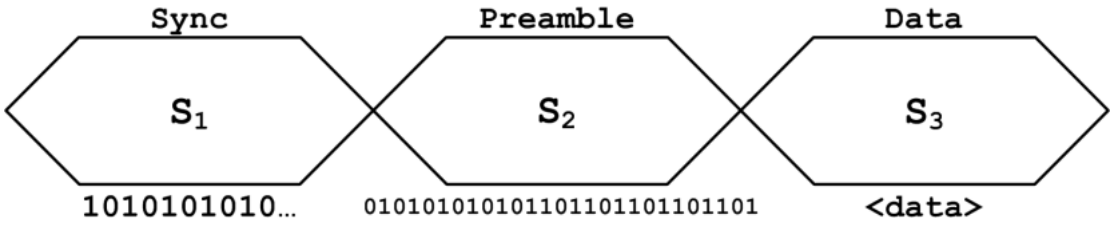

Figure 47: Rx FSM Pattern for Simplified 802.15.6 Packet Reception

Sleep Power

Sleep power performance is critical for a receiver that may spend much of its time inactive on a sensor node. To optimize sleep power efficiency the design incorporates multiple thick oxide sleep headers, despite the resulting increase in active power. Multiple Enable/Sleep headers are used to modularly control which circuits sleep and which remain active which allows for a tradeoff between power consumption and startup time.

Measurements

The IC is fabricated in an IBM 130nm CMOS process. Figure 48 shows the die photo of the fabricated chip. The total size is 0.96mm x 1.25mm, and the receiver occupies 0.28mm^2 without pads. The LNA, PLL-based local oscillator, passive mixer, IF gain stages, and demodulation PLL have all been tested on a breakout chip. The IF PLL and demodulation circuits have not been tested yet because the resistively loaded differential preamplifier of the comparator was sized improperly resulting in an improper bias of the amplifier and preventing the received signal

from properly being demodulated. This was fixed for the full SoC tapeout, but the SoC has not been verified to a point at this time where the radios can be accessed.

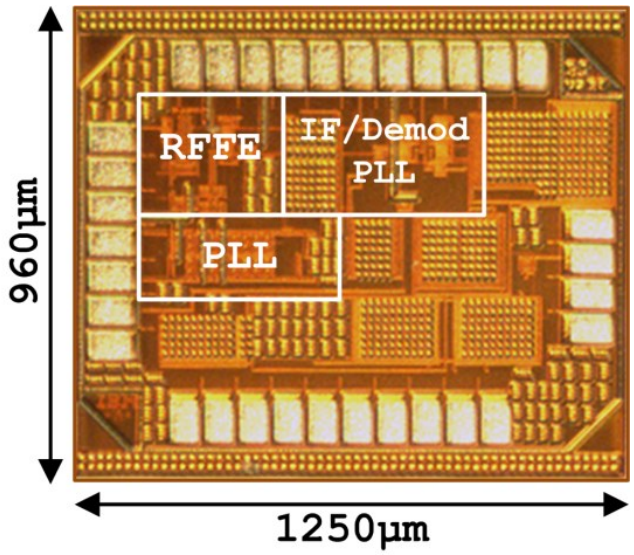

Figure 48: NBRx Die Photo

For the reported receiver, Table 10 shows the power breakdown and performance specifications. The total power in the receiver is 291.6µW with the LNA, PLL, and IF amplifiers consuming 37.2µW, 185.4µW, and 60µW respectively. For the PLL, the VCO and squaring buffers consume 118.8µW and the rest of the PLL consumes 67.2µW. Sleep power is 778nW, which is much higher than anticipated. This is because each of the multiple beta multiplier current references contains a leakage path when in sleep mode and drains power. The data-rate is the 802.15.6-specified 187.5kbps which results in an energy per bit of 1.56nJ/bit. All values are quoted at a sensitivity of –60dBm based on the spectrum of the received signal at IF.

Power [µW]		Performance	
LNA	37.2	Data Rate	187.5kbps
VCO	118.8	Energy/bit	1.56nJ/bit
PLL	67.2	Sensitivity	-60dBm
Mixer	0	Conversion Gain	39dB
IF	60		

Regulation	8.4	Sleep Power	778nW
TOTAL	**291.6**		

Table 10: NBRx performance

In order to work in a multi-user environment as well as meet the WBAN standard, channel selection is critical. Figure 49 shows the frequency content of the PLL-based LO across all 10 channels. Changing channels is digitally controlled by programming the divider through a scan chain. Note that this plot was measured through the LO feedthrough of the LNA so it does not represent an accurate depiction of the noise performance of the PLL. Noise performance can be interpreted from the sensitivity of the entire receiver.

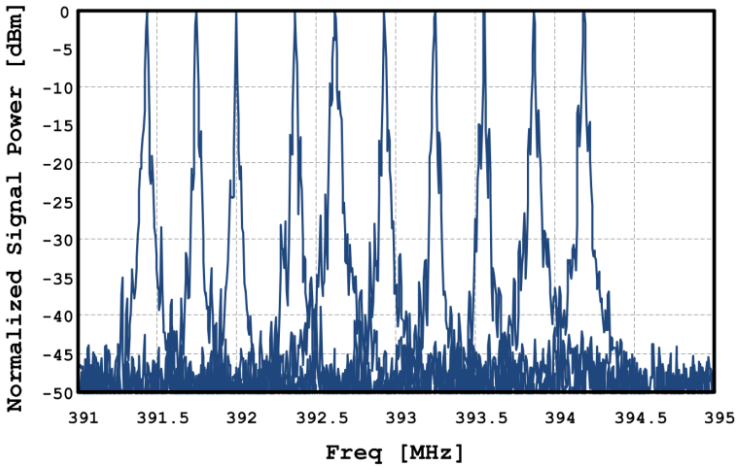

Figure 49: NBRx Channel Selection

Figure 50 shows the normalized output spectrum of the IF filter compared to the IF spectrum generated by a vector signal generator. A $\pi/2$ DBPSK signal was generated with a 1-0-1-0 repeated bit stream, not a random bit stream, which is reflected in the shape of the spectrum. Note that the large noise is largely due to the measurement setup. The measurement was taken from the output of the off-chip filter and had an impedance mismatch with the spectrum analyzer. This approach is acceptable because it shows the trend that the data is being received and is not meant to be an accurate representation of the noise in the system.

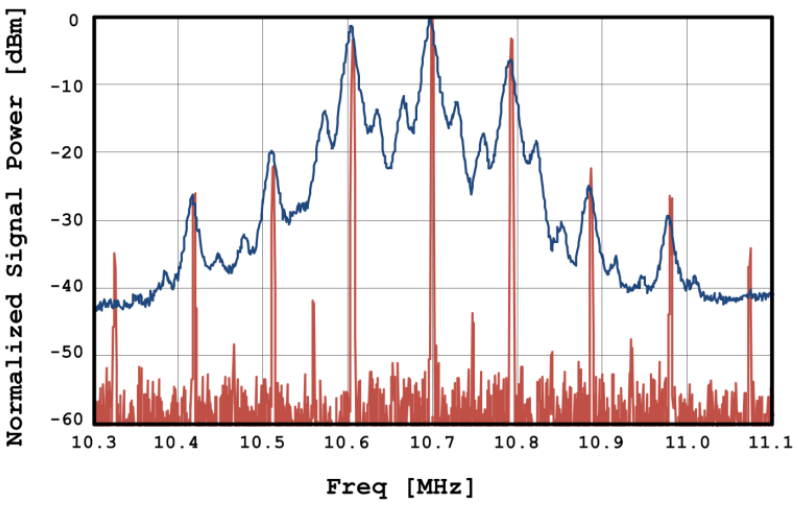

Figure 50: IF Spectrum of Received Data

Since the demodulator could not be tested, transient simulation results are shown in Figure 51 to demonstrate its functionality. The top plot in Figure 51 shows the transmitted and received data, while the middle plot shows the PLL's loop filter (blue) and the effect of charging the storage cap with the secondary charge pump (red). It can be seen that small changes in the loop filter result in measurable changes to the storage capacitor. The black dashed line represents the reset signal that comes from the demodulation PLL which pulls the storage capacitor back to V_{REF}. The final plot shows the demodulated signal at the output of the comparator's pre-amplifier and at the input of the comparator. The preamplifier is a resistive-loaded differential amplifier with the demodulated signal (red line in the middle plot) and V_{REF} as inputs. The comparator is then clocked to make the final decision. The data output is sampled right before the next bit slice (seen as the black dashed reset line in the middle plot).

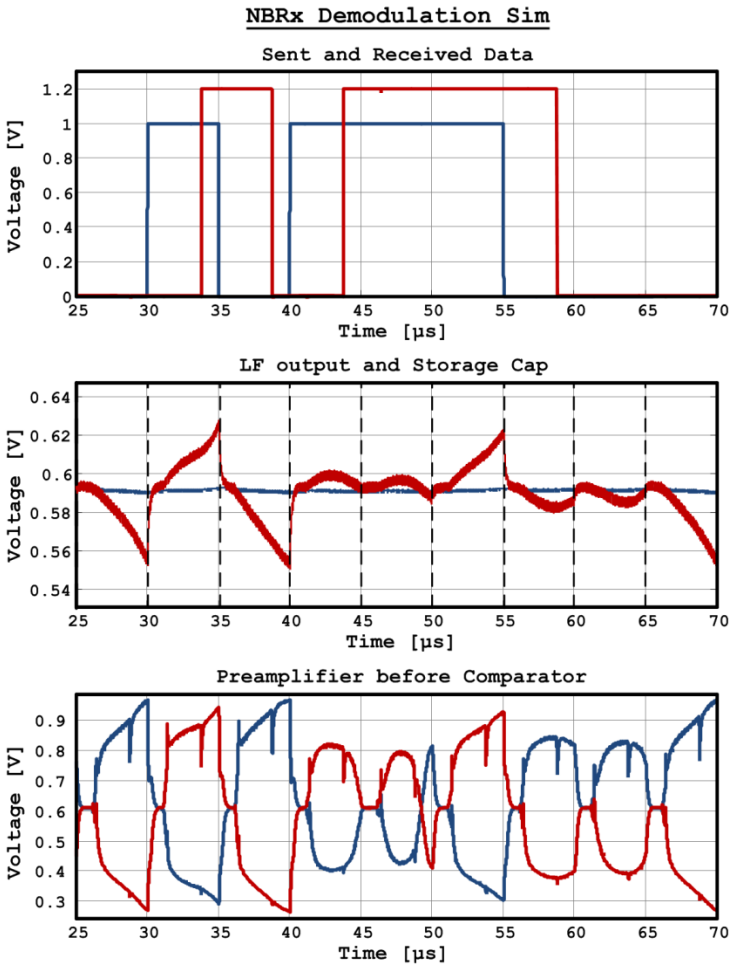

Figure 51: NBRx Demodulator Simulation

Power numbers for the Tx and Rx FSM were not possible to measure since they share VDD with many other blocks, but simulation results show that the Tx FSM consumes 294nW and the Rx FSM consumes 576nW, both at 0.5V and a 187.5kHz clock rate. Successful streaming to the UWB Tx was demonstrated using a logic analyzer to verify functionality of the Tx FSM. For the Rx FSM, data vectors were clocked to emulate the sync and preamble headers and the correct payload was verified in the Rx SRAM buffer.

Table 11 shows a comparison with other low power receivers. It is difficult to compare equally against any single receiver because of design and architectural differences of the NBRx. [32] is the only other PSK-modulated low power receiver in the survey from Figure 7. It uses

injection-locked oscillators to demodulate the PSK data. It achieves a high data rate, but the architecture limits it to low sensitivity performance. [37] is a 3-channel receiver that that has a sensitivity of -67dBm and an energy per bit of 0.18nJ/bit. However, this receiver uses fixed FBAR filters for channel selection that are neither variable or scalable. [70] is a 2.4GHz receiver that has an integrated PLL. As can be seen from the power consumption, added complexity comes at a cost. This receiver has a power consumption of 1.1mW and an energy per bit of 5nJ/bit.

