

Design of a system for cm-range wireless communication

*Simone Gambini
Jan M. Rabaey
Elad Alon*



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2009-184

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-184.html>

December 18, 2009

Copyright © 2009, by the author(s).
All rights reserved.

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

Design of a system for cm-range wireless communications

by

Simone Gambini

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

Doctor of Philosophy

in

Electrical Engineering and Computer Sciences

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:

Professor Jan Rabey, Chair

Professor Elad Alon

Professor Paul K. Wright

Fall 2009

The dissertation of Simone Gambini is approved.

Chair

Date

Date

Date

University of California, Berkeley
Fall 2009

Design of a system for cm-range wireless communications

Copyright © 2009

by

Simone Gambini

Abstract

Design of a system for cm-range wireless communications

by

Simone Gambini

Doctor of Philosophy in Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Jan Rabey, Chair

The continuous growth in the number of mobile phone subscribers, which exceeded 3 billions by 2007, and of the number of wireless devices and systems, led to visions of a near future in which wireless technology is so ubiquitous that *1000 Radios* per person will exist. In this context, ad-hoc area networks between several consumer electronic devices located within a meter of each other will be necessary in order to reduce traffic towards the base station and reduce power consumption and interference generation. As existing air interfaces still both radiate and dissipate an order of magnitude higher power than what this future scenario requires, a new, low-power radio technology must be developed. In this thesis, we develop a low-power transceiver with range of a few centimeters targeted to mobile-mobile data exchange. The same transceiver could also be employed in some implanted applications, as well as in distributed industrial control environments.

The design of the air-interface for this cm-range communication system at the propagation, system and circuit levels. First, we describe an optimization methodology that enables the designer to choose, for any given antenna design, which carrier frequency results in the maximum receiver Signal-To-Noise-Ratio (SNR). We then show how by using impulse-radio signaling, the chosen high-SNR channel can be leveraged to simplify the radio receiver architecture to the simplest possible RF receiver-consisting only of an RF rectifier. To mitigate the known issues of sensitivity to interference for these rectifier-based receivers, a technique to improve selectivity that uses only baseband processing and requires no RF prefilter is introduced.

These techniques are demonstrated by a transceiver test chip, implemented in a 65nm CMOS process. The transceiver dissipates $250\mu W$ in receive mode, and $25\mu W$ in transmit mode when operating at 1Mbps, and it integrates a timing-recovery loop that achieves jitter lower than 2nS while consuming $45\mu W$. These figures correspond to

an energy per bit of 300pJ, which compares favorably with current state-of-the art.

Professor Jan Rabey
Dissertation Committee Chair

To my mother Silvia, my brother Francesco and my father Diego.

If I have seen a little further it is by standing on the shoulders of Giants
I. Newton

Contents

List of Figures	vi
List of Tables	xii
Acknowledgements	xiii
1 1000 Radios per Person And the Case for Ad-Hoc Networking	1
1.1 Short range wireless systems today	4
1.1.1 Commercially available, low-power <i>short range</i> bi-directional communication devices	4
1.1.2 Bi-directional,narrowband short range radios resulting from academic papers	5
1.1.3 Impulse Ultra-wideband radios resulting from academic papers	6
1.1.4 Other short-range, low energy communication systems	7
1.2 Thesis structure	8
2 Link Margin Modeling for cm-range radios	9
2.1 Improving path loss modeling	10
2.1.1 Electromagnetic wave propagation at short range	10
2.1.2 Magnetic Coupling	15
2.1.3 Model Verification	16
2.1.4 Comparison with power transfer	23
2.1.5 Chosen design point and channel measurements with cm-scale antennas	24
2.1.6 Channel impulse response	24
2.2 Antenna Design for a 3-10GHz wideband system	25
2.2.1 Slotted Bowtie Antenna Design	27

2.2.2	Monopole antenna design	30
2.3	Conclusions	31
3	Low Power RF System Design for a cm-range Transceiver	34
3.1	Radio Dynamic Range-Power Tradeoff	34
3.1.1	Overhead power components	36
3.1.2	Summary	38
3.2	Effect of duty-cycling on power dissipation	39
3.2.1	Overhead Costs in duty-cycled systems	41
3.2.2	Duty cycling optimization	41
3.2.3	Inherent Asymmetry of duty-cycled systems	44
3.2.4	Power consumption of VCO and baseband: a more critical look . . .	44
3.3	Radio Architecture Design	45
3.3.1	Technological Constraints	46
3.3.2	Transmitter section and link budget	46
3.3.3	Receiver section	47
3.4	Conclusions	51
4	Baseband and Timer Subsystem Design	53
4.1	Baseband architectures for UWB systems	53
4.2	Improving Interferer Performance in Self-Mixing Radios	55
4.3	Synchronizer and modulation scheme design	66
4.4	Summary	76
5	Circuit Design	78
5.1	Introduction	78
5.2	Transmitter Circuits	79
5.2.1	Architecture selection	79
5.2.2	Oscillator	81
5.2.3	Divider	83
5.2.4	PSK Modulator and Driver	85
5.3	Receiver Circuits	88
5.3.1	Matching network/duplexer	88

5.3.2	Reconfigurable mixer/diode	90
5.3.3	Baseband Filter	97
5.3.4	Comparators	102
5.3.5	Layout and summary	107
5.3.6	Bias Circuits and power gating strategy	107
5.4	Timing Circuits	110
5.4.1	Baseband oscillator	112
5.4.2	Baseband divider and comparator	113
5.4.3	Pulse detection and retiming logic	114
5.4.4	Controller	115
5.4.5	Block Layout and summary	116
5.5	System Layout and summary	116
6	Measurement Results	121
6.1	FIB Issue	121
6.2	Transmitter measurements	122
6.3	Receiver Measurements	126
6.3.1	Sensitivity Measurements	126
6.3.2	Duty-cycled Measurements	128
6.3.3	Interference Rejection Measurements	128
6.3.4	Wireless Tests	131
6.3.5	Receiver Performance Summary	135
6.3.6	Comparison with prior art	135
6.4	Timer Measurements	137
7	Conclusions	142
7.1	Research Directions	143
7.1.1	Power-efficient TCM for wireless sensor networks	143
7.1.2	Circuit Implementation	143
A	Matching Network Impedance Calculations	144
A.1	Receiver Side Impedance	144
A.2	Transmit Side Impedance	146

List of Figures

1.1	Pushpin nodes, a first implementation of the paintable computing concept [66]. Approximate node size is $3.24cm^2$	4
2.1	Schematic of electrical dipoles and frame of reference	10
2.2	Equivalent circuit for monopoles in the near-field operating range	13
2.3	Measured, calculated and simulated return loss for wire monopoles	17
2.4	Calculated and measured path loss as a function of frequency	17
2.5	Measured and simulated path loss as a function of frequency	18
2.6	Measured path loss (including matching network) for the TX-RX Antenna pair at different separations	18
2.7	Experimental setup in the BWRC Lab	19
2.8	Calculated(top) and measured(bottom) loss as a function of frequency for different antennas separations and quality factors	20
2.9	Effect of matching network quality factor on voltage gain	21
2.10	Measured and simulated path loss as a function of range for different frequencies	21
2.11	Antenna impedance quality factor as a function of frequency	22
2.12	Measured Channel S_{21} as a function of TX-RX separation	25
2.13	Received energy as a function of tap number	26
2.14	Channel Impulse Response Sampled at 2GS/s for good channel: I-component (top) and Q component (bottom)	26
2.15	Annotated Bowtie Antenna Schematic	28
2.16	Bowtie Antenna Simulations: Effect of varying the expansion angle (top); effect of varying the inductive stub length(bottom)	29
2.17	Bowtie Antenna Measured Reflection versus HFSS Simulations	30
2.18	Monopole antennas annotated schematic:full monopole(top) and loaded hollow monopole (bottom)	32

2.19	Measured results from monopole antennas (Full monopole (top) and loaded hollow monopole (bottom)	33
3.1	Optimal Radio Range as a function of total radio power budget	35
3.2	Maximum Radio Range as a function of carrier frequency	36
3.3	Power dissipation versus total baseband gain	37
3.4	Transceiver duty cycling concept	39
3.5	Analog Synchronizer	42
3.6	Power Dissipation of TX+RX as a function of duty-cycling	42
3.7	Minimum Power Dissipation as a function of receiver turn-on-time T_m	43
3.8	Desensitization due to reciprocal mixing	45
3.9	Comparison between a receiver comprising of LNA+AM detection (left), and lo power oscillator+passive mixer(right)	48
3.10	Low Power Radio Architectures: a) Heterodyning b) RF Gain plus AM-Detection c) Direct AM Detection	49
3.11	Behavioral model of AM detector	49
3.12	Operation of a single transistor MOS AM Detector	50
3.13	Minimum Input Amplitude s.t. AM-Detector is more power efficient solution	52
4.1	Receiver architecture with linear downconversion and analog correlation . . .	54
4.2	Conversion gain for receivers with different baseband integration times . . .	55
4.3	AM Detection as phasor domain length-measurement	56
4.4	False Positives Rate for system employing high pass filter versus conventional case	57
4.5	Graphical Interpretation of the worst case angle	59
4.6	Complete Baseband architecture for false-negative rate calculation	60
4.7	Bit-Error Rate as a function of power for different frequency offset values . .	60
4.8	Bit-Error Rate as a function of interferer offset of SIR=6dB	62
4.9	Simulated error of the receiver in diode mode for pulse durations of 2nS (blue) and 8nS (red). SIR=6dB	62
4.10	Receiver with linear down converter and high-pass filtering at baseband . . .	64
4.11	Bit-error rate performance in a linear downconverter as a function of interferer frequency offset	64

4.12	Proposed dual-mode receiver with real-time quality of service control(top) and overlaid false negative rates response of linear and AM-detecting down-converters(bottom)	65
4.13	Conceptual Synchronizer Diagram	67
4.14	Block Diagram of Modem employing differential-PPM modulation and merged synchronizer/demodulator	68
4.15	Trellis diagram for differentially demodulated PPM	69
4.16	Effect of jitter on digital timing measurement	70
4.17	Period Estimate in the synchronizer with (red) and without (blue) data-dependent rescaling: input to the filter (left) and period estimate(right) . . .	72
4.18	Waveforms in the synchronizer in the event of a missing pulse	73
4.19	Probability density function(top) and cumulative distribution of input transition	74
4.20	Simulated BER Estimate versus real BER	76
4.21	Comparison of waterfall curves with conventional and proposed amplitude estimation algorithm	77
5.1	Block diagram of transceiver architecture	79
5.2	Bandpass Pulse Generation Architectures: a) filtering; b) Edge Combining; c) Gated Oscillator	80
5.3	Digitally Controlled Oscillator Schematic	82
5.4	Simulated tuning range	82
5.5	Impedance network model	83
5.6	Tuning control shown for the bottom rail. Top rail is symmetric but uses PMOS transistor with 2x width	84
5.7	Johnson divider architecture	84
5.8	TSPC Flip Flop circuit details and transistor sizes	85
5.9	PSK Modulator/multiplexer and control circuits	86
5.10	Simulated Pulser Waveforms	86
5.11	Effect of delayed inputs on glitching at the output of the multiplexer	87
5.12	High-level matching network view	89
5.13	Matching Network Used in The Design	91
5.14	Simulated input matching network for different values of digitally controlled capacitor C_2	92
5.15	Simulated impedance at the transmitter port	92