	This work	[32]	[37]	[70]
Power [μW]	292	120	180	1100
Frequency [MHz]	402-405	300	2400	2400
Modulation	π/2 DBPSK	BPSK	OOK	BPSK
Channels	10	1	3	1
Data Rate [kbps]	187.5*	1000	1000	100
Energy/bit [nJ/bit]	1.56	0.12	0.18	5
Sensitivity [dBm]	-60**	-34	-67	-90
VDD [V]	1.2	1.0	0.7	1.0
Die Area [mm^2]	0.28	0.04	~2	--
Process [nm]	130	90	65	180

Table 11: Comparison Against State of the Art
*** Data Rate determined by 802.15.6 WBAN standard**
**** Estimated sensitivity since NBRx has not been fully tested with SoC**

3.3. Complete WBAN Transceiver

The three main radios, combined with the Tx and Rx FSMs provides the complete transceiver for the EKG WSN. The microcontroller controls any MAC layer functionality. After the EKG WSN collects enough EKG data, or detects atrial fibrillation, the microcontroller with enable the UWB transmitter. Once enabled, the UWB transmitter stays in a low power mode until the Tx FSM sends bits to be converted into pulses and only then does the UWB's oscillator turn on and consume a significant amount of power.

Receiving packets is more complicated. Because of its power consumption the NBRx stays asleep while the CDMA WRX stays on all the time. Once the CDMA WRX detects a packet that signals the need for data reception beyond the CDMA WRX's capabilities the microcontroller will turn on the NBRx. Then the Rx FSM will synchronize during the sync packet and will store any raw data that comes after the preamble. Processing the received data is the responsibility of the microcontroller. Table 12 summarizes the important specs of the entire transceiver.

Spec	Radio	Value	Unit
Power	CDMA WRX	112	nW
Sensitivity		-45	dBm
Data Rate		12.5/15	kbps
Energy/Wake-Up		161	pJ/bit
# Simultaneous Codes		4	
Power	UWB Tx	4.18	µW
Data Rate		187.5*	kbps
Output Power		-28.9	dBm
Energy/bit		39.7	pJ/bit
Power	NBRx	292	µW
Sensitivity		-60**	dBm
Data Rate		187.5*	kbps
Energy/bit		1.56	nJ/bit

Table 12: Summary of Specs for WBAN Transceiver
*** Data Rate determined by 802.15.6 WBAN standard**
**** Estimated sensitivity since NBRx has not been fully tested with SoC**

3.4. Conclusion

This chapter presented a wireless transceiver used to enable wireless communication for an EKG sensor node that runs on harvested energy. Using asymmetric communication the transceiver has a 187.5kbps data rate and has an average transmit power of 4.18µW at 4GHz and a receive power of 292µW in the MICS band. In addition a 112nW CDMA WRX is used for

ultra-low power wireless interrupts and synchronization. Both the receive and transmit sections have 2kB buffers for data storage and integration with the rest of the sensor node. While low power was prioritized, it is recognized that in order to be useful in a deployable environment, advanced functionality is needed, even at the cost of increased power. Because of this the CDMA WRX looks for 4 simultaneous 15-bit wake-up codes and the NBRx can tune to 10 different 300kHz channels within the MICS band in order to allow for multiple users as well as being compatible with the 802.15.6 WBAN standard.

The power density of a TEG is 40μW/cm^2 [71]. Assuming a 1cm^2 TEG and an estimated average power for the SoC to be 19μW, as per the design spec on the untested SoC, which leaves roughly 20μW average power for the transceiver in order to not drain the storage capacitor on the node. The transceiver can both transmit and receive wake-up signals without draining the storage capacitors. Data reception through the NBRx would have to be duty cycled at a rate of 6.6%.

3.5. Acknowledgements

I would like to thank Ben Calhoun and his group at UVA for their help in developing the interface between the transceiver and the integrated microcontroller, specifically Yousef Shaksheer and Alicia Klinefelter. Alicia Klinefelter also wrote the Verilog code implementing the Tx/Rx FSM based on my system specifications. In addition, Prof. Kamin Whitehouse from UVA was instrumental in the design changes in the WRX for this transceiver.

CHAPTER 4

Kilometer Range Communication using Harvested Energy

4.1. Introduction

So far designing RF circuits for IoT applications has been focused on short range communication, which is true both in the circuits presented in previous chapters and in academia in general. However there are several applications where sensor nodes would benefit from having longer range, while still maintaining the qualities of the short range radios. Such a sensor node would not have the size constraint that body area networks impose and therefore more aggressive energy harvesting can be used which allows for a larger power budget. Many commercial companies are making long range communication chips for IoT including Analog Devices [72], Texas Instruments [73], and Semtech [74].

Long range radios would be useful in many different scenarios. A simple example would be for remote environmental sensing applications that do not require high data rates and may need to communicate once in a while. For example, Analog Devices introduced the CN0164 Low Power, Long Range ISM Wireless Measuring Node [75]. This system consists of a low power temperature measurement node that wakes once a minute, measures the temperature, transmits the information, then returns to sleep.

However, it will be difficult to operate any of the mentioned radios off harvested energy supplies. Using the Texas Instruments CC1101 as an example, assume the average power for the

sensor node is required to be 50μW. The CC1101 can be configured to communicate at 1kbps and consumes 48mW in receive mode and 48mW in transmit mode. Duty cycling to make the average power 50μW would lead to an average data rate of only 1.04bps, which is optimistic because it does not take into account startup time. If the sensor node has to transmit a 100-bit packet at 1kbps then 1.6mF of storage capacitance is required in order to prevent the supply voltage from dropping more than 1V.

The goal of this long range transceiver (LRTRx) design is to enable long range communication that runs of harvested energy effectively. An average power consumption of 50μW is the target and the instantaneous power should be kept below 1mW for receive and a few mW for transmit. The LRTRx is designed to communicate chip to chip and also with the CC1101 with greater than 1km range. The next section will look at the feasibility of such a design from a communication theory perspective.

4.2. Communication Theory

In order to improve upon what is commercially available, the first goal is to understand what is possible. One of the principle techniques used in the previous circuits that enabled ultra-low power RF design was sensitivity reduction, but extending the communication range eliminates the ability to use similar techniques. Instead power reduction will come from data rate, bandwidth, and duty cycling. Using Shannon's classic communication capacity theorem [76], shown in Equation 2, a first-order calculation of the power requirements for kilometer range communication can be analyzed.

$$R_b < BW \cdot log_2\left(1 + \frac{P_{RX}}{N_0 BW}\right) \quad (2)$$

Shannon states that the data rate (R_b) is a tradeoff between the bandwidth (BW), the received signal power (P_{RX}), and the noise spectral density (N_0). Equation 3 shows the same formula, but written with respect to transmitted energy [77].

$$\frac{E_b}{N_0} > \frac{2^{Rb/BW} - 1}{Rb/BW} \tag{3}$$

where E_b is the energy per transmitted bit.

To find the optimal data rate and bandwidth for our application, a link budget is calculated that makes the following assumptions: the communication distance is 1 kilometer, the frequency of operation is the 433MHz ISM band, the efficiency between circuit power and power radiated at the antenna is 10%, and the noise figure (NF) of the receiver is 15dB. Looking at Shannon's theorem from a hardware perspective, we can say that the energy per bit to noise ratio (E_bN_0) for the transmitter subtracted by the path loss through the air and the E_bN_0 of the receiver must be greater than 0, as shown in Equation 4 and in more detail in Equation 5:

$$\frac{E_b}{N_{0_{Tx}}} - P_L - \frac{E_b}{N_{0_{Rx}}} > 0 \tag{4}$$

$$(pwr * T_b) - N_0 - 20 \cdot log_{10}\left(\frac{4\pi f}{c}\right) - 20 \cdot log_{10}(1km) - \frac{2^{Rb/BW} - 1}{Rb/BW} > 0 \tag{5}$$

Shown in Figure 52, the data rate (R_b) has a much more significant impact than bandwidth (BW) on circuit power. In order to keep the circuit power in the sub-mW range while keeping the bandwidth low enough for multiple channels within the 433MHz ISM band, we found that optimal specs for power and application performance is a 1kbps data rate and a 200kHz bandwidth, which is 1/8 of the 433MHz ISM channel bandwidth, providing the opportunity for multiple channels. Duty cycling will be utilized so the circuit can operate off harvested energy, but will lower the effective data rate by the rate the duty cycling is applied.

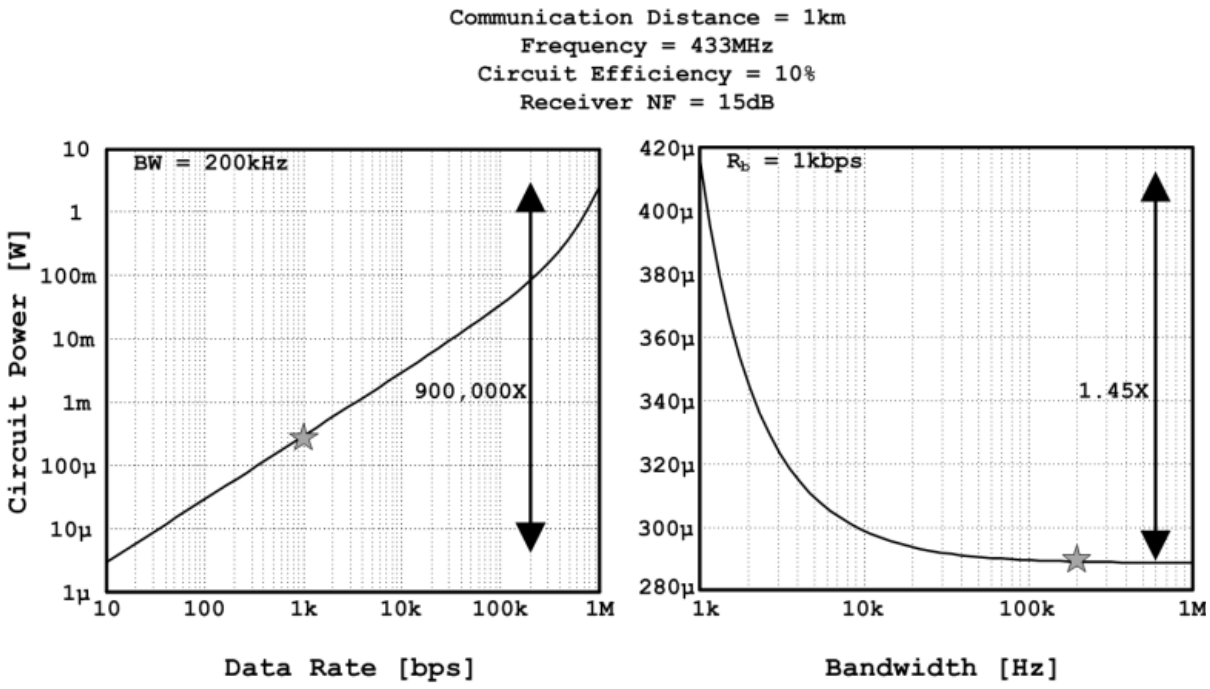

Figure 52: Using Shannon's Theorem to Predict Circuit Power for Kilometer Range Communication

Reducing data rate in an effort to increase communication range is an intuitive tradeoff because it increases the signal energy per bit which directly improves the receiver's ability to detect the transmitted signal. More signal energy improves the signal's energy per bit to noise ratio (E_bN_0) which will result in an improved sensitivity metric in the receiver. To validate this across academia, the same data that was used to plot the low power radio survey in Figure 7 is used again in Figure 53, but this time plotting data rate vs. sensitivity.