5.16	Envelope Detector topology of [52]	93
5.17	Output SNR and Amplitude level comparison for CS and CD diodes	94
5.18	Common-source diode	95
5.19	Implemented Mixer-Diode stage	96
5.20	CAD Layout of RF Front-end	97
5.21	Circuit diagram for lane of interleaved baseband	98
5.22	Loads with high common-mode compliance: diode (left); and diode with level shift(right)	99
5.23	Distributed filter schematic (1 lane)	100
5.24	Dac Unit Cell Schematic. Nominal LSB current is $1.5\mu A$	100
5.25	INL and DNL standard deviation for one the FIR filter bias-DACs	101
5.26	Track and Hold Amplifier	102
5.27	Isolation Amplifier Schematic	103
5.28	Simulated FIR Filter frequency response for different settings	104
5.29	Block diagram of slicing chain	104
5.30	Annotated Comparator Schematic	105
5.31	Comparator simulations: decision threshold versus digital code(left) and uncalibrated input-referred offset distribution (right)	106
5.32	Example of the effect of dual-thesholding and equivalence between dual thresholding and absolute value thresholding: the blue waveform, corresponding to $\delta F = 100MHz$ and initial phase of π , never crosses the positive threshold, and goes undetected by single-comparator scheme.	106
5.33	Schematic of bias circuit for RF front-end	108
5.34	Simulation of bias settling:mixer differential output (left) and bias lines shown in Fig.5.33	109
5.35	CAD Layout of analog baseband	109
5.36	Power-on sequence	110
5.37	Synchronizer Architecture and Synchronous-Asynchronous Partition	111
5.38	Baseband DCO and driver schematic	113
5.39	Baseband divider architecture	114
5.40	Counter Reset Logic	115
5.41	Binary comparator cell	116
5.42	Asynchronous interface building blocks	117

5.43	Pulse Detection Logic	117
5.44	CAD Layout of timer	118
5.45	CAD Layout of transceiver	119
6.1	Die Photo	122
6.2	Location of missing connection	123
6.3	Measured input reflection varying digital tuning of matching network capacitor C_2	124
6.4	Measured Transmitter Output Waveform and Spectrum	125
6.5	Fibbed Transmitter Power dissipation (top) and energy/pulse (bottom) . . .	126
6.6	RF DCO Tuning Curve	127
6.7	Measured received waterfall curve in wired test	128
6.8	Measured power dissipation and error rate as a function of duty-cycling window width	129
6.9	Measured false alarm rate versus interferer carrier frequency in diode mode for two different values of desired signal carrier	130
6.10	Measured missed-detection rate as a function of Signal-to-Interferer Ratio for different offset frequencies	130
6.11	Required SIR versus offset frequency to achieve 5e-3 B.E.R. for diode and mixer mode	131
6.12	Board picture and Skycross and custom monopole antennas	132
6.13	Test Setup	133
6.14	Required SIR versus offset frequency to achieve 5e-3 B.E.R. for diode and mixer mode	134
6.15	Wireless Link Measurement with TX-RX pair	134
6.16	Alternate Bias Scheme	136
6.17	Measured Timer Power Dissipation versus timing resolution at different supply voltages	139
6.18	Transmit mode cycle-to-cycle jitter	140
6.19	Measured distribution of window opening to pulse arrival delays	140
6.20	Demonstration of error-detection functionality	141
A.1	Initial L-matching network	145
B.1	Probability density function of time measurement for different values of input jitter variance σ_{in} to quantization resolution T_s	148

List of Tables

1.1	Commercial Radio Performance Summary	5
1.2	Table 1.2: Recently published narrow-band low power radios	6
1.3	Table 1.3: Recently published impulse UWB transceivers	7
2.1	Operating regions in the size-separation-wavelength plane	12
2.2	Comparison of calculated and measured power and voltage transfer parameters	24
3.1	Receiver Link Margin Calculations	47
5.1	Power consumption breakdown in pulse generator	88
5.2	Mixer-Diode Performance	95
5.3	Baseband circuits performance breakdown	107
5.4	Baseband circuits power consumption breakdown	107
5.5	Synchronizer power consumption breakdown	116
5.6	Synchronizer performance summary	118
5.7	Transceiver performance summary	120
6.1	Summary of measured transceiver performance	136
6.2	Transceiver Performance Comparison with 0-1GHz system [11]	137
6.3	Performance Comparison with published UWB Receivers	138
6.4	Transmitter Performance Comparison with published UWB Receivers	138

Acknowledgements

I would not have made it through these last five years at Berkeley if it was not for the many people that, at different times, worked with me, helped me, supported me or came out drinking with me. First of all, I would not have even started without the support of my mother, my brother, my aunt Lucilla, my grandmothers Clara and Edra and the rest of the family in Italy , who always encouraged me to follow my ambitions, irrespective of how much they missed me or how difficult their life was at the time.

Once arrived here, without the input of many fellow graduate students who shared part of this trip of with me, this thesis would not have existed. I was lucky in having B.Otis and J.Vanderhaegen as first mentors and ,over the years, I had the opportunity to work with many other exceptionally talented individuals. Many discussions with Davide Guermandi, Mubaraq Mishra,Louis Alarcon,David Sobel, Amin Arbabian, Debopriyo Chowdury, Ehsan Adabi,John Crossley, and Renaldi Winoto contributed to shaping my research choices directions. A special thanks to Davide for all his mentoring during his year at Berkeley, and to John for going out of his way to help me verify and tape out the chip described in Chapter 4 .

A large percentage of my success, both in the academic and the sports arenas at Berkeley, has been shared with Nathan Pletcher. Being a physicist, Nate was equipped with just the right amount of common sense and detail-orientation to keep our joint projects from derailing on some of my in-practical ideas. Were this not enough,he also scored the decisive goal in *Atletico Manero*'s glorious intramural soccer C-league victory. Thanks.

Many of the ideas on interference rejection described in this thesis matured during an internship at Intel Research in Hillsboro through my discussions with the members of Radio Integration Laboratory. If such title existed, Yorgos Palaskas, Stefano Pellerano , Hasnain Lakdawala and Ashoke Ravi should all be honorary members of this dissertation committee .

My efforts towards this degree would similarly have been vane without the contribution of several Berkeley faculty and staff members. I would like to thank my advisor, Prof. Jan Rabaey, for promoting , shaping with his vision and funding the research described in this thesis , as well as for founding BWRC. I am also very grateful to Professor Elad Alon for our weekly meetings and his un-quenchable thirst for understanding every engineering problem, and then argue about it until fully convinced, forcing me to achieve his same level of understanding. A great thanks also to Prof.Bernhard Boser and Paul Wright for serving on my qualifying exam and (respectively) M.S. and Ph.D. thesis committees and to Prof. Alberto Sangiovanni Vincentelli 's for the many times that I took advantage of his remarkable wisdom and research vision.

I would also like to thank Tom Boot, Fred Burghart, Brian Richards, Kevin Zimmerman and Susan Mellers for making my job as a researcher so much easier through their daily support work; and my " local Mom" Ruth Gjerde for looking over my shoulder in all the departemental matters.

My life at Berkeley would have been unbearable without the friends I met outside of lab.Thanks to Pietro Toscano, Devis Contarato and Lorenzo " Ataturk" Orecchia I always had someone to talk and explore San Francisco with. Thanks also to Costas and Maria-Daphne

for making me an honorary member of the Berkeley (and now Boston) Greek community early on, and to Max Fratoni, Alessandro Abate, Pinto and the rest of *I Vitelloni* and *Atletico Manero* for keeping me into the Italian one as well. Thanks to my roommates Carmelo, Thøersten and Maura for putting up with my disorder on a daily basis.

Thanks also to my friends in Italy, Stefano, Lorenzo, Federico Luca, Marco, Sara and the rest of the Rain Day, have done their best in making me as if nothing had changed every time I went home.

Last, but certainly not least, I would like to thank my girlfriend Rikky (whom to be fair, I should have also named amongst the brilliant colleagues), who has been filling with love almost every hour of my life for more than a year now. Having met her is by itself a good enough payoff for the sacrifices of the last five years. I hope this is nothing but the beginning of our story together.

Chapter 1

1000 Radios per Person And the Case for Ad-Hoc Networking

The last ten years have witnessed a tremendous growth in wireless systems. After saturating industrialized countries, cell phone penetration has become prominent in the developing world, increasing the number of subscribers to 4 billions by 2009. Wireless LAN cards are now provided in virtually every one of the over 100 millions laptops sold each year worldwide and are also included in several smart-phones. Around 2002, researchers ([22]) identified the next trend in wireless to be one in which many chips assembled in ultra-small scale systems and equipped with wireless communication capability are distributed in the environment, providing several sensing and actuation functions without human intervention. The first examples of wireless sensor networks have already seen, or will soon see commercialization in the fields of automatic meter monitoring and industrial control . Beyond these original wireless sensor network deployments, growth in low-power and small size distributed wireless systems could be fueled by *ultra-short-range* radios introduced to replace connections that are now wired, or established directly by the user. Indeed, as technology scaling made transistor switching speed higher and an increasingly large fraction of RF chips are realized in low-cost CMOS , RF modules have shrunk their size and power consumption so that it has become possible to integrate several radios in a single consumer electronic device . In the near future, the percentage of electronic appliances including wireless connectivity is likely to keep increasing at a fast pace, providing users of more and more markets with greater ease of use and mobility performance. In this *1000 Radios per person* scenario, ad hoc, peer to peer wireless interconnects between consumer electronic devices will gain increasing importance both as a mean of extending battery life, as well as to mitigate interference. Energy efficient wireless links with somewhat different specifications, but sharing a sub-1m range could be employed by any of the following applications:

- Handheld-laptop synchronization.
A high-data rate (10Mbps-1Gbps),cm-range point-to-point/ star wireless link could be used to avoid cables in docking stations, or mobile device synch operation,leading to greater ease of use. In synchronization between a laptop and an handheld device it would be desirable to transmit large files, leading to long packets and high data rate requirements, even at the expense of power. However, in the scenario of data transfer between handhelds , likely packet lengths shrink and power dissipation takes priority over data rate.
- Collaboration channel for multi-terminal/collaborative MIMO.
It has been shown in ([45]) that the total capacity of a set of nodes scattered in a finite area can scale linearly with the number of nodes (implying no interference limitation), if the nodes exchange co-operation information and synchronize their transmissions. Introducing a separate short-range coordination channel could reduce the power overhead associated with the coordination process and ultimately lead to higher capacity. This same reasoning can also be applied to the cooperation of an ensemble of sensor nodes that coordinates before relaying data to a base-station with known location using distributed beamforming.
- Distributed industrial control.
Wireless links can be used to replace wires in an industrial automation scenario, avoiding mechanical wear, simplifying assembly and reducing cost. This application demands point-to-multipoint links with moderate data rate (1-10Mbps), but capable of almost error-free operation
- Wireless packaging. Wireless board to board/ module to module interconnects in tightly integrated microsystems could make expensive connectors unnecessary, decreasing system cost and size. Point- to-point links with reliable operation and low-channel to channel interference are required . Data rates of these links are highly dependent on the application. For example, replacing the video connector in a laptop requires an interface with a throughput exceeding 1Gbps; while simple business card/contact info exchange between mobile devices could be accomplished with rates of the order of 1Mbps.
- Body-area networks/ implanted sensor to reader data transmission: the possibility of realizing connections between sensor nodes distributed in various part of a human body and a base station situated in a neighborhood of the body itself, such as a cell phone or a dedicated super-node[17] open exciting possibilities, such as real-time, free-behaving patient monitoring and seamless drug administration. While these systems are typically envisioned to operate with a star connection, i.e. with the sensor nodes collecting vital sign data and relaying them to the base station, actuation or control packets can also be exchanged, requiring a low energy downlink as well. This application demands sub-meter radio range, although in a difficult propagation environment such as the the human body [16] . Low power,Mb/s links could also be used to communicate in a variety of electronic implants, such as cochlear implants [33] , brain-machine interface circuits [43] or retinal prosthetics [63]