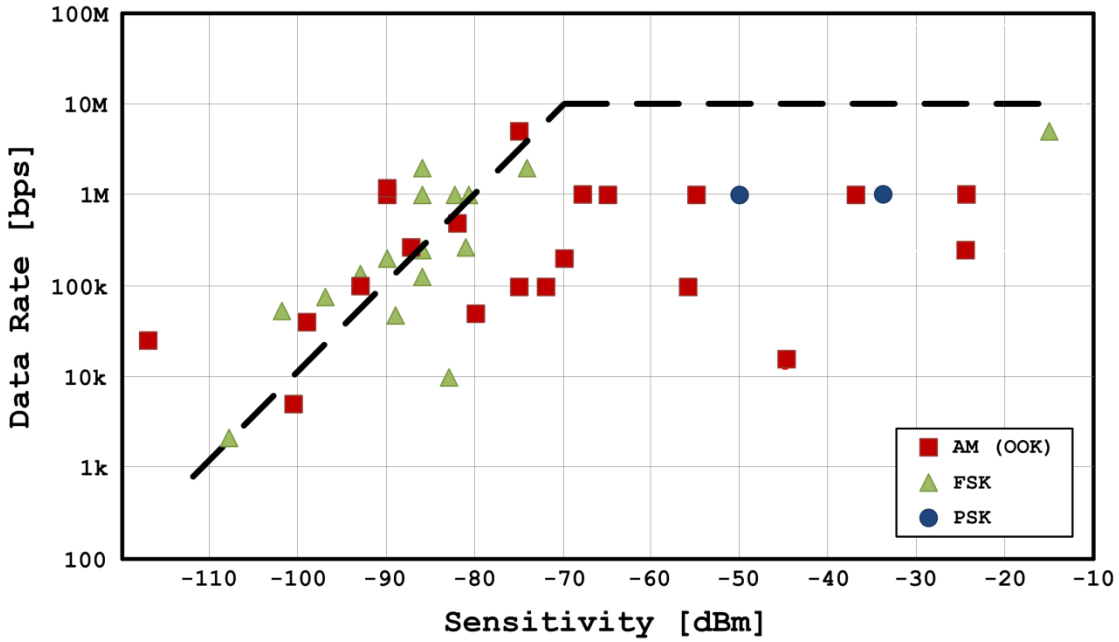

Figure 53: Low Power Radio Survey: Data Rate vs. Sensitivity

At sensitivity levels from -70dBm and to the left, there is a correlation showing that a 10dBm change in sensitivity results in about a 10X change in data rate, which makes sense because the faster data rate has less energy per bit that can be detected at the receiver. Much like in Figure 7, there is an empirical limit, this time at -70dBm, where the energy per bit is high enough that other factors dominate the data rate performance. This shows that low power radios are generally good for sub-1Mb/s data rates. Looking at the type of modulations listed, it also shows that FSK is the most consistent modulation for excellent sensitivity and trades off well with data rate.

Figure 54 shows the same data as Figure 53, but with power included via data point colors. Looking at sensitivity from -70dBm and to the left, where there is a 10dB sensitivity to 10X power correlation, the majority of data points along the fit line are FSK and OOK modulation radios, but the FSK-modulated radios in general consume much less power.

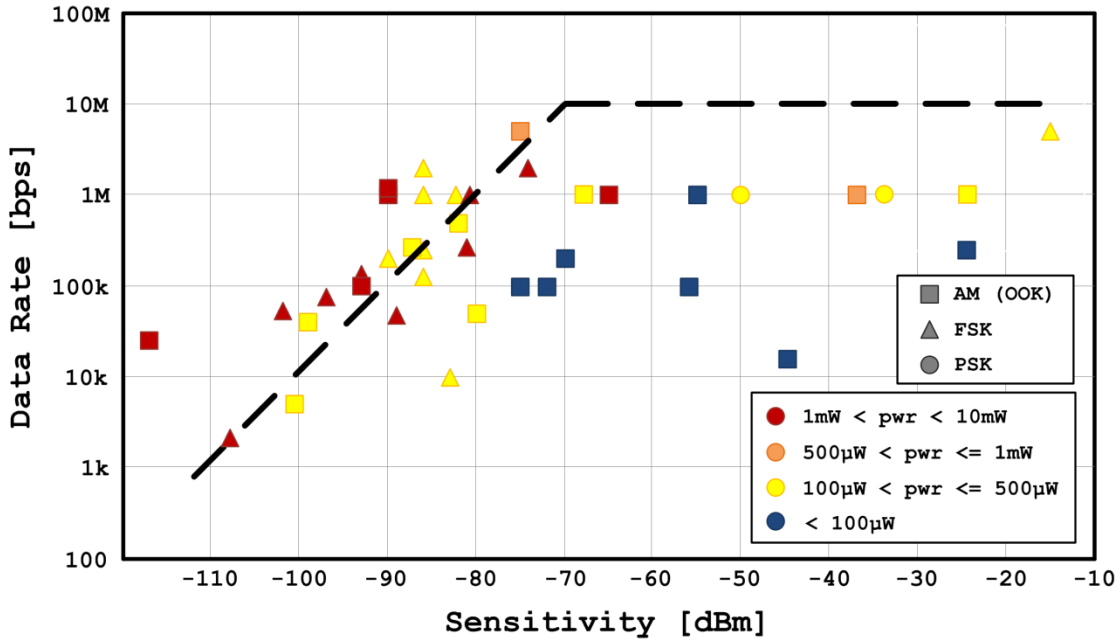

Figure 54: Low Power Radio Survey: Data Rate vs. Sensitivity vs. Power

4.3. Long Range Transceiver Chip Design

Based on the results from the previous section, the transceiver specifications were chosen to make it operable with the TI CC1101.

Texas Instruments CC1101 Configuration

The Texas Instruments CC1101 is a sub-GHz transceiver designed specifically for long range and low power applications, intended mainly for the ISM band, and is a commercial chip that closely resembles the LRTRx design goals. The LRTRx was designed to be compatible with the CC1101 under the following configuration settings in Table 13 which are derived from the results of the previous section and the available configuration settings in the CC110:

Specification	Value	Unit
Frequency	433	MHz
Modulation	2-FSK	
Data Rate	1001.2	Baud
Channel BW	203.125	kHz

FSK Frequency Deviation	50.781	kHz
IF Frequency	863.281	kHz

Table 13: Texas Instruments' CC1101 Long Range, Low Power Transceiver Configuration

The 433MHz band has a 1.74MHz total channel bandwidth so the transceiver will have 8 total channels and these values line up closely with the theoretical values computed above. These values were calculated from the CC1101 datasheet [73].

The low power CC1101 consumes 16.0mA in receive mode at 433MHz at 1.2kbaud with register settings set to reduced current and the input signal at the sensitivity limit of the receiver, as stated by the datasheet. The transmitter consumes 16.0mA when outputting 0dBm at 433MHz. VDD is 3.0V, bringing the total power for transmit and receive to 48mW each.

System Overview

The LRTRx is designed to comply with the CC1101 specs in Table 13. The link budget for the LRTRx is shown in Table 14 which computes the required receiver sensitivity and noise figure to achieve kilometer range communication.

Spec	**Value**	**Unit**
Center Frequency	433	MHz
Data Rate	1001.2	bps
Bandwidth	203.125	kHz
Distance	1000	Meters
Transmit Power	0	dBm
Path Loss	85.18	dB
Antenna Gain	0	dB
Noise Power	-143.99	dBm
Min. Eb/No	12	dB
Link Margin	10	dB
Receive Sensitivity	**-85.18**	**dBm**

Receiver Noise Figure	36.81	dBm

Table 14: LRTRx Link Budget

The difference between the link budget and the previous capacity calculation is that power between transmit and receive must be reasonable based on hardware constraints. Reducing the output power of a power amplifier only works as long as the increased sensitivity burden placed on the receiver is reasonable and does not dramatically change the overall receiver power. And since the receiver spends much of its time in the active state listening for packets, low receiver power is critical to the overall system design. Setting the transmit power to 0dBm sets the receiver's required sensitivity to be -85dBm for kilometer range communication. While the transmitter's power will be high, greater than 1mW, the receiver's power will be in the μW range. Communicating with the CC1101, with its potential +10dBm transmit output power and -112dBm receive sensitivity should enable kilometer range communication with 10dB design margin on the receiver sensitivity and about 25dB design margin for the transmitter. A block diagram of the system architecture can be seen in Figure 55.

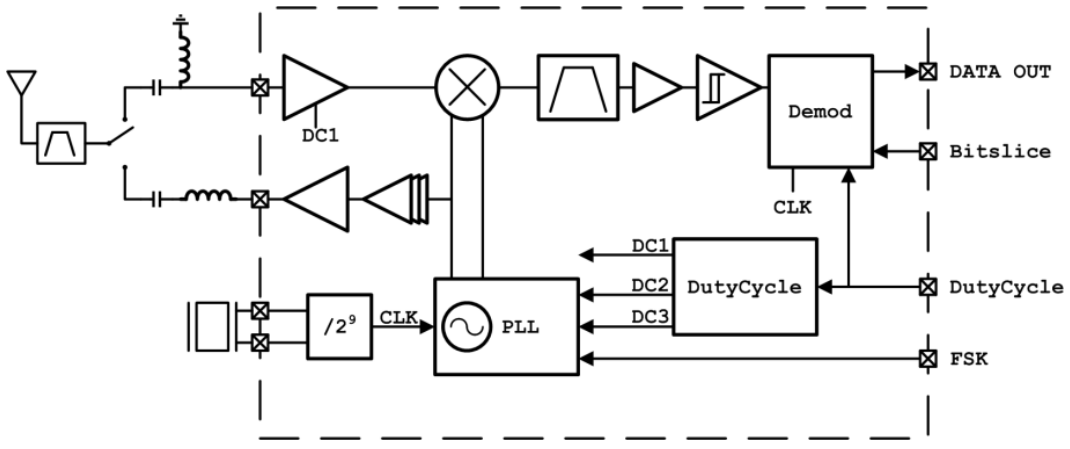

Figure 55: Block Diagram of the Long Range Transceiver

The receive portion of the LRTRx consists of a single-ended common source cascoded LNA with a passive mixer and transimpedance amplifier to improve linearity. The IF stage contains a

6^{th} order gm-C filter which provides gain and channel selection and is followed by another gain stage. The signal is then squared up with a comparator and driven into a small baseband processor which uses the IF frequency as a counter and determines the bit value based on the number of counts in a certain bit slice. Bit slicing and synchronization is controlled through an off-chip FPGA. In scenarios with sufficient SNR margin and low enough data rate, bit-level duty cycling is possible to reduce transmit and receive power and is also controlled from an off-chip FPGA.

The transmitter consists of a Class-E power amplifier running off a 0.5V supply, but is clocked using buffers running at 1.2V. The larger supply on the buffers increases their drive strength, allowing them to clock the large PA switch more efficiently, improving overall efficiency in the transmitter.

A 26MHz crystal with on-chip oscillator is used as a reference clock; the same reference frequency used by the CC1101, which will improve synchronization between the CC1101 and the LRTRx. Because the LRTRx does not use a fast clock for power saving purposes, the 26MHz signal is immediately divided by 2^9 to establish a reference frequency of 50.78kHz for the PLL providing the LO signal. The reference frequency divider is made up of 9 cascaded D-FFs and the output of the last D-FF enables a final D-FF stage which is clocked by the 26MHz crystal, eliminating the accumulated jitter from the previous 9 stages.

Dividing down the 26MHz reference clock is a design decision that will trade-off PLL settling time and power. This decision will put an upper limit on the data rate possible for bit-level duty cycling. If the required data rate is higher than the upper limit then only packet-level duty cycling will be possible.

Transmission consumes mA of current and receiving consumes hundreds of μA, which are relatively high power for harvested energy circuits. Storing energy on a bank of off-chip capacitors, about 1-2μF is required for every bit transmitted if the power management can only allow for a 1V drop on the supply.

I was responsible for all the system-level analysis and architecture design as well as all circuit design with the exception of the g_mC filter and the crystal oscillator.

4.4. Circuit Descriptions

Receiver Front-End

The RF front-end consists of a 433MHz antenna followed by a SAW filter and an off-chip RF switch. The RF switch allows for the use of independent off-chip matching networks for both the power amplifier (PA) and the LNA. On-chip the LNA is an AC-coupled single-ended cascode amplifier with a tunable current source. Because of the relatively low RF frequency an off-chip inductor load is used and resonates with an on-chip switchable cap bank.

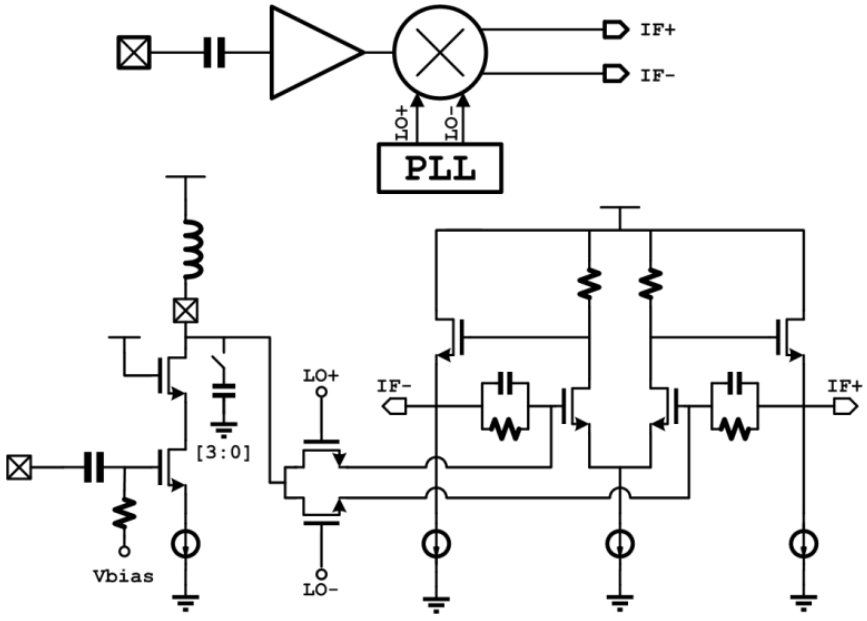

Figure 56: Receiver Front End

The mixer consists of a single-balanced passive switching stage which is controlled by the differential PLL square wave output. The output from the single balanced mixing goes into a transimpedance amplifier (TIA) to improve linearity and provide gain. The front end can be seen in Figure 56.

IF stages

Following the mixing operation, the signal is filtered by a 6-order g_m-C filter. The filter has a channel bandwidth of 200kHz and provides 20-40dB of programmable gain. Following the g_m-C filter is another gain stage and a comparator with tunable hysteresis for noise resilience.

After the comparator, the signal goes into a synthesized baseband core that is used for demodulation. The squared-up IF signal clocks a counter in between bit slice pulses, which are controlled by the off-chip FPGA. Once the bit slice pulse is asserted the baseband demodulator compares the given count with a user-programmed threshold. If the count is greater than the threshold then then demodulated data is 1, and is 0 if the count is less than the threshold. For example, if the bit rate is 1kbps then the bit slice time is 1ms. With an IF frequency of 863MHz and a 2-FSK frequency deviation of 50.78kHz, a transmitted 0 would have a count of about 813 and a 1 would have a count of 913. Setting the threshold at the center of the IF frequency would

put it at 863 and provide a count buffer of 50 counts before a bit error would be made.

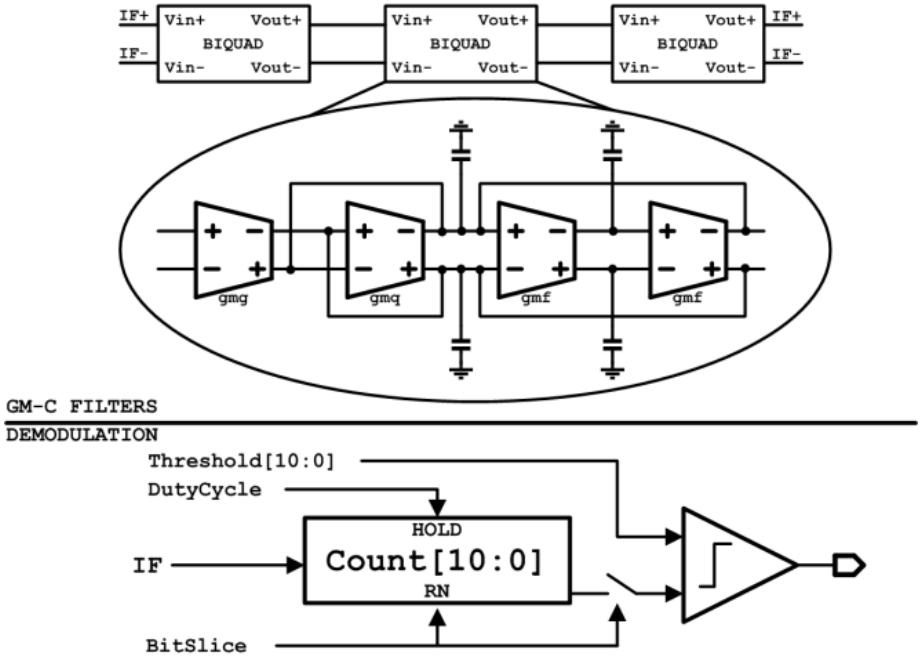

Figure 57: GmC Filter and Demodulation Block

In addition, the baseband demodulator can also be duty cycled. If the duty cycle bit is asserted the counter freezes where it is and waits until the bit slice command is asserted before comparing with the threshold. This way, in an environment with good SNR, the receiver doesn't need to receive for the entire 1ms bit slice duration to make a decision. Under the same conditions as before, but employing a 10% duty cycle, a 0 would be 81 counts and a 1 would be 91 counts. The demodulator block consumes 400nA of current in simulation. The architecture of the g_m-C filter and demodulator can be seen in Figure 57.

Transmitter

The transmitter consists of a Class-E power amplifier operating at 0.5V preceded by a 4 stage clock buffer created with standard cells, as shown in Figure 58. 2-FSK modulation is achieved from the PLL providing the LO signal. The Class-E NMOS switch is open drain connected and

the RF choke and matching network passives are off-chip, again due to the size of the passives as a result of the 433MHz RF frequency.

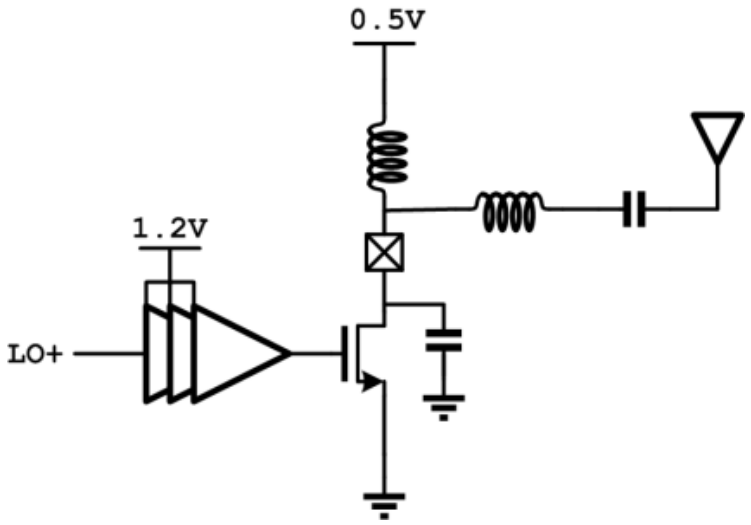

Figure 58: Class-E Transmitter

PLL-based LO

A charge pump-based PLL was designed to act as the LO for the entire system. Figure 59 shows the architecture of the PLL's main circuits. Running off the divided down 26MHz crystal oscillator, the reference frequency for the PLL is 50.78kHz, the same frequency as the 2-FSK deviation specifies, allowing for integer-N division. Because the divide value between the reference frequency and the ~433MHz output frequency is so large, the same technique that is applied to the crystal oscillator divider is applied to the PLL's divider, namely that it counts up to N-2 (1 clock cycle to remove accumulated jitter and another clock cycle for resetting the counter) and then the Nth clock is directly clocked by the VCO which removes the accumulated jitter from the rest of the divider chain. In addition a 2-bit FSK value is controlled from an off-chip FPGA and is added to the user programmable divide value. This allows for effective 2-FSK modulation. The programmable N-2 divider allows the LRTRx to frequency hop between the specified 8 different physical channels in the 433MHz ISM band.

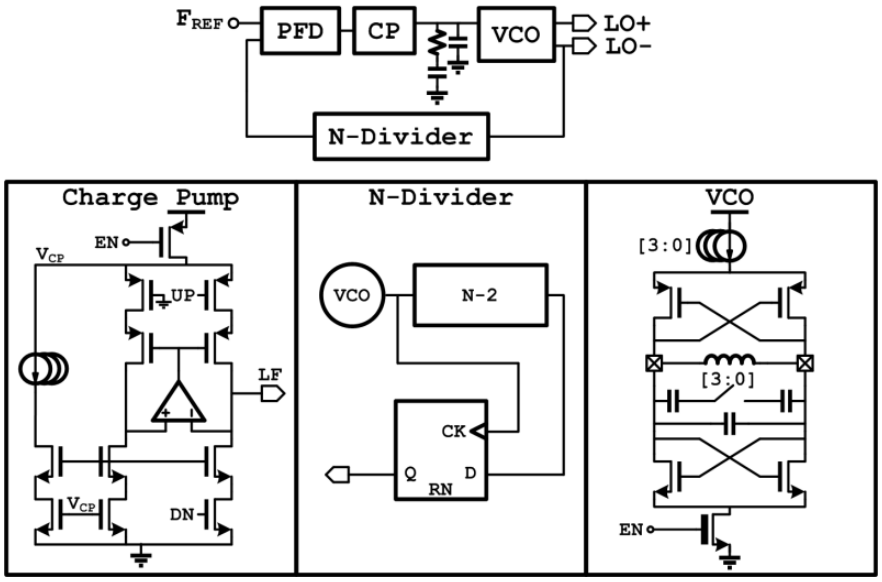

Figure 59: Integer-N PLL

Bit-Level Duty Cycling

In an environment with good SNR, bit-level duty cycling can be employed to save power, which is controlled by an off-chip FPGA. When duty cycling is asserted, a 3 stage sleep sequence is initiated with a period of 1/50.78kHz between each stage. Upon duty cycle assertion the baseband demodulator freezes its count value. On the next rising clock edge the receiver front-end is put into sleep mode. The next stage holds the PLL's PFD so that no UP/DN pulses are triggered and the last step turns off the VCO buffers and divider. Because the reference frequency to the PLL is so low and the PLL's startup time is slow, the charge pump stays on to maintain the VCO's tuning state. When duty cycling is de-asserted the baseband demodulator remains frozen (until the next bit slice is asserted) and the receiver front end, PFD, and PLL VCO are all woken up in reverse order. Since the VCO will be close in frequency, but not in phase, there will be a necessary time to lock, but that time will be faster than the original startup time of the PLL. This sequence is shown in Figure 60.

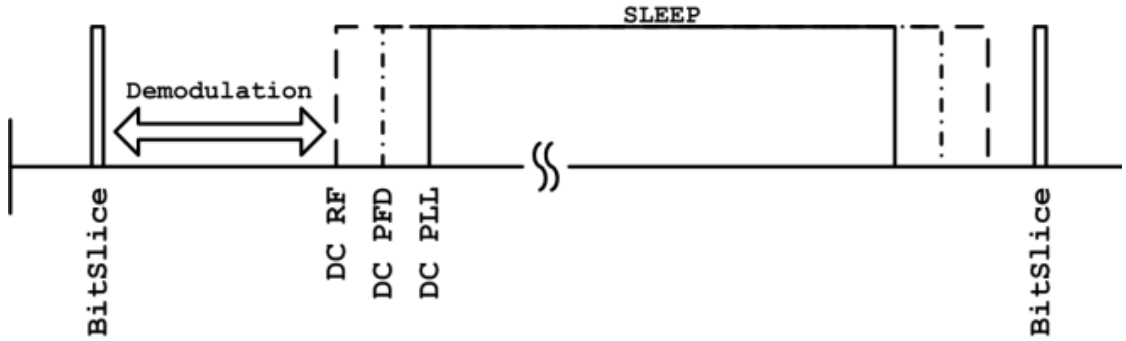

Figure 60: Bit-Level Dutycycling

Packet-level Duty Cycling

When the data rate is greater than 500bps then packet-level duty cycling is more efficient than bit-level duty cycling due to the PLL lock time. For example, at 1kpbs the PLL has a lock time less than 1 bit. Further details will be discussed in the Measurements section.

4.5. Measurement Results

The LRTRx was taped out using an IBM 130nm process node and has an area of 1.1mm^2, including pads. The entire chip runs off a 1.2V supply with the exception of the power amplifier which runs at 0.5V. A die photo can be seen in Figure 61.