These systems are characterized not only by short range, but also by a prevailing point-to-point or multi-point to point connectivity, operation in a environment rich of electromagnetic scatterers and sources of electromagnetic interference, and the requirement for a data rate of at least 1 Mbps. It is also fair to say that across the board, systems with extremely small form factor are mandatory to facilitate unobtrusive integration in different devices. Due to the different specifications, a single design cannot efficiently satisfy the demands of all these applications. As a result , the focus of this thesis is to realize a solution that is tailored at a specific use case, while developing a design methodology that could be extended to others. The application targeted in this work is that of paintable computing [66]. In this scenario, several lightweight computing devices are *sprayed* on a surface and equipped with a short range radio allowing them to communicate with their close neighbors and run distributed algorithms. For the purpose of this work, we assume that these radios are distributed in such a way that the distance between nearest neighbors is not larger than 5 cm and that the nodes already have a power source available (i.e. power does not have to be provided through wireless coupling). Peer-to-peer connectivity is required in this case, and since the sensor nodes are now the core of an isolated system , low cost and high integration are mandatory . In order to understand the power dissipation and data rate requirements of this transmission system, let us consider the power dissipation and clock rate of the microprocessors at its core. A generic 90nm low-power ARM core dissipates about $60\mu W/MHz$ clock rate ([20]). In specific applications, simpler cores such as Phoenix ([55]) could be used, leading to lower power consumption of about $3\mu W/MHz$. While the bandwidth requirements could ultimately be derived by first fixing the message-passing scheme of the nodes, and budgeting the communication power to be a certain percentage of the computation power, it is clear that unless the aggregate communication bandwidth is much smaller than the processor clock frequency, the communication power will be dominant. If the Phoenix core is taken as an example, and we assume that the communication bandwidth required between any two nodes is the same as the clock frequency, a link energy dissipation of 3pJ/bit is required in order for the communication power to be equal to the computation power. Since this energy dissipation is of the same order of magnitude as that of a state-of-the-art high speed processor-memory link, this task is clearly very challenging. A workable solution for a commercially available ARM core would require, under the same conditions, an energy-per-bit of 60pJ, which is still very challenging. As a result, it is clear that any likely solution to this problem will involve both an algorithmic and a circuit-level portion. For the circuit-level part, which is the focus of this work, we set out a *best effort* target, with the intent to achieve an energy consumption at least as low as 60pJ/bit. The data rate requirement will be calculated considering the available energy source for the nodes. To guarantee 1 year of operation out of a 320mAh Zinc-Air battery and assuming intermittent usage with a duty cycle of 10%, an active power consumption of $360\mu W$ is required for the system. If this budget is divided equally between computing, sensing and communication $120\mu W$ are available for each task, leading to a 2MHz system clock, and due to our assumptions, 2Mbps communication rate. The target specifications of our communication system can therefore be summarized as follows:

- Low PCB footprint of $1cm^2$
- Range of 5 cm



Figure 1.1. Pushpin nodes, a first implementation of the paintable computing concept [66]. Approximate node size is 3.24cm^2

- Timing Integration : to minimize cost and area, we strive to avoid off-chip timing references such as crystals.
- Power dissipation lower than $120\mu W$ at 2Mbps

1.1 Short range wireless systems today

Low power, short range ($\leq 10m$) wireless communication devices developed to date fall in one of three categories: narrowband, standard-compliant transceivers; narrowband transceivers with proprietary PHY specifications, and UWB transceivers. To date, most of the industrial development effort has focused on narrowband systems utilizing ZigBee or Bluetooth; while academic research has focused on both narrowband and wideband solutions with often proprietary air interfaces. In addition, ultra short range links with applications in wireless packaging [34],[19], contact-less payments [7] and implantable electronics for vital sign monitoring /biological study [18] have emerged.

1.1.1 Commercially available, low-power *short range* bi-directional communication devices

Table 1.1 below reports data on commercially available communication devices, as extracted from data-sheets or ISSCC publications produced by companies between the year

2004 and the year 2009. Virtually all of these products adhere to ZigBee or Bluetooth specifications, and obtain a power consumption in the range of 10 to 30mA at radiated power levels between -15 and 10dBm.

Ref.	Output Power [dBm]	Sensitivity [dBm]	Power (TX/RX)	Data Rate	Link Margin	Protocol
[51]	3	-96	18mA/20mA	2Mbps	92	BlueTooth
[57]	3	-96	29.4mA/23mA	2Mbps	92	Bluetooth
[60]	0dBm	-98	25.8mA/22.5mA	250Kbps	103	ZigBee

Table 1.1. Commercial Radio Performance Summary

One common feature of these products is that they are characterized by a link margin in excess of 90dB-corresponding to a free space range of 400m. Such large link margin guarantees robust operation in the presence of small and inefficient antennas as well as fading or obstacles but it also limits the power consumption in transmit mode, and further demands high selectivity in the receiver to reject signals due to nearby transmitters that stand above the noise floor. For example, the ChipCon 2520 radio features an adjacent channel rejection of approximately 50dB for 500KHz bandwidth channels. As a result, TX and RX frequencies need to be precisely aligned (better than 200ppm) to ensure reliable communication, demanding a crystal reference that limits system cost and size. Furthermore, the narrow-bandwidth allocation dictates the use of spectrally efficient modulation schemes such as GFSK, demanding more complex modulator/demodulator architectures and limiting the receiver power consumption.

While it is conceivable that one of these radios be used as coordination channel in the cooperative MIMO application, their power dissipation is too high for the other mentioned scenarios. In a situation such as paintable computing for example, the high transmitter output power would result in interference across multiple links. Even though in principle the receivers could be redesigned so that only nearest neighbors are within the transmitter range, the power consumption of the transmitters itself would be of several milliwatts, which is orders of magnitude too high compared to our target.

1.1.2 Bi-directional,narrowband short range radios resulting from academic papers

The above-mentioned limitations of radios available commercially have already resulted in a large amount of academic research in non-standard compliant transceivers with order-of-magnitude lower power dissipation than industry counterparts. The PicoRadio and Smart-Dust projects at U.C. Berkeley, and the $\mu Amps$ project at MIT focused on this development. A summary of contributions from these groups as well as from other research organizations is shown in Table 1.2. In order to reduce power consumption, these devices mostly rely on a simplified modulation schemes (OOK or FSK)in order to simplify the demodulator/modulator

Reference	Output Power	Sensitivity	Power (TX/RX)	Data Rate	Link Margin	Carrier Frequency
[44]	0	-99	2mW/0.4mW	5kbps	92	1.9GHz
[8]	-3	-96	1mA/0.4mA	400 Kbps	90	2.4GHz
[49]/[62]	10	-96	20mA./1mA	24Kbps	93	434MHz
[30]	N.A.	-104	5.4mW	1Mbps	N.A.	1.5GHz
[39]	N.A.	-50	65 μ W	250Kbps	N.A.	1.9GHz
[40]	N.A.	-72	52 μ W	100Kbps	N.A.	1.9GHz
[21]	-16	-93	400 μ W/350 μ W	40 Kbps	64	400MHz
[59]	-16	N/A	N/A / 400 μ W	40Kbps	N/A	400MHz

Table 1.2. Table 1.2: Recently published narrow-band low power radios

architecture at the cost of spectral efficiency. Also, RF-MEMS passives are introduced in papers [44],[39],[40] to obtain selectivity in the front-end at low power.

Receivers [39] and [40], compared to both other academic papers and industrial developments, reduce power dissipation in the receiver by an order of magnitude, and show that power dissipation can be exchanged for sensitivity (and to a certain extent, selectivity), indicating the possibility of realizing ultra-short-range wireless interconnection at reduced energy cost. To significantly reduce the power consumption of a complete link however, both TX and RX power dissipation should be scaled. Despite much effort, reducing transmitter power consumption substantially below a milliwatt has not been possible in radios employing narrowband signaling, due to the constraints on the frequency synthesis imposed by traditional modulation schemes. For example, recent papers [21] and [59], both targeting the MICS standard at 400MHz, radiate -15dBm while consuming over 300 μ W in the transmitter. As a result, further innovation is necessary to enable a peer-to-peer network with average energy consumption lower than 100pJ/bit in both transmit and receive mode.

1.1.3 Impulse Ultra-wideband radios resulting from academic papers

Impulse-radio UWB signaling has also received much attention, mostly from the academic community. Impulse-UWB signals are defined by FCC as those having an absolute bandwidth greater than 500MHz, or a fractional bandwidth (relative to the carrier) greater than 25%, and represent an extreme point in the tradeoff between energy and spectral efficiency. This type of signaling has potentially important benefits in a low-power network, as the large bandwidth and loose spectral mask utilized make the design of the frequency synthesis subsystem very simple, reducing the energy consumption in both TX and RX modes. An additional benefit lies in the ranging ability of these systems. Several recent publications on UWB transceiver chips/chipsets are reported in table 1.3. Energy consumption of the order 25pJ/pulse on the transmit side [48], and 0.5nJ/pulse [9] on the receive side have been

achieved recently by researchers at MIT. Papers [36] and [25] have also shown complete implementation of UWB systems with energy dissipation well below 1nJ/pulse. Despite the attractive numbers reported, significant problems remain open. The most important limitation of virtually all reported UWB radios is that they require spreading gain to mitigate interference from narrowband radios, as the wide front-end bandwidth makes filtering in the analog domain challenging. For a given data-rate, the use of spreading multiplies the system energy dissipation by a factor equal to the spreading gain (while improving sensitivity at the same time). When this is taken into account, UWB system energy consumption becomes similar to that of narrowband systems even in transmit mode. In addition, the use of spreading increases synchronization time as well as baseband clock stability requirements. In the system described in [42] for instance, the use of a spreading dictates the use of a medium quality quartz as time reference, increasing cost and board size. Another limitation, specific of the systems([36] , [42]) that operate in the 0-1GHz band, is the tradeoff between antenna miniaturization, bandwidth and efficiency. When operating with miniaturized antennas, the link-margin requirements of these transceivers become extremely challenging, leading to complex receiver architectures that are power-inefficient.

Reference	Output Power[Pk]	Sensitivity [dBm]	Power (TX/RX)	Data Rate	Link Margin	Carrier Frequency
[27]	N/A	-99	250 μ W	100kbps	N/A	3.5-4.5GHz
[25]	-10	-78	36mW/640 μ W	16Mbps	68	3.5-5 GHz
[36]	N/A	-52	4.5mW	40Mpps	N/A	0-1GHz
[42]	13dBm	-62	1.8mW/0.5mW	10Mbps	75	0-1GHz
[9]/[48]	0dBm	-75	22.5mW/320 μ W	16Mbps	75	3.5-4.5GHz

Table 1.3. Table 1.3: Recently published impulse UWB transceivers

At a different extreme, UWB systems developed recently for chip-to-chip IOs ([34],[19]) show excellent energy dissipation and reliability performance. Since these systems operate in close proximity, the transceiver circuits are designed under the assumption that interference is negligible and employ an exceedingly simple architecture. These systems lead to think that an hypothetical UWB system capable of reliable communication without using spreading could be a very power efficient solution for several of the ad-hoc connectivity scenarios. Understanding whether this supposition is realistic and how to engineer such a system constitutes most of the remainder of this dissertation.

1.1.4 Other short-range, low energy communication systems

It is finally worth mentioning that recently, low-power mm-wave systems aimed at short range data-transmission have appeared that achieved energy/bit close to our target. For example, the system in [32], and that in [46] consume less than 300mW while delivering 5Gbps, while the system in [1] (albeit incomplete) enters at 20mW of power dissipation for 2Gbps. While these works achieve extraordinary energy-efficiency, their suitability to a

situation characterized by short packets and long sleep times remains to be demonstrated, and will likely require some further development. Furthermore, the increase path loss at mm-wave demands more complex architectures. It is unquestionable however that thanks to the smaller antenna size, the opportunity of using mm-wave interconnects is indeed very promising and may finally lead to integration of the antenna on chip as well, with the associated system size reduction.

1.2 Thesis structure

The remainder of the thesis is organized in four main parts, corresponding to four main questions we try to answer .

- What is the path loss incurred by such an ultra-short range system, and what frequency should this centimeter link be operated at? (Chapter 2)
- What is the radio architecture most suitable to realize the link ? (Chapter 3 and 4)
- What circuits should be used to implement the target architecture ? (Chapter 5)
- What is the measured performance of the prototype (Chapter 6)

Chapter 2

Link Margin Modeling for cm-range radios

The goal of this chapter is to develop a model for propagation in near field coupled systems, so that we can avoid costly overdesign of our radio system. We have two objectives: first, to find an expression for power and voltage gain as a function of antenna size, field wavelength and antenna separation. We distinguish power from voltage loss because, differently from far-field coupled systems, systems operating in the near field present antennas with large quality factors, that can be exploited to achieve resonant gain. The second goal is to establish the impulse response of the channel in order to be able to predict multi-path effects. We will see that in both respects, short distance links are quite different from conventional ones. We'll use the results of this model to select an optimal (i.e. voltage-loss-minimizing) transmission frequency for our radio. Before delving deeper into the modeling issues, we'll use an example to highlight the limitations of Friis' Formular (Eq. 2.1) in this context. To understand this, let us consider transmission over a 400MHz carrier at a distance of 5 cm.

$$P_{rx} = P_{tx} G_{tx} G_{rx} \frac{\lambda^2}{4\pi r^2} \quad (2.1)$$

In the simplified case of $G_{tx} = G_{rx} = 1$, we find that $\lambda = 75cm$, and $P_{rx} = 17.9P_{tx}$, which is physically impossible as it indicates that the power received is larger than the power transmitted. This result makes the need for a better path-loss model obvious.




Figure 2.1. Schematic of electrical dipoles and frame of reference

2.1 Improving path loss modeling

The absurd result obtained above is a consequence of 2 main assumptions underlying Friis' Formula that, while usually satisfied in radio-systems, are not valid at cm-range. These 2 assumptions are

- The antenna separation is much greater than the antenna physical dimensions, so that transmitting and receiving antennas can be thought of as infinitesimal
- The antenna separation is much larger than the wavelength of electromagnetic waves transporting information, so that $1/r^2$ dependance is assumed and the radial and angular spatial dependences appear in a product form

In order to understand link budget at short distances, a new model is required that encompasses these effects. We'll first derive a simple model based on analytical expressions for elementary dipoles. Later we'll extend this model to optimize coupling using S-parameter data.

2.1.1 Electromagnetic wave propagation at short range

To understand the electromagnetic propagation in these short-range systems, we begin our analysis with elementary dipoles, as for this simple structure, exact expressions for the fields are available in the whole space (except at the origin, where the antenna is placed). Furthermore, elementary dipoles are approximate models of the wider class of electrically

small antennas [3]. In the following, we concentrate on electrical dipoles; however, due to the duality of Maxwell's equations the same results also apply to magnetic dipoles. Here, we call $k = \frac{2\pi}{\lambda}$ the wave number, and I the sinusoidal current of angular frequency $\omega = \frac{2\pi c}{\lambda}$ driven by our test source in the first dipole and h is the total length of the dipole. Equations 2.3-2.4 below ([3]) express the electric field \vec{E} , magnetic induction \vec{B} and Poynting vector \vec{S} in an arbitrary point of the space (excluding the origin) for an electric dipole of infinitesimal length h .

$$\vec{E} = -\frac{j\omega\mu I h e^{jkr}}{4\pi r} \left(\hat{r} \left(\frac{j}{kr} + \left(\frac{j}{kr} \right)^2 \right) 2 \cos(\theta) + \hat{\theta} \left(1 + \frac{j}{kr} + \left(\frac{j}{kr} \right)^2 \right) 2 \sin(\theta) \right) \quad (2.2)$$

$$\vec{H} = -\frac{\hat{\phi} j\omega\mu I h e^{jkr}}{4\pi r} \left(1 + \left(\frac{j}{kr} \right) \right) \quad (2.3)$$

$$\vec{S} = \eta \left[\left(\frac{k I h}{4\pi r} \right)^2 \left(\hat{r} \left(1 - \left(\frac{j}{kr} \right)^3 \right) \sin^2(\theta) - \hat{\theta} \left(\frac{j}{kr} \right) \left(1 - \left(\frac{j}{kr} \right)^2 \right) \sin(2\theta) \right) \right] \quad (2.4)$$

It is instructive to simplify these equations in two limiting cases. In the *far field* region ($r \rightarrow \infty$), terms with dependance on higher power of r can be dropped, leaving

$$\vec{E} = -\frac{j\omega\mu I h e^{jkr} \sin(\theta)}{2\pi r} \hat{\theta} \quad (2.5)$$

$$\vec{H} = -\frac{\hat{\phi} j\omega\mu I h e^{jkr}}{4\pi r} \quad (2.6)$$

$$\vec{S} = \eta \left[\left(\frac{k I h}{4\pi r} \right)^2 \hat{r} \sin^2(\theta) \right] \quad (2.7)$$

Which is the well known plane-wave solution used also to derive Friis' Formula.

When $\frac{r}{\lambda} \ll 1$, the system is said to be operating in the *near field* (More precisely the far-field, near field boundary is usually taken at $\frac{r}{\lambda} = \frac{1}{2\pi}$). Keeping only the terms in r^{-3} in \vec{E} and in r^{-2} in \vec{H} :

$$\vec{E} = -\frac{j\omega\mu I l}{4\pi r^3} \left(\hat{r} \left(\frac{j}{k} \right)^2 2 \cos(\theta) + \hat{\theta} \left(\frac{j}{k} \right)^2 \sin(\theta) \right) \quad (2.8)$$

$$\vec{H} = -\frac{\hat{\phi} j\omega\mu I l}{4\pi r} \left(\frac{j}{kr} \right) \quad (2.9)$$

The equations give the field-components that are sometimes referred-to as *Quasi-Static*. While we characterized near and far field regions in terms of the ratio of wavelength to distance between transmitter and receiver, due to the relationship $\frac{c}{nf_c} = \lambda$ (n denotes the refraction index in the medium) the same conclusions can be stated in terms of the product of carrier frequency to TX-RX wave propagation time $t_p = \frac{r}{cn}$. According to literature, for 5 cm distance in air ($n=1$), carrier frequencies smaller than $\frac{1}{2\pi t_p} = 900 MHz$ result in a near-field link, while carrier frequencies exceeding this value correspond to far field propagation. Obviously, the separation between these two regimes is not quite abrupt, and a *radiating near field* region exists in between them.

Considerations similar to the ones above can also be made for what concerns the ratio

of antenna size h to wavelength. An antenna such that the phase shift along its physical dimension is smaller than 36 degrees ($\frac{h}{\lambda} \leq \frac{1}{10}$), is called an *electrically small antenna* [26]. This is to be contrasted to the classic $\lambda/4$ condition for resonating dipoles ($\frac{h}{\lambda} = \frac{1}{2}$). Electrically small antennas are known to have a high-quality factor and therefore their radiation resistance can be easily swamped out by external losses due to connectors and traces [3].

	Near Field ($\frac{r}{\lambda} \leq 0.16$)	Far Field ($\frac{r}{\lambda} \geq 0.16$)
Electrically Small ($\frac{h}{\lambda} \leq 0.1$)		$\frac{h}{r} \leq 0.6$
Resonant ($\frac{h}{\lambda} \approx 0.5$)	$\frac{h}{r} \geq 3$	

Table 2.1. Operating regions in the size-separation-wavelength plane

Table 2.1 summarizes the different operating conditions of the link (columns) and the antennas (rows).

We see that in order to operate a link with electrically resonant antennas and near field coupling, the antennas should be larger than the physical separation between them, which is rarely practical. Operating electrically small antennas in the far-field coupling regime on the other hand requires an antenna size much smaller than half of the antenna separation, and therefore barely fits the requirements the system under design, where $\frac{h}{r} \approx \frac{1}{5}$. The system we are designing is therefore better described as falling into the *radiating near field* case.

In order to arrive at the voltage gain of the channel, and examine the different merits of these operating regimes, we start by calculating a 2-port equivalent circuit model for the wireless link. It is easy to convert the field equations into a Z-parameter description of the antenna pair, as

$$V_{oc} = \int_{\Delta L} \vec{E} dL \rightarrow Z_{12} = \frac{V_{oc}}{I} \quad (2.10)$$

Using the expressions above, the Z-matrix for a pair of identical dipoles can be calculated for all values of frequencies and antenna separations (To be accurate, one can only approximating a dipole of finite length with an elementary dipole when $h \ll \lambda$, so our results for trans-impedances will be inaccurate when $r \geq .1\lambda$). Using the expressions for small antenna and near fields, we find

$$Z_{11} = Z_{22} \approx \frac{O}{j2\pi\epsilon h\omega} \quad (2.11)$$

$$Z_{12} = Z_{21} = \frac{h^2}{j4\omega\pi\epsilon r^3} \nu \quad (2.12)$$

$$\nu = \langle (\hat{\theta} 2 \cos(\theta) + \hat{\theta} \sin(\theta)) \mid \hat{a} \rangle \quad (2.13)$$

Where a is the diameter of the dipole wire, $O = 2 \log(\frac{2h}{a})$ is a measure of the wire thinness and \hat{a} is a unit-length vector indicating the direction of the receiving dipole in a frame of




Figure 2.2. Equivalent circuit for monopoles in the near-field operating range

reference having the other at the origin. The resulting equivalent network is shown in figure 2.2 and comprises only of capacitors. Capacitor values for this network can be extracted from the parameters noting that

$$j\omega Z_{11} \approx \frac{1}{C_{11}} \rightarrow C_{11} = \frac{2\pi\epsilon h}{O} \quad (2.14)$$

$$j\omega Z_{12} \approx \frac{C_c}{2(C_1^2 + C_1 C_c)} \rightarrow C_c = \frac{h^4 \pi \epsilon \nu}{2O^2 r^3 - h^3 O \nu} \quad (2.15)$$

The case of resonant antennas and far-field operation occurs when $\lambda = \frac{h}{2}$. In this case

$$Z_{11} = Z_{22} \approx 77\Omega \quad (2.16)$$

$$Z_{12} \approx -\frac{j\omega\mu h e^{jkr}}{4\pi r} \hat{\theta} = \frac{j\sqrt{\frac{\epsilon}{\mu}} h \exp(jkr)}{4r} \quad (2.17)$$

We can now compute the path loss under these two special conditions. As mentioned before, since in this system impedance level at transmitter and receiver can be freely chosen¹, we'll search expressions for the voltage gain of the channel, as opposed to the power gain. For the case of near field coupling, assuming $\nu = 1$, maximum voltage gain is obtained by series-resonating the input port, and parallel resonating the output port. The transfer in this case can be approximated by

$$A_v^{NF} \approx \frac{C_c}{C_{11} + C_c} Q^2 = Q^2 \frac{h^3 \nu}{4O r^3 - h^3 \nu} \approx \frac{Q^2}{4O} \left(\frac{h}{r}\right)^3 \quad (2.18)$$

Where we assumed the matching network has a quality factor Q , and the antenna itself is lossless. Note that this voltage gain is independent of frequency, as long as this is low enough that the antennas are electrically small, and the link is operating in the near field. The case of resonant antennas is more difficult. In fact, the term Z_{12} has a phase that depends on the TX-RX separation expressed in wavelengths, and an absolute value depending on antenna separation to physical size ratio. In short range systems where $r/h \leq 10$, we find $|Z_{12}| \geq 9\Omega$, which is hardly negligible compared to the dipole 77Ω impedance. While

¹As we are targeting full integration and the system comprises no Ceramic/SAW filters

if the TX-RX separation was known one could devise a matching network that takes into account this component, this approach is impractical, due to the strong dependence of the impedance phase on separation embodied in the complex exponential. As a result, we argue that the only practical strategy consists in driving the transmitting dipole with an ideal voltage source, and receiving on an open-circuited termination. Doing so results in a path loss

$$A_v^{FF} \approx \left| \frac{Z_{12}}{Z_{11}} \right| = \frac{377h}{77 \cdot 4r} \approx \frac{h}{r} \quad (2.19)$$

We see that as known from the literature, the voltage gain of a near-field channel exhibits a $1/r^3$ dependence, as opposed to the $1/r$ behavior of the far-field ones. As a result, for a given antenna geometry and technology, there will be a critical antenna separation at which far-field coupling become more effective than near field coupling in transferring information. Equating 2.18 to 2.19, and assuming $r \geq h$, we find that radiative transfer becomes more effective when

$$r \geq h \frac{Q}{2\sqrt{O}} \approx hQ \quad (2.20)$$

This approximate analysis establishes that the relative advantage of using near field versus far-field coupling only depends on the ratio of antenna physical size to TX-RX separation. *Normalized range* ($\frac{h}{r}$) is therefore the critical link design parameter taking into account both size and range constraints.