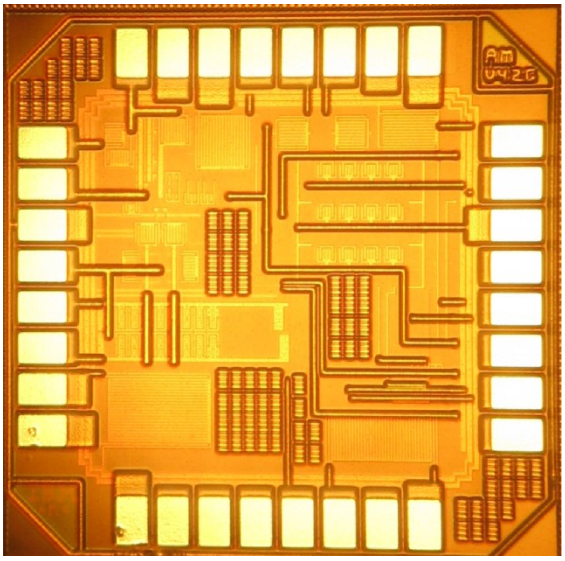

Figure 61: LRTRx Die Photo

Transmitter

The power amplifier was measured to have a 0.04dBm output power with a PAE of 47.49%. Total power on the supply for the PA is 2.13mA and the PLL and clock buffer consume 266µW and 180µW, respectively. Figure 62 shows the spectrum output of two adjacent channels both modulating a 1010 data pattern at 1kbps next to each other. A drawback from choosing the lower power integer-N PLL with a slow reference frequency is the location of spurs in the frequency regions of interest. Even so, adjacent channels see a -16.7dB dropoff. Thick oxide headers were used on the clock buffer to save sleep power, but not on the PA to maintain efficiency.

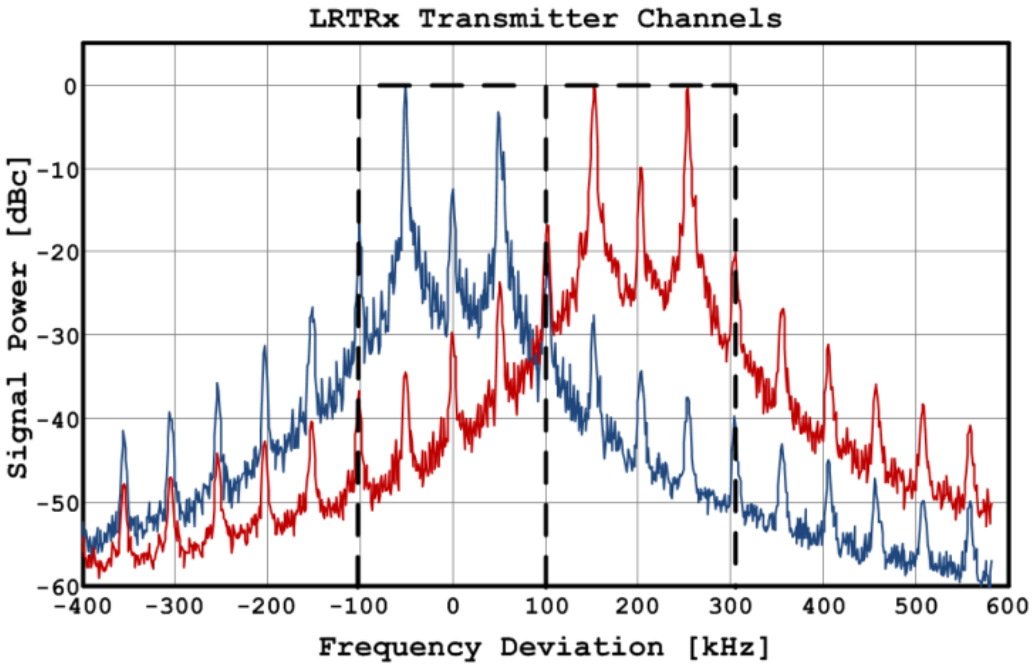

Figure 62: PLL Modulating in Adjacent Channels

Receiver

The receiver has a total power consumption of 370µW. All circuits have thick oxide headers to reduce sleep power.

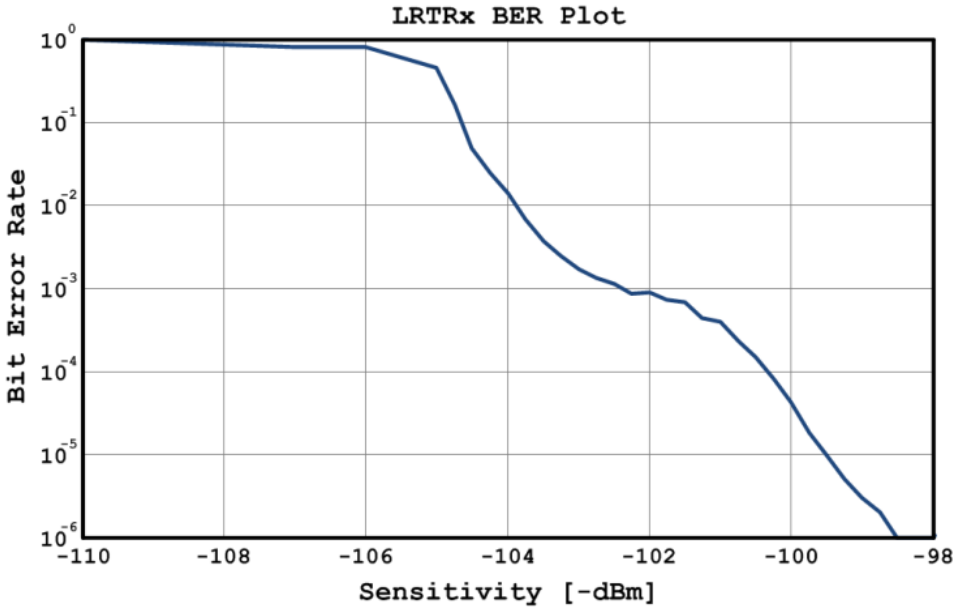

Figure 63: BER plot for LRTRx Rx

The receiver's sensitivity of -102.5dBm is measured at a BER of 10^{-3}, as shown in Figure 63. Assuming an antenna gain of 0dB, the theoretical communication range for chip to chip communication is greater than 5 kilometers based on the Friis free space path loss equation. Communicating LRTRx to CC1101 extends that range to over 20 kilometers, well beyond the design spec. In addition, Figure 64 shows how speed and data rate can be traded based on the required range of the application and the SNR of the surrounding environment.

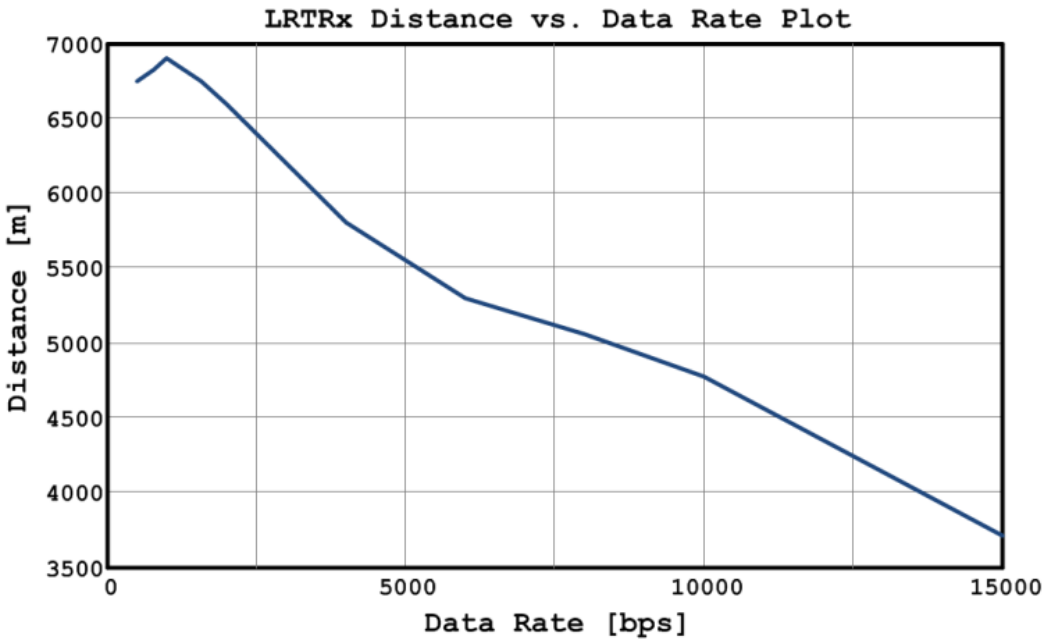

Figure 64: Data rate vs. Distance

Figure 65 shows adjacent channel rejection across all channels. In the next adjacent channel, channel rejection is measured to be 14dB and two channels over is measured to be 34dB. Adjacent channel rejection was measured by setting the desired signal power so that the BER was 10^{-4}. An interfering signal was placed in the FSK0 position (worst case frequency of the adjacent channels) and its signal power was swept until the BER fell to 10^{-3}. Channel 0 was the reference channel.

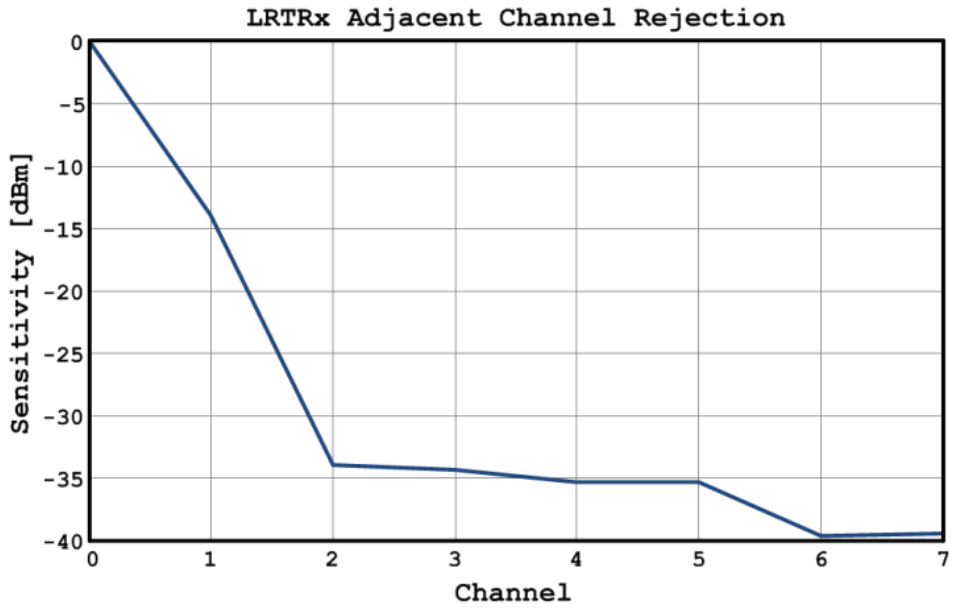

Figure 65: Adjacent Channel Rejection

Bit-level duty cycling can be employed if the required data rate is very slow. Figure 66 shows the sensitivity vs. power trade-off based on the percentage of time during each bit window that the receiver is asleep. Sensitivity and power are both normalized to the receiver performance without duty cycling. Due to the slow lock time of the integer-N PLL bit-level duty cycling only works on data rates less than 500bps.

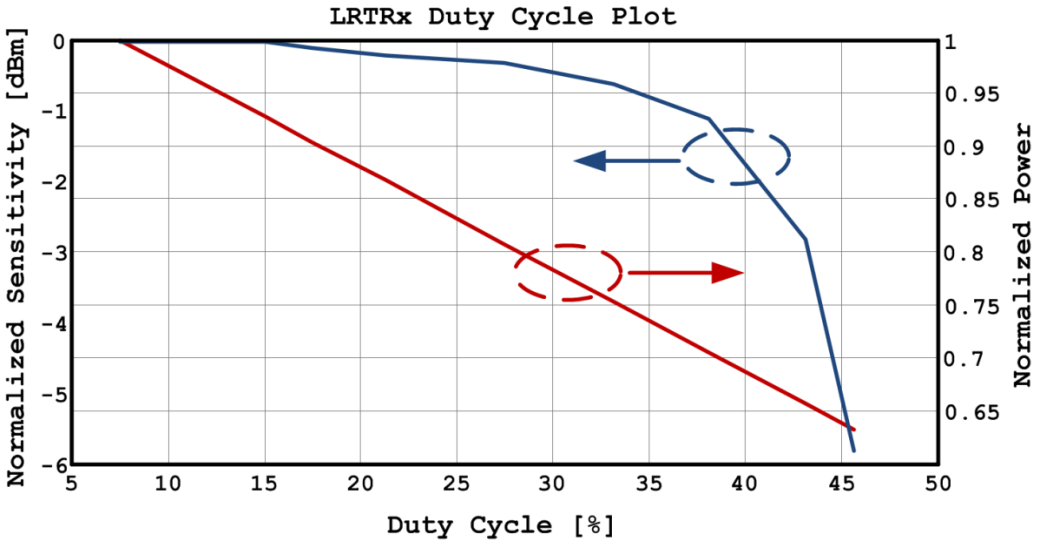

Figure 66: Normalized LRTRx Bit-Level Dutycycling (250bps)

For data rates above 500bps, packet-level duty cycling is employed. The PLL startup time dictates the startup time of the entire receiver and is measured to be 0.6 bits. This makes sending packets of 100 or more bits very efficient because the startup time per packet is less than 1 bit.