In addition, the higher the matching network quality factor the larger the *normalized range* at which it is advantageous to use near-field coupling over far field coupling. In other words, large matching network quality factors result in a smaller optimal operating frequency; while at low quality-factors, the lower intrinsic propagation loss of far-fields is preferred.

Bandwidth and component value considerations

In the previous section, we found that the signal gain in the near field region is essentially independent of carrier frequency. Since for a fixed signal reducing the carrier frequency reduces power consumption, one would then be tempted to operate with $f_c \approx 0$. Obviously however, even if a high-Q matching network is available, there are limits to what is the minimum carrier frequency that can be used. First, for a given quality factor Q, reducing the carrier frequency reduces the available signal bandwidth. For data transmission with bandwidth BW, defined the intrinsic quality factor $Q_I = \frac{f_c}{BW}$, in order to prevent signal loss due to bandwidth limitations, the link should be operated at a frequency high-enough that $Q_I \gg Q$. In addition, lowering the operating frequency requires larger values of capacitance or inductance, which could be unavailable for either a technological or a cost limitation.

Consider for example the case of $h = 1cm$, $r = 5cm$, $Q=10$, and consider a data-rate $R_b = 1MBps$. Since $\frac{r}{h} \leq Q$, we should operate the link in the near field and with electrically small antennas, which results in $f_c \leq 900MHz$. To ensure $Q_I \geq 10Q$, we are also required $f_c \geq 100MHz$. Since lower carrier frequency, for the same path loss, results in lower power, the link should be operated with $f_C = 100MHz$.

If on the other hand under the same conditions we choose operate at a data-rate $R_b =$

100Mbps, we find that $f_c \geq 10GHz$, which implies that the link should be operated in the far-field regime.

2.1.2 Magnetic Coupling

Due to the symmetry of Maxwell's equations, the same derivation holds unchanged also for magnetic loops, provided the electric dipole momentum U_e is replaced with a magnetic dipole momentum U_m . There are however three main differences between inductive links and capacitive links

1. The designer of inductive links has an extra degree of freedom given by the number of turns in the inductor. While it can be proven ([26]) that this parameter does not impact the low-frequency transfer, it can be used to increase the electrical size of the loop given a physical size, or equivalently to shrink the physical size of a loop for a given electrical size.
2. The low-frequency impedance level of an inductive link is much lower than that of a capacitive link.
3. Last, and perhaps most important, one needs to consider that when targeting hybrid implementations, where the antenna is off-chip and the matching network is on chip, inductive link take advantage of the higher quality factor of integrated capacitors over integrated inductors at low-frequencies.

As a result, magnetic antennas are clearly at an advantage when matching network integration is planned, the antenna is off-chip, and we are operating in the near field. Magnetic antennas can also result in lower operating frequencies when operating in resonant antenna mode for a given size constraint.

Radiating Near Field

Several questions remain open in the *Radiating Near Field* case. We would like to know whether operating in the far-field region of the link but below the antenna resonant frequency can be attractive compared to the two simple cases treated above. In addition, since a range of frequencies fall into the region $\frac{h}{\lambda} \leq 0.16$, $\frac{r}{\lambda} \geq 0.16$, we would like to know whether it is better to operate closer to the upper or to the lower limit of this region. Unfortunately, no simple formula is available for this operating conditions, and to gain understanding in this regard we'll have to resort to numerical techniques discussed in the next section.

2.1.3 Model Verification

We resort to FDTD simulation and measurements to verify and refine our existing model. We focus our verification effort on monopoles antennas for their ready availability and simplicity; however, the results are valid in a broader context. In order to readily understand and process simulation or measurements results, which are typically in S-parameter format, we need a procedure to calculate the link budget that takes measured or simulated S-parameters and optimizes the matching network for voltage gain at any given frequency.

² The equations used for this S-parameter modeling are:

$$Y_s = jB_s(1 - j\frac{\text{sign}(B_s)}{Q}) = 50\frac{1 + \Gamma_s}{1 - \Gamma_s} \quad (2.21)$$

$$Y_l = jB_l(1 - j\frac{\text{sign}(B_l)}{Q}) = 50\frac{1 + \Gamma_L}{1 - \Gamma_L} \quad (2.22)$$

$$\Gamma_{in} = S_{11} + \frac{S_{12}S_{21}\Gamma_L}{1 - S_{22}\Gamma_L} \quad (2.23)$$

$$\frac{V_r}{V_s} = \frac{S_{12}(1 - \Gamma_s)(1 + \Gamma_L)}{2(1 - \Gamma_L S_{22})(1 - S_{11}\Gamma_{in})} \quad (2.24)$$

Where Q is the quality factor of the matching network and B_s and B_l are the source and load susceptances. The values B_s^{opt}, B_l^{opt} for B_s and B_l maximizing this function are determined numerically at each frequency point. This allows choosing the optimal terminations even in the face of the finite Q constraint, which is not trivial analytically.

To verify our calculations, we performed both FEM simulations and measurements on pair of wire monopoles. The scattering-matrix of the antenna pair was obtained as a function of frequency in three different ways: through calculations (using a short dipole model), HFSS simulations and measurements. To reduce artifacts due to the surrounding connectors and metal objects, the measurements were taken on a scaled-up model of the desired antennas constituted by 5cm long copper-wire monopoles. The monopoles have a calculated resonant frequency of 1.56GHz (measured value 1.4GHz). In order to take measurements over a ground plane at different ranges, a custom measurement setup, consisting of a 14 inch X 14 inch copper clad board with 1-inch pitch holes large enough to fit an SMA connector, was used. The holes do not affect the measurement as they are much smaller than the wavelength in the frequency range of interest ($f \leq 2GHz$). S-parameters for collinear monopoles were then obtained using a Vector Network Analyzer (VNA) for different values of TX-RX antenna separation (See figure 2.7). Calculated, simulated and measured results are compared in figure 2.3 to 2.6 and show good agreement.

We can now discuss how the optimal frequency selection for data transmission derived from measurements compares to the calculated one. The computed and measured optimal coupling frequencies (normalized to antenna resonant frequency) versus range (normalized to antenna size) are reported in figure 2.8. We notice that for large values of Q ($Q \geq 100$), and very small normalized range the optimal point lies at the low end of the frequency range.

²In principle, this procedure could also be based on Y-Parameter equations and rely on an S-Y parameter conversion, however, this conversion is prone to numerical errors as it involves inverting a matrix that for is almost singular for S-parameter values close to 1, and gives inaccurate results at low-frequency.




Figure 2.3. Measured, calculated and simulated return loss for wire monopoles

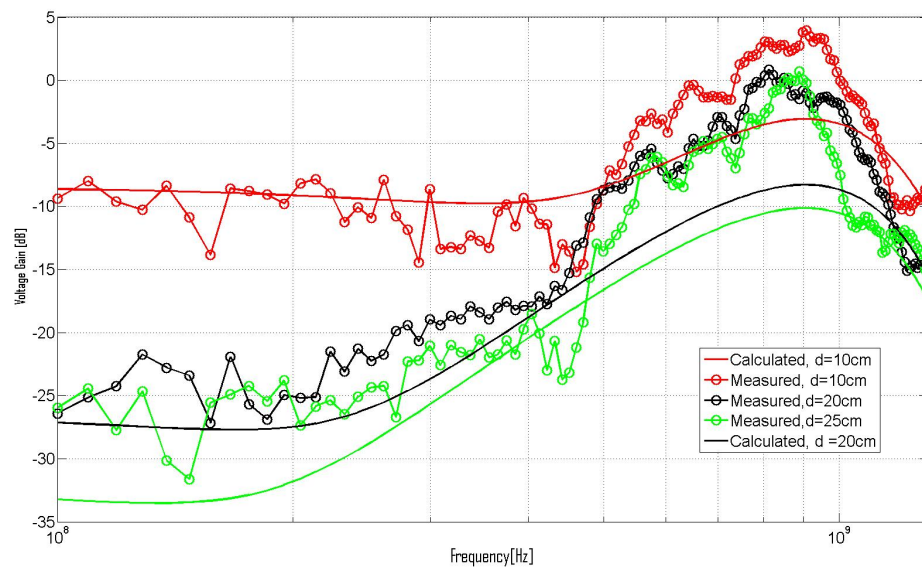


Figure 2.4. Calculated and measured path loss as a function of frequency




Figure 2.5. Measured and simulated path loss as a function of frequency

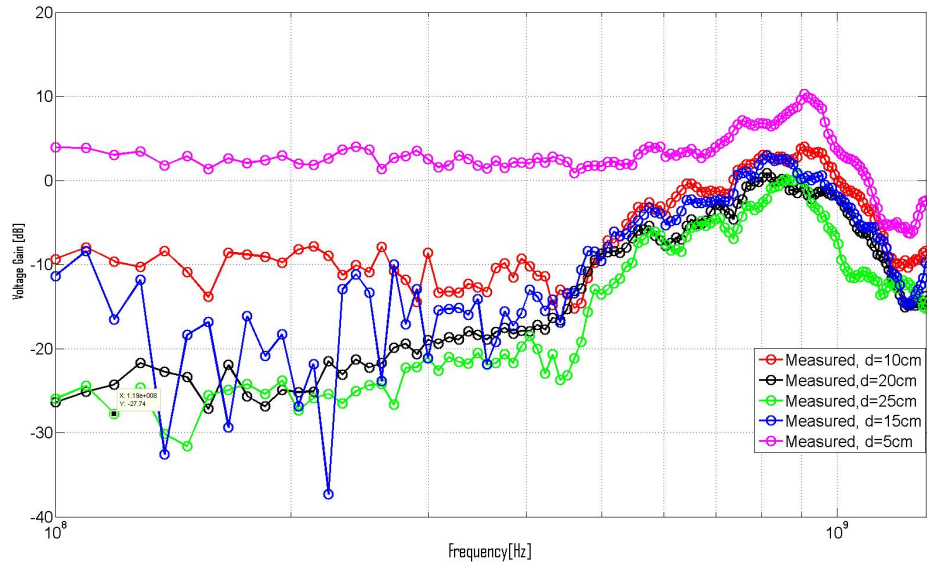


Figure 2.6. Measured path loss (including matching network) for the TX-RX Antenna pair at different separations

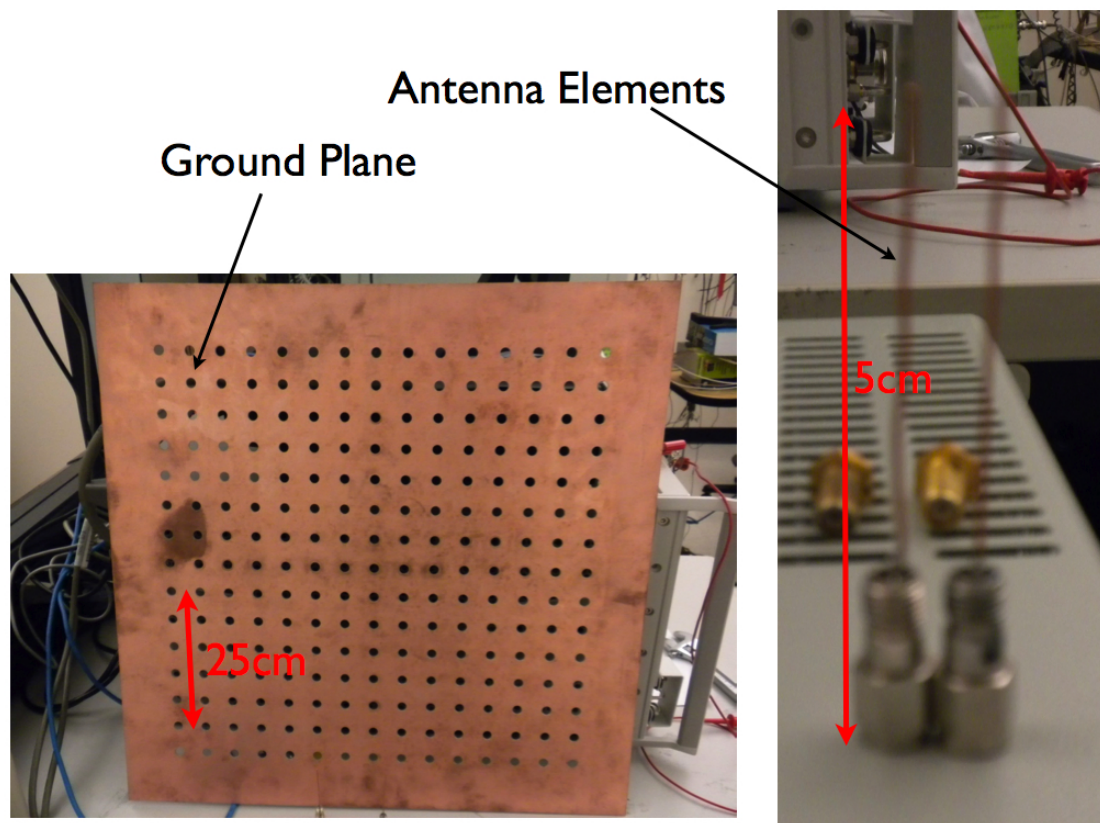


Figure 2.7. Experimental setup in the BWRC Lab

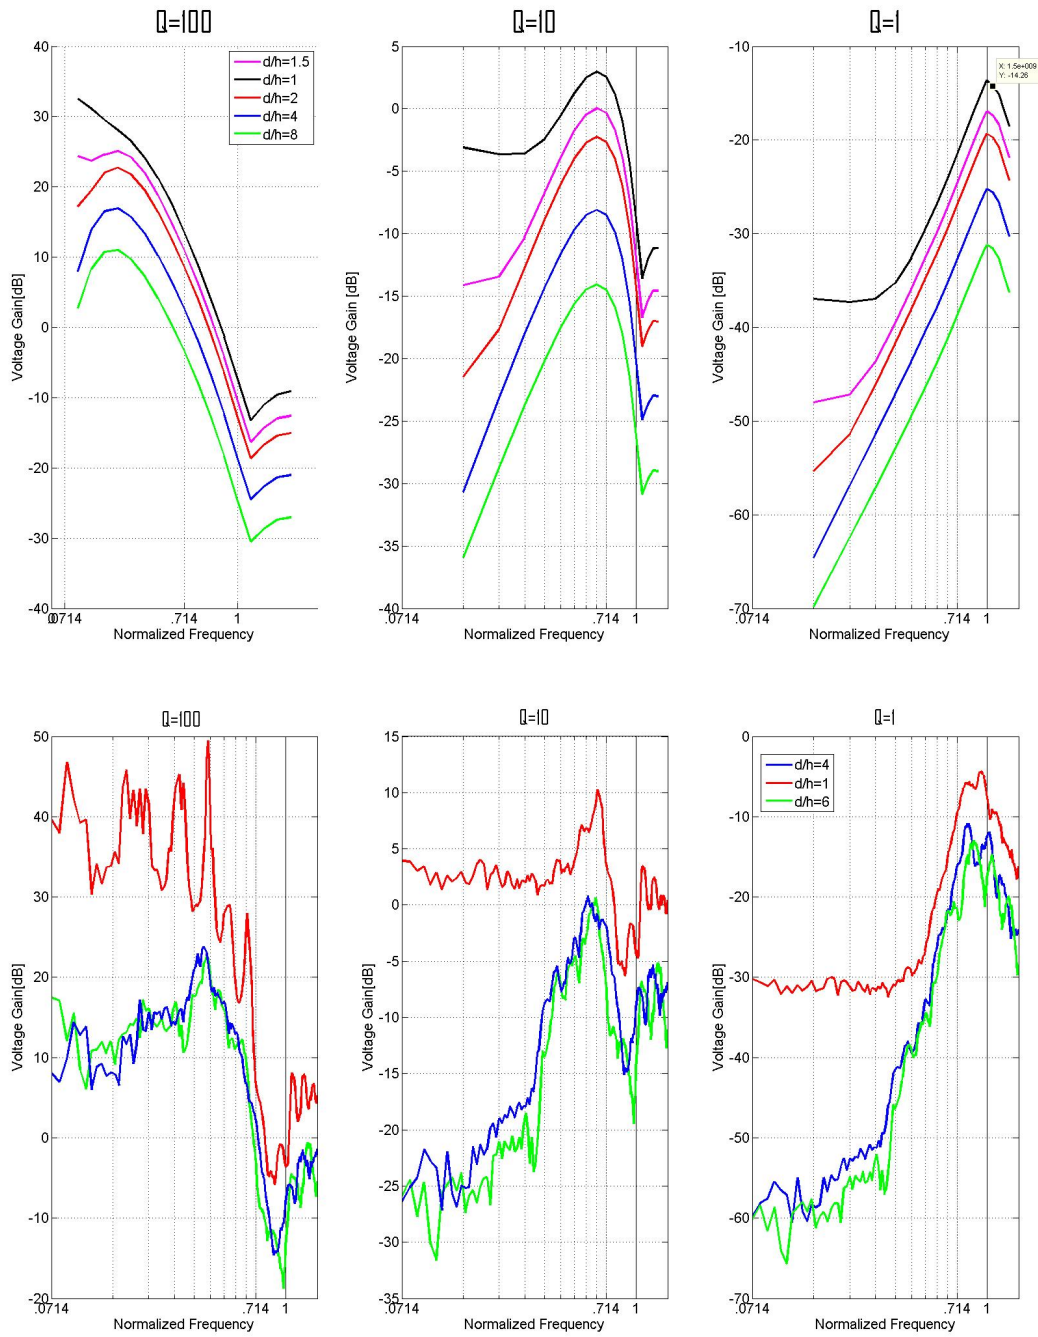


Figure 2.8. Calculated(top) and measured(bottom) loss as a function of frequency for different antennas separations and quality factors




Figure 2.9. Effect of matching network quality factor on voltage gain




Figure 2.10. Measured and simulated path loss as a function of range for different frequencies

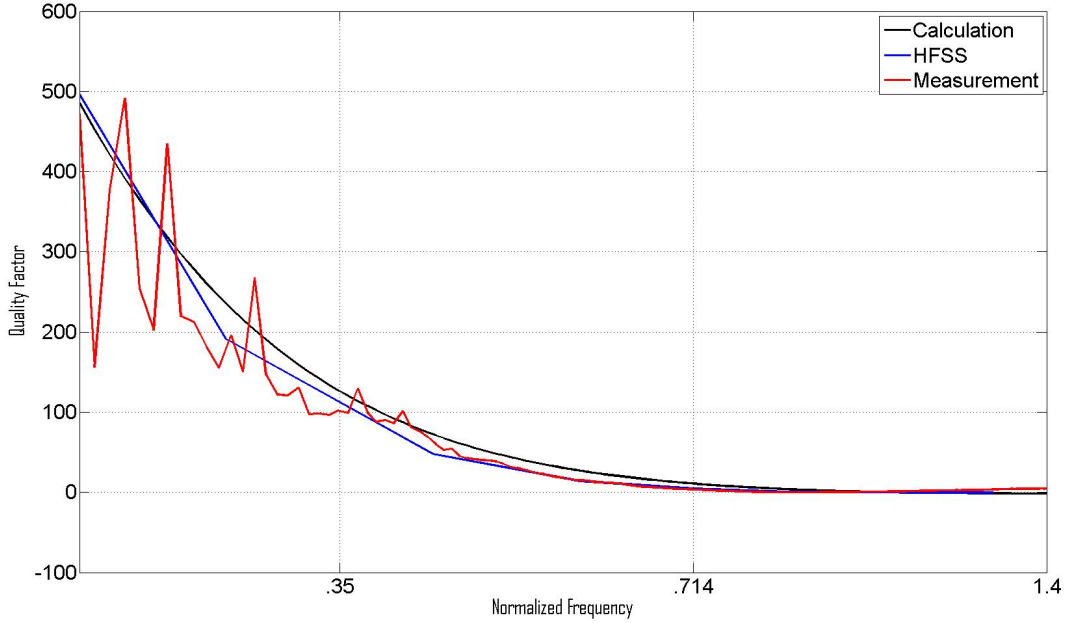


Figure 2.11. Antenna impedance quality factor as a function of frequency

For very low Q values ($Q = 1$), the resonant frequency of the antenna is the optimal choice. For quality factors $10 \leq Q \leq 100$, a frequency lower than the antenna resonant frequency is optimal. At this frequency, the coupling is largely radiative, and yet the antenna operates off-resonance. This could not be anticipated in our initial simplified derivation, as we only looked at the extremes of low-frequency and resonant coupling. We find instead that operating in the *radiating near field* ([2]) is advantageous.

Measured, calculated and simulated results are all in qualitative agreement. From a quantitative standpoint, the discrepancy between the calculated and measured values is smaller for smaller TX-RX separation and for lower frequency. This is expected since at lower frequency the antenna structure more closely resembles the elementary dipole assumed in the calculations. More accurate calculations could be performed using numerical integration, but the increased accuracy does not justify the loss in intuition. HFSS simulations and measured values on the other hand appear to match within 5dB over a wide frequency range. This is an acceptable agreement especially given the relative simplicity of the simulation model used, and the presence of reflections in the measurement.

Figure 2.10 shows the path loss as a function of distance for 100MHz, 650MHz and 1.3 GHz, for a matching network quality factor of 10. We see that the path loss at 1.3GHz and that at 100MHz become equal when the antenna separation equals about twice the antenna size. As the TX-RX separation is increased beyond this value, operating the antenna at resonance results in lower loss than operating in the near field regime.

However, the channel gain obtained operating the antenna off-resonance ($f = 70\%f_R$) never becomes smaller than the coupling at resonance. As a result, it is always advantageous to operate the link off-resonance.