LRTRx

Table 15 shows a power breakdown for the LRTRx. The crystal oscillator (XO) and PLL are shared between both the transmitter and receiver and consume 31µW and 226µW respectively. On the receive side, the LNA and mixer consume just under 100µW while the IF gain and filtering stages plus the digital demodulation block consume 22µW. On the transmitter side, the power amplifier consumes 2.13mW in order to produce almost exactly 1mW output power at the antenna. The standard cell clock buffer consumes 182µW. Thick oxide sleep headers were used through the design to reduce sleep power, even at the expense of active power, with the one exception being the PA in an effort to optimize efficiency. The entire receiver has a sleep power of 10.2nW and the transmitter has a sleep power of 98.9nW, with 93nW of that coming from the PA. However, this can be reduced since an LDO or some other power management circuit will be used to regulate a 0.5V supply off the 1.2V supply. Disabling the power management circuit will have the same effect as a header would and will reduce the measured transmitter sleep power.

Power Consumption			
Rx		Tx	
LNA+Mixer	99µW	Power Amplifier	2.13mW
IF+Demodulation	22µW	Clock Buffer	182µW
PLL	226µW	PLL	226µW
XO	31µW	XO	31µW
TOTAL	**378µW**	**TOTAL**	**2.57mW**
SLEEP	10.2nW	SLEEP	98.9nW

Table 15: LRTRx Power Breakdown

Based on the power measurements, packet-level duty cycling can be analyzed. Measured startup time from sleep to receiving data is on average 3ms, or 3 bits and the turn-off time is less than 1 bit. The power between packets is measured to be 10.2nW for the receiver and 99nW for the transmitter. Returning to the examples provided earlier, assume the application requires a temperature measurement be sent once a minute and it is a 100-bit packet. Turn on and off time increase that length to 104 packets which is 104ms every minute. The power averaged over a minute would then be 4.55μW for transmission and 665nW for receive. The transmitter can send 11 packets per minute, or 1100 bits, and stay within the target 50μW average power budget. The receiver can actively listen for nearly 8 seconds out of every minute to stay within the 50μW power budget.

Table 16 compares the LRTRx with two commercial radios [72],[73] and one academic radio [53]. Finding direct comparisons is difficult. The commercial radios have much more overhead compared to the LRTRx so similar settings were chosen for comparison. The academic radio is a transceiver, but the modulation and frequency are different. General trends can be seen between the commercial radios and LRTRx. The main transceiver specifications match up, but the power in the Tx and Rx for the LRTRx is much lower. The penalty for the significant power reduction is seen a bit in the adjacent channel rejection and also in the PLL settle time, but PLL settle time is not significant when the data rate is so low. Comparing with the academic paper, the frequency and data rate are both much lower, which were chosen to enable kilometer range communication. The slow data rate hurts the receiver's energy/bit, but the LRTRx is lower power, higher sensitivity, and has a more efficient transmitter in addition to multiple physical channels which are important for communicating in a multi-user environment.

		[72]	[73]	[53]	This Work
Frequency [MHz]	TRx	433	433	915	433

Data Rate [kbps]		1	1	1000	1
# Channels		8	8	1	8
PLL Startup Time [µs]	PLL	N/A	712µs	none	600µs
PA output power [dBm]	Tx	0	0	-2.2	0.04
PAE [%]		--	--	--	47.49%
Total Power [mW]		57.3	48	9.1	2.57
Sleep Power [nW]		300	600	N/A	98.9
Sensitivity [dBm]	Rx	-119	-112	-65	-102.5
Adjacent Channel Rejection [dB]		27	--	--	14
Energy/bit [nJ/bit]		43000	48000	2.5	378
Total Power [µW]		43000	48000	2500	378
Sleep Power [nW]		300	600	N/A	10.2
VDD [V]	Si	3.0	3.0	0.8|1.0|1.4	0.5|1.2
Process [nm]		--	--	180	130
Die Area [mm^2]		--	--	1.82	1.1

Table 16: Comparison with State of the Art

4.6. Conclusion

In the introduction section of the chapter a hypothetical scenario was discussed using a CC1101 in which a 100-bit packet is communicated at 1kbps. To run off harvested energy this resulted in 1.6mF of storage caps per packet and would require an average data rate of 1.04 bps to bring the overall average power consumption down to 50µW. Under the same conditions the LRTRx in receive mode will be able to receive an average 145bps and in transmit mode will be able to send 20bps. The LRTRx in Rx mode can send a packet using 28.7uF of storage cap while that increases to 120uF for Tx mode.

Increasing the average bit rate 20X for the transmitter and 145X for the receiver enables wireless sensor nodes to communicate more information over kilometer ranges. In addition, the lower instantaneous power consumption reduces the amount of storage capacitance needed which eases the design burden and allows energy harvesters to charge the node quicker. Bit-level duty cycling increases these improvements. In addition, the LRTRx demonstrates

interoperability by being designed so that it can communicate with an existing commercial TI CC1101 radio.

4.7. Acknowledgments

I would like that thank Michael C. Kines for his design of the g_mC filters and the crystal oscillator as well as his help with test automation. Also a thanks to David Moore and Elnaz Ansari for their help implementing Verilog for the digital demodulator and the duty cycle controller.

CHAPTER 5

Conclusion

The Internet of Things (IoT) revolution will result in the presence of billions of sensors all around us sensing, computing, and communicating information, even without our specified involvement. The ubiquitous deployment of these sensors will enable productivity increases that are expected to be unparalleled to date and will change the way we interact with other people and other machines. Wireless sensor nodes, living on the "edge of the cloud" will be the backbone that makes the entire IoT system work. Due to the quantity of sensor nodes to be deployed, the energy constraints placed on these nodes are significant. Current technology is not sufficient to realize billions of sensor nodes and many companies are working to meet the design challenges for these sensor nodes at the software, hardware, and system levels. One area of extreme interest is in wireless communication, which is known to be a significant hurdle for low energy design. This research looked at RF integrated circuit design in an effort to minimize both power and/or energy while maintaining the functionality of higher power radios to enable them to communicate in a crowded multi-user environment.

One of the main objectives when designing each radio was to make it robust, or capable of performing in a realistic environment and not just a lab bench. Many academic low power radio publications stress low power so much that they strip away power-intensive features that reduce their usefulness outside a lab. Robust, in this sense, then does not refer to yield or reliability, which are difficult metrics to measure with the limited time and quantity of silicon used when

testing. The "robust" radios presented here utilize several different techniques to ensure they can properly handle interference and support multiple users. Techniques include CDMA codes, automatic threshold control, and low power PLLs that can switch into different physical frequency channels.

Table 17 shows the techniques presented in this work that were used to reduce power while maintaining performance and functionality. Each will be described in more detail below.

Low Power Technique	WRX	WBAN SoC	LRTRx
Sensitivity Reduction	✓		
Subthreshold Analog Design	✓	✓	✓
Subthreshold Digital Design	✓	✓	
Asymmetric Transceiver		✓	
Standard Compatible in 1 Configuration		✓	
Low Reference PLL		✓	✓
Off-chip, High-Q Inductors		✓	✓
Low Power PSK Demodulator		✓	
Low Power FSK Demodulator			✓
Bit-Level Duty Cycling			✓

Table 17: Summary of Low Power Design Techniques

The first contribution of this research is nanowatt radio design which resulted in a nanowatt wake-up radio used for node synchronization and even low data rate communication for low energy environments. The low power WRX remains active and listening for packets while the higher power radios remain asleep to conserve power. In order to be practical the power budget of the WRX needs to be on the order of the sleep power of a sensor node, which is less than 1μW. A second contribution to this space is the development of robust radio design techniques like CDMA codes and automatic threshold control which together create a WRX that has good interference avoidance and multi-user capabilities. The WRX presented has a power of 116nW and listens for OOK-modulated CDMA packets before waking up, providing the radio with

resilience against both in-band and out-of-band interferers. In addition, multiple codes were used to wake up different nodes uniquely or to send different messages to the same WRX.

Several techniques were used to reduce the power to 116nW. First, the sensitivity design target was reduced to -40dBm because body worn communication has very short range requirements. This reduced the dependency of amplification at RF. Ultra-low power rectifiers were used to replace the LNA as the RF front-end. The first version of the WRX used an active rectifier biased in deep subthreshold, while the second rectifier was passive, while also using subthreshold techniques. The difference between the two was a tradeoff between power and speed. The digital baseband power was reduced by using subthreshold design.

The second contribution was a transceiver suited for environments that require small size, and limited energy harvesting ability and no battery. The transceiver consists of a CDMA WRX that looks for four 15-bit CDMA codes simultaneously and serves as the synchronizing receiver which allows the higher power communication receiver to remain asleep as long as possible. The main communication receiver (NBRx) is a MICS-band receiver and consumes 292µW while demodulating 802.15.6 WBAN packets at 187.5kbps in 10 different physical channels. Demodulating WBAN packets at low power bridges the power vs. performance gap between academic low power receivers that are often sub-mW and higher powered standard compliant receivers that consume 10's of mW. The ultra-wideband (UWB) transmitter operates at 4GHz and has the same 187.5kbps data rate as the NBRx. Average power for the UWB transmitter is 4.18µW with a peak output power of -28.9dBm. All measurements were taking from breakout chips because the full SoC is still being tested at UVA at this time.

Numerous techniques were used to reduce power in the transceiver. From a system perspective, the nanowatt WRX synchronization receiver saves a lot of power compared to more

traditional asynchronous communication methods. The asymmetric design allowed independent optimization of both the transmit and receive radios. Also, restricting the design to communicate over only a specific configuration of the 802.15.6 WBAN standard removed the need for significant configuration and saved power. On the transmit side, the UWB transmitter utilized a low power ring oscillator to keep the instantaneous power reasonable and the UWB architecture then naturally keeps the average power low. The narrowband receiver maintained robustness through the use of a low power PLL that was capable of receiving in 10 different frequency channels. The PLL reduced power by using a low frequency reference. Standard inductors, but with higher Q than integrated inductors, were used to reduce power in the LNA and VCO and were necessary due to the 402MHz design frequency. To demodulate 802.15.6 WBAN packets a PLL-based demodulator was designed at lower power than a traditional PSK demodulator.

The third contribution was the design of a transceiver that enabled long range communication while maintaining low overall energy, enabling the use of energy harvesting techniques. The operating frequency was in the 433 ISM band and the data rate was designed to be 1kbps in order to increase the amount of signal energy per bit. FSK demodulation was handled using digital techniques instead of a more traditional envelope detection approach which enabled the use of bit-level duty cycling. Instantaneous power is 378μW for the receiver and 2.57mW for the transmitter. At those power numbers 5 kilometer communication range is theoretically possible. Lastly, the LRTRx is designed to be compatible with a commercial radio to increase interoperability and commercial viability.

This transceiver used several of the low power techniques from the previous transceiver to keep power low. A similar low reference frequency PLL was used to provide multiple physical channels and off-chip inductors were also used, again because at 433MHz inductors are not

realizable in silicon. The digitally-assisted FSK demodulator is lower power than traditional FSK demodulation techniques and also enables the ability to use bit-level duty cycling which reduces power even further. VDD reduction on the transmitter's power amplifier significantly reduced power on the transmit side.

Appendix A

Some of this information is repeated from the motivation section of Chapter 2, but this appendix gives deeper detail into channel modeling and how it can impact short range communication.

802.15.6 Channel Models CM3 and CM4

The 802.15.6 (WBAN) channel model document [79] is a collection of experiments spanning different defined channels, including body-to-body or body-to-off body communication as well as line-of-sight (LOS) and non-line-of-sight (NLOS) variations. The purpose of these channel models is for evaluating potential physical layer proposals more than producing all-encompassing models. The ones included below describe WBAN CM3, which is body-to-body communication in the 900MHz and 2.4GHz bands for both LOS and NLOS communication. Three different experiments are described, resulting in two path loss modeling equations. Experiment A in the Channel Modeling document [80] is:

$$\boldsymbol{PL(d)[dB] = a \cdot log_{10}(d) + b + N} \tag{6}$$

where a and b are coefficients of linear fitting, d is the Tx-Rx distance in mm, and N is a normally distributed random variable with σ_n. This experiment measures path loss around a stationary subject in a hospital room. S_{21} is measured between two antennas using a vector network analyzer.