The quality factor of the antenna input impedance as a function of frequency is shown in figure 2.11. At low frequency, the antenna quality factor and impedance both increase dramatically. This results in poor accuracy for a 50Ω S-parameter measurement, since the sensitivity of the reflection coefficient to changes in impedance for such high impedance values is quite small, and explains the noise in the lower frequency portion of Fig.2.11.

At the optimal frequency (1GHz, or 0.7 Normalized to the resonance) a quality factor Q of 6 is obtained from the antenna, making matching with on chip elements easy and guaranteeing a roughly 10% fractional transmission bandwidth. Assuming these results scale when frequency and antenna size are scaled so as to keep their product constant, for a 1cm linear dipole, and a matching network Q of 10, an optimal transmission frequency of 5GHz is found, resulting in a 3dB bandwidth of about 500MHz. Unfortunately, this analysis was not completed before the system design phase was completed and as a result we chose to design for a straight radiative channel (i.e. using resonant antennas), paying a price of approximately 10dB in propagation loss, while having the benefit of a broader available transmission bandwidth.

2.1.4 Comparison with power transfer

In recent years the problem of determining the optimal carrier frequency for power transmission in a short range radio link has been tackled in the context of biomedical implants ([10],[2]). While these problems share the common substrate of a trade-off between antenna efficiency and path loss, they are differentiated by the different propagation medium (tissue versus free-space) and in the quantity of interest (voltage gain in our case, power gain in the case of remotely-powered implanted systems). To get further verification of our methodology, the case of wireless power transfer in free-space was also analyzed. In this case, an S-parameter formulation can also be used, but the quantity of interest become now the transducer power gain G_T , which can be expressed as

$$\Gamma_{in} = S_{11} + \frac{S_{12}S_{21}\Gamma_L}{1 - S_{22}\Gamma_L} \quad (2.25)$$

$$G_T = \frac{1}{1 - \Gamma_{in}\Gamma_S} |S_{21}|^2 \frac{1}{1 - S_{22}\Gamma_L} \quad (2.26)$$

We optimize this function at each frequency numerically, taking into account a finite quality factor of the matching network, Q. The results of this optimization are shown in table 2.2, and indicate that power and voltage optimal frequencies are very close to each other. Also calculated power transfer values are within 2-3 dB of the experimental results-further confirming the analysis procedure.

Range (normalized)	Quality Factor	Opt.Frequency for Power [Calc/Meas]	Opt. Frequency for Voltage [Calc./Meas.]	Power Gain [dB]	Voltage Gain [dB]
1	10	0.71/0.71	0.71/0.71	-11.43/-5.25	3/10.3
1	100	0.07/0.07	.07/.07	-0.9/-1.1	33/39
1.5	10	0.71	0.57	-14/-12	0/3
1.5	100	0.28	0.28	-9 / -8	25/26.17

Table 2.2. Comparison of calculated and measured power and voltage transfer parameters

2.1.5 Chosen design point and channel measurements with cm-scale antennas

The previous analysis indicates that operation slightly below the antenna resonance minimizes channel power loss. In a pulse-based system, time dispersion is however also a concern. For a 1cm-long dipole in particular, a resonant frequency of 7 GHz is predicted, and the path-loss optimal propagation frequency should be approximately 5GHz. As we see in the next chapter, we also utilize impulse-based signaling to reduce transmit power. According to our calculations, operating with a pulse width of 2nS and an antenna+matching network quality-factor of 6 results in acceptable time dispersion for this carrier value. However, since we could not verify in measurement this conjecture prior to the prototype design, we made a conservative design choice of using a broadband antenna, such as the one commercially available from SkyCross , operating at a 6GHz carrier. The choice of a broadband antenna makes the design less prone to dispersion, but sacrifices the voltage gain attainable by incorporating the antenna into the matching network. In addition, since the SkyCross antennas have a size of approximately 1.8x2 cm, we also designed a custom circular monopole antenna to demonstrate the feasibility of size-reduction, as described in the next section. Figure 2.12 reports the magnitude of S_{21} as a function of TX-RX separation for a pair of Skycross broadband antennas. Note that since these antennas exhibit broadband impedance match, S_{21} is a good approximation of the channel power gain. For 5 cm range (leftmost end of the plot) the measured loss under the test conditions the expected channel loss is between 24dB and 28dB depending on the antenna orientation. The angle θ measures rotation with respect to the axis of the SMA connector. A best fit to the data using a power law is shown in black. Through this fitting, a loss exponent close to 2 was found, indicating that the effect of multi-path on this channel should be small. This assertion is confirmed by the experiments described in the next section.

2.1.6 Channel impulse response

To complete our wireless channel characterization, we further investigated the time-domain characteristics of transmissions and the relevance of multi-path phenomena at this




Figure 2.12. Measured Channel S_{21} as a function of TX-RX separation

short range. Intuitively, given the short distance between transmitter and receiver, there should be many circumstances and applications in which a strong line-of-sight path can be assumed, and hence multi-path effects can be expected much weaker when compared to a general-purpose wireless channel. In fact, considering a simple ray-tracing model and loss-less reflections, only obstacles in a 25 cm range from the transmitter will result in echos to main path ratios higher than -20dB. Furthermore, only reflections with distance from the transmitter smaller than 30cm ($= \frac{cT_p}{2}$) will have the chance to arrive before the main path has extinguished and hence to cause fading. The channel impulse response was measured through omni-directional commercial antennas and a logic analyzer in the BWRC lab. The figures below (2.14 and 2.13) show the received energy as a function of the utilized number of taps at 2GS/s, for a *Bad* Channel (i.e. with a nearby obstacle) and with a *Good* channel (suspended). We see that the first 3 taps of the channel approximately contain 90% of the energy in the impulse response, with the first contributing between 40 and 65%. While this is significantly fewer taps than required by a generic indoor environment, this channel is still not completely flat. However, the simplicity gain by avoiding a rake receiver still justifies the loss, given that in this channel SNR is high and complexity is costly.

2.2 Antenna Design for a 3-10GHz wideband system

In view of the previous considerations, we tackled the design of an antenna for radiative communication. The system necessitates approximately a 1GHz bandwidth around 6GHz




Figure 2.13. Received energy as a function of tap number




Figure 2.14. Channel Impulse Response Sampled at 2GS/s for good channel: I-component (top) and Q component (bottom)

for impedance, time dispersion and radiation pattern; and as a result we choose 2GHz as a minimum bandwidth design target in order to accomodate manufacturing tolerances and ensure low time dispersion. Our goal was to demonstrate that FR-4 based printed antennas with $1cm^2$ size can achieve these performance figures. A 6 GHz carrier frequency translates into a free space wavelength of $5cm$. As a result, a dipole length of $2.5cm$ would be required to ensure resonant operation, which does not fit into our area budget. Even taking account a material effective dielectric constant of 2 for the FR4 substrate, we still find a required size of $1.76cm$. Furthermore, fractional bandwidth of dipoles typically does not exceed 10%, which is a factor of 3 short of our target of 30% bandwidth.

A literature survey of wide-band small antennas was conducted and slotted bow-tie antenna and a circular monopole antenna were identified as promising candidates. We then fabricated several antenna structures in order to verify whether their performance would be sufficient for our application, feeding back into our design methodology. Based on the design process and measurements that we report in what follows, we conclude that the allocated area budget of $1cm^2$ and the chosen operating frequency and bandwidth are compatible. However, correlation between simulation is poor as a result of the proximity of many components to the radiating element itself, and further work should be done to improve this aspect.

2.2.1 Slotted Bowtie Antenna Design

The bow-tie antenna is the simplest broadband extension of a dipole, and corresponds to a planar approximation to the frequency-independent bi-conical antenna [3]. By tapering the open wires, broader radiation bandwidth is achieved. A slotted bow-tie is chosen in this design ([38]) as it provides a single-ended feed and thus avoids a balun. The degrees of freedom in the bow-tie antenna design are the angle and the length of the taper. Since inductive loading is used in this design to reduce the physical size of the antenna, the placement and length of the loading stubs become a variable as well. Starting from an initial design with a taper angle of $60degrees$ and an arm length of 4mm per arm, the antenna was sized through numerical simulations performed using ANSOFT HSS software (See figure 2.16). The goal of this design was to demonstrate a center frequency of 8 GHz, reflecting an initial belief that it would not be possible to achieve the required bandwidth at 6 GHz. The final annotated schematic is shown in figure 2.15.

The bowtie antenna was fabricated on an 63 mil thick FR-4 substrate with 1oz (35μ) thick copper. In order to avoid error-prone manual layout translation from HFSS to Cadence Allegro, the HFSS solid model was exported to a DXF file using HFSS, and the DXF file translated into a gerber file by the software E-Machine-Shop. The gerber file can then be imported by the PCB design tool Allegro and submitted for fabrication. The measured response of the bowtie antenna is shown in figure 2.17. It is clear that the performance is significantly degraded compared the simulations. This discrepancy is attributed to the effect of the taper and the presence of the SMA connector in proximity to the antenna. While a setup using a longer feed, and potentially anti-reflective coating of the connector would contribute to isolating the antenna performance, it is clear that when operating in an actual system, the antenna would be located in close proximity to several conducting bodies (chip,




Figure 2.15. Annotated Bowtie Antenna Schematic




Figure 2.16. Bowtie Antenna Simulations: Effect of varying the expansion angle (top); effect of varying the inductive stub length(bottom)

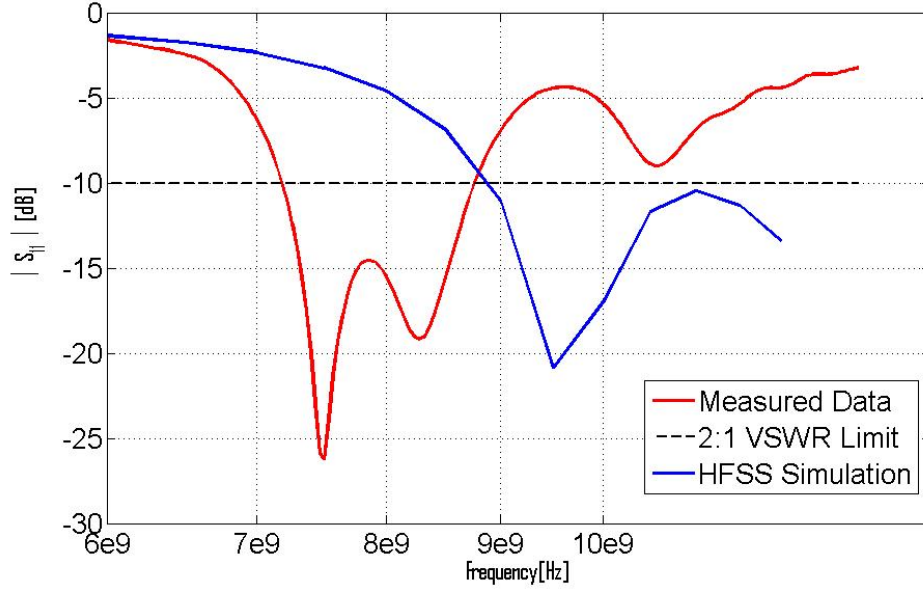


Figure 2.17. Bowtie Antenna Measured Reflection versus HFSS Simulations

power supplies, etc.). As a result, a test-setup with a close connector might represent more realistic a use-case.

Because this experiment did not provide the desired degree of model hardware correlation, we could not build enough confidence in the design methodology to operate the system off-resonance. We therefore chose use broader band antenna structures (see next), in a effort to over-design around the potential dispersion issues of the resonant antennas and demonstrate functionality of a complete system, including the antenna. Further antenna optimization is left as future work to improve performance.