Equation B in the Channel Modeling document [81] derives the following equation:

$$\boldsymbol{PL(d)[dB] = -10log_{10}\left(P_0 e^{-m_0 d} + P_1\right) + \sigma_p n_p} \tag{7}$$

where P_0 is the average loss close to antenna, m_0 is the average decay rate in dB/cm for the surface wave traveling around the perimeter of the body, P_1 is the average attenuation of components in an indoor environment radiated away from the body and reflected back toward the receiving antenna, σ_p is the log-normal variance in dB, and n_p is a zero mean unit variance Gaussian random variable. This equation accounts for signal fading. Antennas are placed horizontally around the torso as well as vertically along it and S_{21} is measured similar to Experiment A. The test subject is standing still during the experiments.

Experiment C [82] observes subject movement, but does not develop a new equation. Test subjects are observed standing, walking, and running in place in an office environment using BPSK modulated signals at 820MHz and 2360MHz. The channel response is captured on a vector signal analyzer. Results show the most significant fading effects are due to movement and the change in distance and alignment of the antennas. Variation also increases with increased movement from the test subject. Finally, channel stability over time is observed and assigned a value, the channel variation factor, which is the ratio between the standard deviation and the RMS power of the sequence.

Many different distributions such as lognormal, normal, or Weibull, [83] fit different experimentation scenarios within a BSN. The most commonly used distribution for a static channel is lognormal, supported by both the S_{21} and RSSI data collection. Different experiments show different adjustment factors to the basic lognormal equation. The lognormal result is explained in [83] by the large number of contributing effects to the attenuation of the transmitted signal which are multiplicative, or additive in the log domain. In addition, movement was shown to increase the variability of the channel, which is an important observation for BSNs. Several

papers cite the significant impact of antenna angles as well as influences from the environment (multipath) and the size and shape of the user [84].

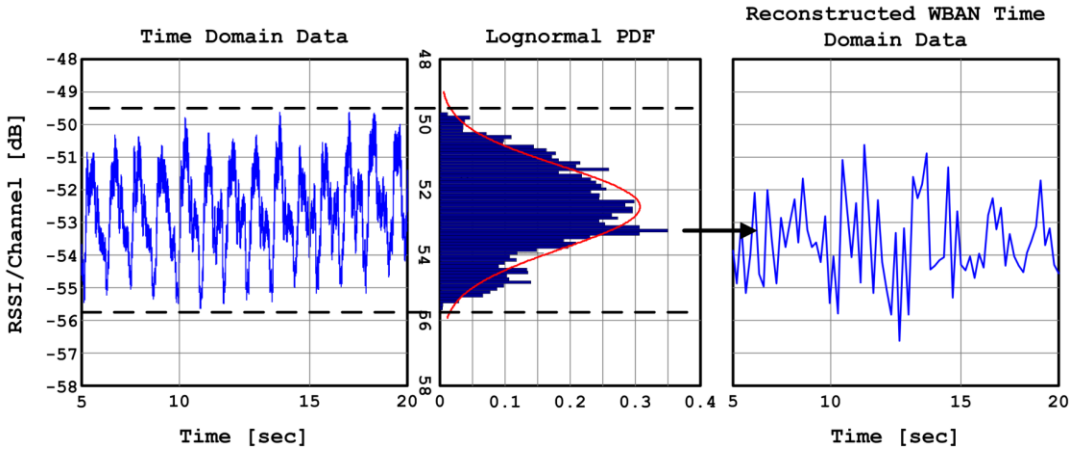

Figure 67: Channel Periodicity being Erased from WBAN Equations

Key to IoT analysis, but not included in the above equations or experiments are the channel characteristics in the time domain. Figure 67 illustrates how this periodicity is erased when using only path loss statistics, much like the previous equations do. A measured time domain channel response is shown (left) with its respective distribution (middle), which is then reconstructed using (1) from the WBAN task group channel model (right). It is obvious that (1) does not reproduce the structured periodic channel response in the original measured data, nor is that its intention. The purpose of (1) is to give a single path loss estimate for a channel. The time domain information is lost.

Dynamic Channel Modeling

In order to understand and design for the temporal characteristics of an IoT application, a simplified model that translates intuitive variables into a channel periodicity is needed. Such a model can be used to complement the already existing WBAN path loss models. Details on the experiments and hardware used in this analysis can be found in [57].

The periodic behavior of the channel seen in the data is relatively sinusoidal and can be represented as such. The proposed path loss of the periodic BSN channel model with respect to time is:

$$\boldsymbol{PL(d)[dB] = PL(d) + 2\sigma_{PL}\sin(2\pi f_{PL} t)} \quad (8)$$

where PL(d) is the path loss in dB calculated from one of the 802.15.6 [79] models (e.g. (1) or (2)), and the second term adds the modeled channel periodicity where σ_{PL} is the standard deviation, in dB, and f_{PL} is the fundamental path loss repetition frequency. This model complements the WBAN equations that calculate path loss, by adding a periodic dimension to the equation which accounts for the periodicity measured in the channel.

Numerous individual factors contribute to the path loss and periodicity of the channel including antenna placement and direction, multipath, and body shape. To integrate all these factors would produce a complicated model that would not be practical for BSN design. To generalize the results of the channel modeling measurement campaign two variables are introduced that will account for all the different factors in a simple and intuitive way. Even though it is simple by design, these factors are shown to predict the channel response with accuracy sufficient for body to body communication. The two variables are Activity Factor (AF) and Location Factor (LF).

$$Activity\ Factor\ (AF) = amount\ of\ user\ movement$$

Activity Factor is a qualitative number between 0 and 1 that is approximated by knowing what a person is doing. If a person is completely still the Activity Factor is 0 and a full sprint is 1. Activity between these extremes is qualitatively assigned an intermediate number based on a best guess as to the relative level of activity. For example, AF=0.25 is someone working at a computer, AF=0.5 is someone walking, and AF=0.75 is someone jogging. The frequency component of the path loss in (3) is then calculated as:

$$f_{PL} = f_{max} \cdot AF \quad (9)$$

where f_{max} is the maximum capable frequency of the user's movement, of which 2Hz is a good general assumption. Under normal circumstances, it is highly unlikely that an average person would exceed a fundamental repetition rate greater than 2Hz.

σ_{PL} in (3) is dependent on Activity Factor as well as a Location Factor (LF). LF is a qualitative measure of the relative motion between two sensors, normalized between 0 and 1. An LF of 1 represents sensor locations that have a lot of movement relative to each other, e.g. wrist to wrist communication. An LF of 0 represents sensor locations with little to no relative movement, e.g. two sensors on the chest. Location factor between these extremes is qualitatively assigned an intermediate number, e.g. LF=0.25 for hip to ankle, LF=0.5 for chest to wrist, and LF=0.75 is hip to wrist. σ_{PL} used in (3) is then calculated as:

$$\sigma_{PL} = k_{\sigma} \cdot LF \cdot log_{10}(1 + f_{PL}) \quad (10)$$

where k_{σ} is a data fitting parameter. Based on the controlled experiments, ranging from standing to running with sensors at various locations, $k_{\sigma} = 15$ results in the best overall fit.

If one assumes that the transmitted power and antenna gain are 0dBm and 0dB, respectively, then the sensitivity of the receiver is equivalent to the measured channel path loss. Using the periodic channel model in (3), a tradeoff between communication time and receiver power can be determined using representative receivers from literature: a low power receiver [51] and a wake up receiver [39]. Also included is a commercial Bluetooth low energy radio [59] for comparison.

Dynamic Channel Modeling Impact on Radio Design

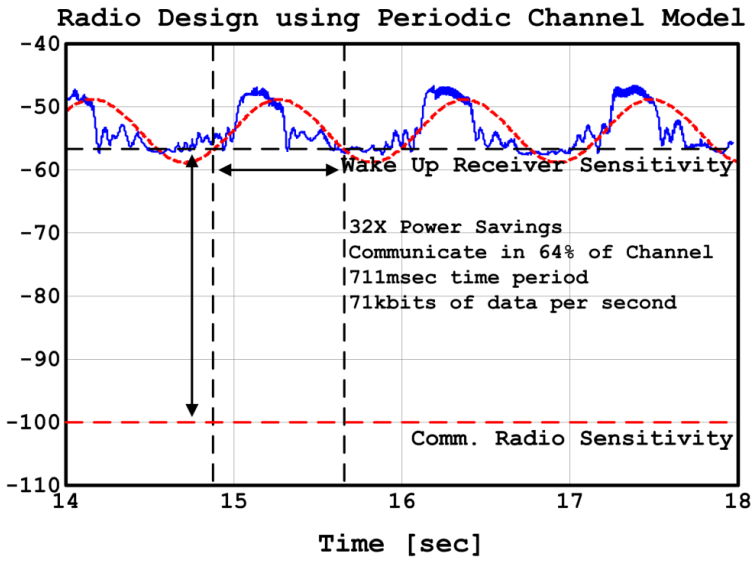

Figure 68: Radio Design with Dynamic Channel Models

Using AF, LF and PL(d) values of 0.45, 0.6, and -52dB respectively, the performance of the two low power radios can be compared. A radio [51] with a power of 2.1mW and a sensitivity of -100dBm would be the primary receiver in the BSN and always remain active. At 100% operation the total power would be 2.1mW and it can communicate 100% of the time. Using the wakeup receiver [39] with a sensitivity of -56dBm, one can turn off the main receiver to save energy while leaving the wake up receiver always listening for packets, or even utilize the wake up receiver for low speed communication. By applying the variables determined from our channel model we calculate the percentage of time the wake up receiver can communicate as:

$$T_p = \frac{\pi/2 - \arcsin\left(S_{RX} - \frac{PL(d)}{2\sigma_{PL}}\right)}{\pi} \tag{11}$$

where S_{RX} is the sensitivity of the receiver and σ_{PL} is calculated from (9) and (10), respectively. Using (11), we determine the wake up receiver will communicate 64% of the time at a power consumption of 65μW. Utilizing channel periodicity by reducing the receiver sensitivity results in a 32X power savings over the standard low power radio while being able to communicate in

64% of the channel compared to 100%. Comparing against the commercial Bluetooth radio, with a power of 58.6mW and a sensitivity of -87dBm, this results in a 900x power savings.

References:

[1] Bell, C. G., "The Mini and Micro Industries," *Computer* , vol.17, no.10, pp.14,30, Oct. 1984

[2] Bell, G., "Bell's Law for the Birth and Death of Computer Classes: A theory of the Computer's Evolution," *Solid-State Circuits Society Newsletter, IEEE* , vol.13, no.4, pp.8,19, Fall 2008

[3] Moore, G.E., "Cramming More Components Onto Integrated Circuits," Proceedings of the IEEE , vol.86, no.1, pp.82,85, Jan. 1998

[4] Cisco Visualizations, "The Internet of Things," Web http://share.cisco.com/internet-of-things.html.

[5] Murphy, Chris. "Cisco CEO: Internet of Things Will Reshape IT" Information Week. 02 Oct. 2013. Web. < http://www.informationweek.com/interop/cisco-ceo-internet-of-things-will-reshape-it/d/d-id/1111792? >.

[6] Cattle Temperature Monitoring. Bella Ag, Web. <http://www.bellaag.com/>.

[7] Endler, Michael, "Cisco CEO: We're All In on Internet of Everything," Information Week, Feb. 2013. Web. < http://www.informationweek.com/software/information-management/cisco-ceo-were-all-in-on-internet-of-everything/d/d-id/1108801?>

[8] "Nike FuelBand Teardown." *YouTube*. 31 May 2012. Web. <http://www.youtube.com/watch?v=7xdajSS_cOU>.

[9] Texas Instruments, "Bluetooth and Dual Mode Controller," CC2560 datasheet, July 2012, [Revised Oct. 2013].

[10] StickNFind. Web. March 05, 2014. <https://www.sticknfind.com>

[11] Energizer Holdings Inc, "ENERGIZER CR2016," datasheet

[12] M. Yoshio, R. J. Brodd, and A. Kozawa, eds., Lithium-Ion Batteries: Science and Technologies (Springer, 2009).

[13] Fan Zhang; et al, "A batteryless 19μW MICS/ISM-band energy harvesting body area sensor node SoC," ISSCC, 2012 IEEE, 19-23 Feb. 2012

[14] Pandey, J.; et al, "A 120μW MICS/ISM-band FSK receiver with a 44μW low-power mode based on injection-locking and 9x frequency multiplication," ISSCC, 2011 IEEE, 20-24 Feb. 2011

[15] Bluetooth SIG, "Specification of the Bluetooth System", Bluetooth 4.0 standard, Volume 0, 30 June 2010

[16] IEEE Standard for Local and metropolitan area networks--Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)," IEEE Standard for Local and metropolitan area networks--Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)

[17] IEEE Standard for Local and metropolitan area networks--Part 15.6: Wireless Body Area Networks

[18] Nordic Semiconductor, "Single-chip Bluetooth Low Energy Solution," nRF8001 datasheet, Jan. 2012 [Revised Aug. 2013].

[19] Peiris, V.; et. al, "A 1 V 433/868 MHz 25 kb/s-FSK 2 kb/s-OOK RF transceiver SoC in standard digital 0.18 μm CMOS," *ISSCC*. Feb. 2005

[20] Otis, B.; et. al, "A 400 μW-RX, 1.6mW-TX super-regenerative transceiver for wireless sensor networks," *ISSCC*. Feb. 2005

[21] Pletcher, N.M., et. al, "A 2GHz 52 μW Wake-Up Receiver with -72dBm Sensitivity Using Uncertain-IF Architecture," *ISSCC*. Feb. 2008

[22] Drago, S.; et. al, "A 2.4GHz 830pJ/bit duty-cycled wake-up receiver with −82dBm sensitivity for crystal-less wireless sensor nodes," *ISSCC,* Feb. 2010

[23] Xiongchuan Huang; et. al, "A 2.4GHz/915MHz 51μW wake-up receiver with offset and noise suppression," *ISSCC,* Feb. 2010

[24] Vidojkovic, M.; et al, "A 2.4GHz ULP OOK single-chip transceiver for healthcare applications," *ISSCC,* Feb. 2011

[25] Jeongki Choi; et. al, "An interference-aware 5.8GHz wake-up radio for ETCS," *ISSCC,* Feb. 2012

[26] Papotto, G.; et. al, "A 90nm CMOS 5Mb/s crystal-less RF transceiver for RF-powered WSN nodes," *ISSCC,* Feb. 2012

[27] Huang, X.; et. al, "A 915MHz 120μW-RX/900μW-TX envelope-detection transceiver with 20dB in-band interference tolerance," *ISSCC,* Feb. 2012

[28] Hancock, T.M.; et. al, "A Sub-10mW 2Mbps BFSK Transceiver at 1.35 to 1.75GHz," *RFIC*, June 2007

[29] Bernier, C.; et. al, "An ultra low power 130nm CMOS direct conversion transceiver for IEEE802.15.4," *RFIC,* June 2008

[30] Nezhad-Ahmadi, et. al, "A 2mW 400MHz RF transceiver SoC in 0.18um CMOS technology for wireless medical applications," *RFIC, 2008.*

[31] van Langevelde, et. al, "An ultra-low-power 868/915 MHz RF transceiver for wireless sensor network applications," *RFIC, 2009*

[32] Yan, H.; et. al, "A 120μW fully-integrated BPSK receiver in 90nm CMOS," *RFIC,* May 2010

[33] Yong-Il Kwon; et. al, "A fully integrated 2.4-GHz CMOS diversity receiver with a novel antenna selection," *RFIC,* May 2010

[34] Seungkee Min; et. al, "A 2mW CMOS MICS-band BFSK transceiver with reconfigurable antenna interface," *RFIC,* May 2010

[35] Brown, J.K.; Wentzloff, D.D., "A 1900MHz-band GSM-based clock-harvesting receiver with −87dBm sensitivity," *RFIC,* June 2011

[36] Rong-Fu Ye; et. al, "Highly sensitive and low power injection-locked FSK receiver for short-range wireless applications," *RFIC,* June 2012

[37] Nadeau, P.M.; et. al, "Multi-channel 180pJ/b 2.4GHz FBAR-based receiver," *RFIC,* June 2012

[38] Jia-Yi Chen; et. al, "A 3.6mW 2.4-GHz multi-channel super-regenerative receiver in 130nm CMOS," *CICC,* Sept. 2005

[39] Pletcher, N.; et. al, "A 65μW, 1.9 GHz RF to digital baseband wakeup receiver for wireless sensor nodes," *CICC,* Sept. 2007

[40] Ayers, J.; et. al, "A 0.4 nJ/b 900MHz CMOS BFSK super-regenerative receiver," *CICC,* Sept. 2008

[41] Jae Hyuk Jang; et. al, "A wireless sensor node for condition monitoring powered by a vibration energy harvester," *CICC,* Sept. 2011

[42] Ishizaki, H.; et. al, "A 2.4GHz ISM-band digital CMOS wireless transceiver with an intra-symbol adaptively intermittent Rx," *VLSI Symposium*, June 2007

[43] Bohorquez, J.L.; et. al, "A 350 μW CMOS MSK Transmitter and 400μ W OOK Super-Regenerative Receiver for Medical Implant Communications," *JSSC*, April 2009

[44] Dan Shi; et. al, "A 5GHz fully integrated super-regenerative receiver with on-chip slot antenna in 0.13μm CMOS," *VLSI Symposium*, June 2008

[45] Gambini, S.; et. al, "A fully integrated, 300pJ/bit, dual mode wireless transceiver for cm-range interconnects," *VLSI Symposium*, June 2010

[46] Ayers, J.; et. al, "A 2.4GHz wireless transceiver with 0.95nJ/b link energy for multi-hop battery-freewireless sensor networks," *VLSI Symposium*, June 2010

[47] Ronghua Ni; et. al, "A 2.4GHz hybrid PPF based BFSK receiver with ±180ppm frequency offset tolerance for wireless sensor networks," *VLSI Symposium*, June 2012

[48] Saito, A.; et. al, "An all 0.5V, 1Mbps, 315MHz OOK transceiver with 38-μW career-frequency-free intermittent sampling receiver and 52-μW class-F transmitter in 40-nm CMOS," *VLSI Symposium*, June 2012

[49] Sayilir, S.; et. al, "A −90dBm sensitivity wireless transceiver using VCO-PA-LNA-switch-modulator co-design for low power insect-based wireless sensor networks," *VLSI Symposium*, June 2013

[50] Daly, D.C.; Chandrakasan, A.P., "An Energy-Efficient OOK Transceiver for Wireless Sensor Networks," *JSSC*, May 2007

[51] Wong, A. C W; et. al, "A 1V wireless transceiver for an ultra low power SoC for biotelemetry applications," *ESSCIRC,* 2007

[52] Ayers, J.; et. al, "An Ultralow-Power Receiver for Wireless Sensor Networks," *JSSC*, Sept. 2010

[53] D.C. Daly and A.P. Chandrakasan, "An Energy-Efficient OOK Transceiver for Wireless Sensor Networks," *IEEE J. Solid-State Circuits*, vol.42, no.5, pp.1003-1011, May 2007

[54] JaeHyuk Jang; et. al, "A Wireless Condition Monitoring System Powered by a Vibration Energy Harvester," *TCAS*, April 2013

[55] Yoonmyung Lee, et al, "A 660pW multi-stage temperature-compensated timer for ultra-low-power wireless sensor node synchronization," *ISSCC 2011*, 20-24 Feb. 2011

[56] B. H. Calhoun, et. al, "Body Sensor Networks: A Holistic Approach from Silicon to Users," *Proceedings of the IEEE*, Vol. 100, No. 1, 2012, pp. 91-106

[57] N.E. Roberts, et. al, "Exploiting Channel Periodicity in Body Sensor Networks," *IEEE Emerging and Selected Topics in Circuits and Systems*, vol.2, no.1, pp.4-13, March 2012

[58] P.S. Hall, et. al, "Antennas and Propagation for On-Body Communication Systems," *IEEE Antennas and Propagation Magazine*, June 2007

[59] Texas Instruments, "2.4GHz Bluetooth low energy System-on-Chip," CC2540 datasheet, Oct. 2010, [Revised June 2013]

[60] F. Assaderaghi, et. al, "Dynamic threshold-voltage MOSFET (DTMOS) for ultra-low voltage VLSI," *IEEE Trans. Electron Devices*, vol.44, no.3, pp.414-422, Mar 1997

[61] S. Miller and L. MacEachern, "A nanowatt bandgap voltage reference for ultra-low power applications," *IEEE Int. Sym. Circuits and Systems*, pp.645-648, May 2006

[62] Q. Xinbo and T.H. Teo, "A low-power comparator with programmable hysteresis level for blood pressure peak detection," *TENCON*, pp.1-4, Jan. 2009

[63] J.F. Dickson, "On-chip high-voltage generation in MNOS integrated circuits using an improved voltage multiplier technique," *Solid-State Circuits, IEEE Journal of* , vol.11, no.3, pp. 374- 378, Jun 1976

[64] Seunghyun Oh and D.D. Wentzloff, "A −32dBm Sensitivity RF Power Harvester in 130nm CMOS," *RFIC 2012,* June 2012

[65] J.K. Holmes, "Spread Spectrum Systems for GNSS and Wireless Communications, " Artech House, ISBN 978-1-59693-093-4, 2007

[66] Oporta, Hector, "An Ultra-low Power Frequency Reference for Timekeeping Applications," *Master's Thesis*, Oregon State University, December, 2008.

[67] N.E. Roberts, and D.D. Wentzloff, "A 98nW Wake-Up Radio for Wireless Body Area Networks," *RFIC 2012,* June 2012

[68] Seunghyun Oh; et. al, "A 116nW multi-band wake-up receiver with 31-bit correlator and interference rejection," *CICC,* Sept. 2013

[69] Lacaita, Andrea, et. al, Integrated frequency Synthesizers for Wireless Systems. Cambridge University Press, 2007. Print.

[70] Song, Taeksang, et al., "A Low-Power 2.4-GHz Current-Reused Receiver Front-End and Frequency Source for Wireless Sensor Network," JSSC, vol.42, no.5, May 2007

[71] Roundy, et al., “Power sources for wireless networks,” Proceedings of 1st European Workshop on Wireless Sensor Networks, 2004

[72] Analog Devices, “High Performance, ISM Band, FSK/ASK Transceiver IC,” ADF7020 datasheet, June 2005, [Revised Aug. 2012]

[73] Texas Instruments, “Low Power Sub-1GHz RF Transceiver,” CC1101 datasheet, April 2007, [Revised Nov. 2013]

[74] Semtech, “137MHz to 1020MHz Low Power Long Range Transceiver”, SX1276/77/78 datasheet, Sept. 2013.

[75] Analog Devices, “Low Power, Long Range, ISM Wireless Measuring Node,” CN-0164 datasheet, Oct. 2010, [Revised Feb. 2011]

[76] Shannon, Claude E., and Warren Weaver. The mathematical theory of communication. Urbana: University of Illinois Press, 1964. Print.

[77] University of Michigan, Electrical Engineering & Computer Science 455 Class Notes

[78] Texas Instruments, “Low Power Sub-1GHz RF Transceiver,” CC1101 datasheet, April 2007, [Revised Nov. 2013]

[79] “Channel modeling for body area networks,” IEEE P802.15-08-0780-12-0006 Nov. 2010.

[80] “Channel models for WBANs-NICT,” IEEE P802.15-08-0416-04-0006 Nov. 2008.

[81] “Channel models for WBAN-holst centre/IMEC-NL,” IEEE P802.15-08-0418-01-0006 Jul. 2008.

[82] “Narrowband channel characterization for body area networks,” IEEE P802. 15-08-0421-00-0006 Jul. 2008.

[83] D. Smith et al., “Characterization of the dynamic narrowband on-body to off-body area channel,” in IEEE Int. Conf. Commun., Jun. 2009, pp. 1–6.

[84] Y. Zhao et al., “A simulation environment for subject-specific radio channel modeling in wireless body sensor networks,” Wearable Implantable Body Sensor Netw., pp. 23–28, Jun. 2009.

[85] Roberts, N.E.; et al, "Exploiting Channel Periodicity in Body Sensor Networks," Emerging and Selected Topics in Circuits and Systems, IEEE Journal, March 2012