2.2.2 Monopole antenna design

To improve system robustness and demonstrate that a broadband antenna could also be realized in the 1cm^2 area, different instances of circular monopoles were also designed [23]. These antennas are well known to have a broad radiation bandwidth, and a relatively constant gain over a wide range of frequencies. The design equations for circular monopoles suggest that the monopole diameter d should be such that ([29])

$$\frac{d}{\lambda} = 0.23 \quad (2.27)$$

Where λ is the electromagnetic wavelength corresponding to the lowest operating frequency. Setting this frequency equal to 5GHz results in $d=1.4\text{ cm}$, while a diameter of 1cm corresponds to a 7GHz lower cutoff. Circular monopoles with radius of 4.5mm , were built on FR4 boards to compare their performance to calculations and understand wether they

could constitute an attractive alternative. Their annotated mechanical drawings are shown below (Figure 2.18). The measured results (Figure 2.19) from these antennas show broad impedance bandwidth between 6 and 12 GHz . This shows that this antenna structure is adequate for the system specifications we are targeting. Measurements also revealed that impedance bandwidth was consistently better for structures fed using microstrip feed as opposed to coplanar waveguides. As for the bow-tie antennas, the performance of these antennas appears to be significantly different than simulated results, most likely due to the presence of the SMA connector in close proximity to the radiating structure.

2.3 Conclusions

We have developed a design-oriented model for the wireless propagation at cm-range . This model provided the key result that for a given antenna size, a low carrier frequency (such that the transmitter and receiver antennas are electrically small and located in the near field of one another) is advantageous for distances that are small compared to the product of this size and the available matching network quality factor. When the TX-RX separation is larger than the product of antenna size and matching network quality factor, using a carrier frequency such that antennas operate as efficient radiator results instead in lower loss. We verified and refined these predictions using a methodology based on S-parameter measurements of a given antenna pair, such as the one proposed in [43]. This more accurate analysis indicates instead that operating the antenna slightly off-resonance , in the so called *Radiating Near Field* region , results in even higher voltage gain thanks to the combined gains from the matching network as well as the low-loss propagation.

Finally, we also demonstrated the feasibility of broad-band antennas with 1cm^2 area and 6GHz center frequency in a low-cost FR-4 based PCB process. Using such antennas, a voltage path loss of approximately 28 dB is obtained under worst case conditions at 5 cm range. In the next chapter we will investigate how to convert this low loss into a low power transceiver system.





Figure 2.18. Monopole antennas annotated schematic: full monopole (top) and loaded hollow monopole (bottom)

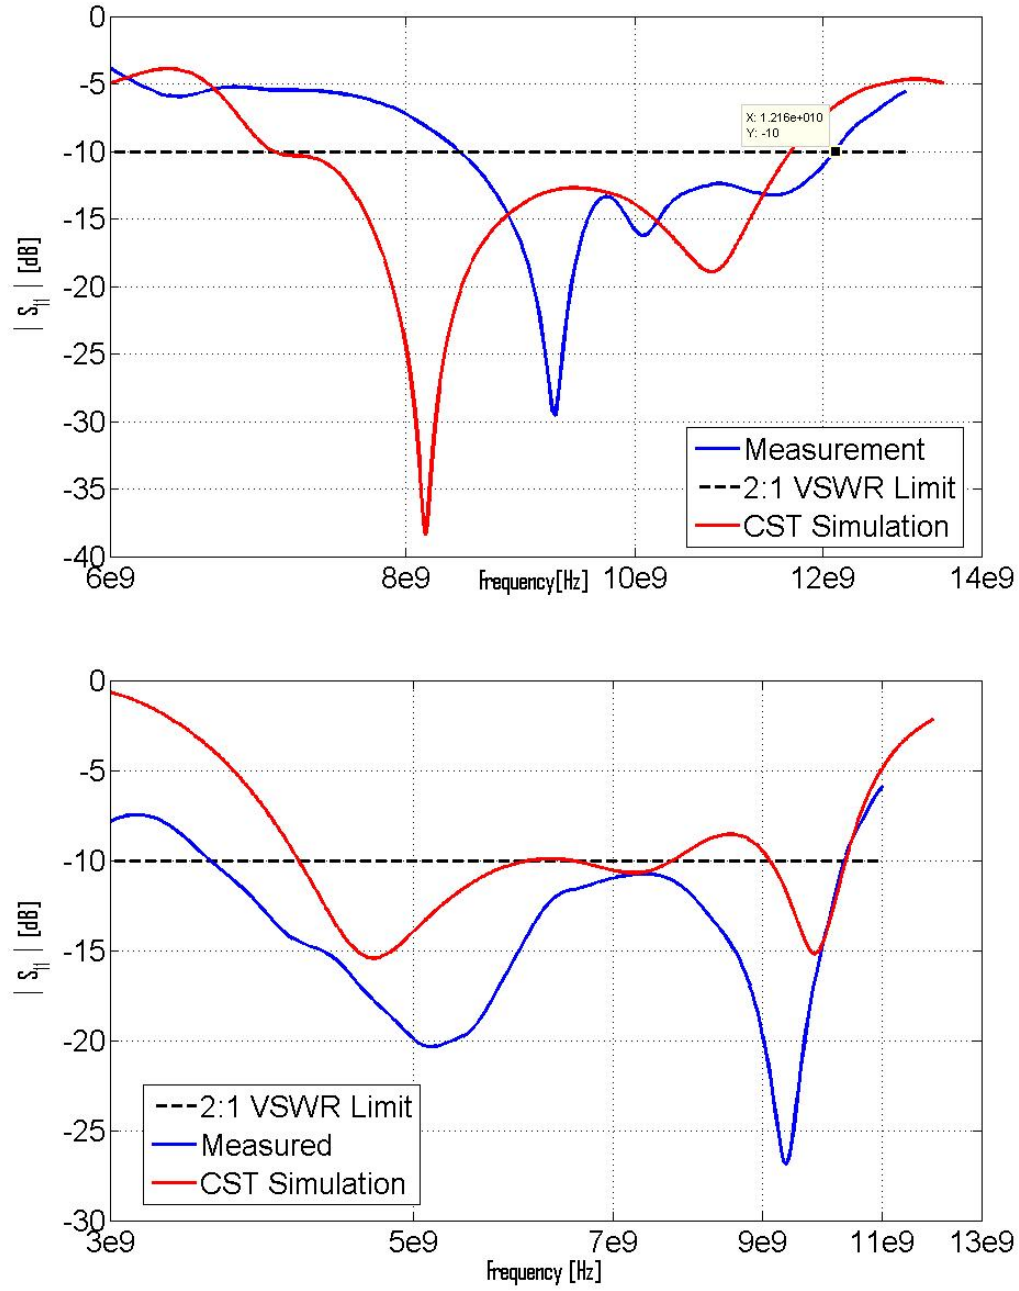


Figure 2.19. Measured results from monopole antennas (Full monopole (top) and loaded hollow monopole (bottom))

Chapter 3

Low Power RF System Design for a cm-range Transceiver

We would now like to understand how to turn the large received signal strength that characterizes this cm-range link into power reduction. To this end, we start this chapter by reviewing the dependence of power dissipation on sensitivity specifications for a generic radio. We first consider the effect of thermal noise only, and then briefly extend to gain and selectivity. The analysis leads to the conclusion that in order to take full advantage of the range reduction, we should use an AM-radio architecture and employ duty-cycled signaling.

3.1 Radio Dynamic Range-Power Tradeoff

In this phase, we model a generic radio transceiver as composed of three parts : a Power Amplifier (P.A.) a Low Noise Amplifier (LNA) and overhead power. Even when, as in [4], we postulate that a noise-free, zero-power receiver is available, a Power Amplifier is necessary, as thermal noise floor and path loss always determine the minimum energy that should be radiated in order to be able to communicate. In a low-power transceiver, a low-noise amplifier is not always necessary [5], as noise figure specifications are often relaxed. When this block is omitted, the considerations in this section hold then for the first amplifying stage . Finally, all the other power spent to generate the RF carrier and demodulating the signal at baseband is lumped into *Overhead Power* P_{OH} .

While overhead power actually depends on selectivity specifications, we'll regard it as a constant in this analysis, which is entirely focused towards understanding limitations due to thermal noise. For a given PA efficiency ϵ , LNA noise factor F_{LNA} , carrier frequency f_c and




Figure 3.1. Optimal Radio Range as a function of total radio power budget

path-loss exponent β , signal-to-noise requirement SNR_{lin} , signal bandwidth BW and factors

$$\alpha = SNR_{lin} BW KT \quad (3.1)$$

$$\gamma = (F_{LNA} - 1)P_{LNA} \quad (3.2)$$

$$\lambda = \frac{3e8}{f_c} \quad (3.3)$$

[4] calculates an optimum value for the maximum radio range attainable and the optimum PA/LNA power ratio, which are expressed by Eqns.3.4-3.5 as a function of the total "signal path" radio power, $P_{sum} = P_{LNA} + P_{PA}$.

$$r_{max} = \left(\left(\frac{\lambda}{4\pi} \right)^2 \frac{e}{\alpha} \frac{(P_{sum} + \gamma - \sqrt{\gamma^2 + \gamma P_{sum}})(-\gamma + \sqrt{\gamma^2 + \gamma P_{sum}})}{\sqrt{\gamma^2 + \gamma P_{sum}}} \right)^{\frac{1}{\beta}} \quad (3.4)$$

$$\frac{P_{LNA}}{P_{PA}} = \frac{-\gamma + \sqrt{\gamma^2 + \gamma P_{sum}}}{P_{sum} + \gamma - \sqrt{\gamma^2 + \gamma P_{sum}}} \quad (3.5)$$

Figure 3.1 plots these values as a function of design parameters f_c and P_{sum} for a values $\gamma = 10^{-3}$ (Corresponding to a power dissipation of 1mW for an LNA to obtain a 3dB noise figure on 50 Ω) and BW=1GHz and SNR=20dB. According to this model, even assigning a $1\mu W$ value to P_{sum} , (meaning that the total of LNA and PA power dissipation is only $1\mu W$) a range of 30cm is obtained for a carrier of 6GHz and our desired range of 5cm can only be obtained using a carrier value in the 90GHz range.

For a realistic power consumption of $P_{sum} = 100\mu W$, the maximum range obtained is 0.9m at 6GHz, which is in great excess of the requirement. The conclusion of this analysis is that power dissipation for this radio scenario will not be dominated by low-noise and power amplifiers, but by *overhead* components. To give an idea of to which extent overhead




Figure 3.2. Maximum Radio Range as a function of carrier frequency

components limit dissipation, consider the systems [8] and [40]. In [8], the 2.4GHz LC RF DCO accounts for over 90% of the $400\mu W$ receiver power consumption. In [40], a ring oscillator is used to improve this figure to 40% ($20\mu W$ s out of a $52\mu W$ total). In the 65nm process in use, a 6GHz ring oscillator, excluding drivers, has a power consumption of $120\mu W$. If CMOS clock distribution is used, the power rapidly increases above a 1 mW, which is unacceptable. Techniques to reduce P_{OH} in general and especially the portion due to the RF VCO must then be investigated.

3.1.1 Overhead power components

Determined that the design will be overhead-limited, it is worth trying to analyze in more depth the reasons limiting over-head power dissipation. For both the transmitter and the receiver side, overhead is mainly constituted by the cost of generating a stable RF carrier, and by the cost of conditioning and demodulating the baseband signal.

Baseband power

The baseband stages of a receiver have the goal to condition the downconverted signal for A/D conversion or slicing. The power cost associated with this operation is spent both regenerating the signal as well as rejecting out-of band interference. In general, determining the minimum power consumption of a baseband filter for a given set of noise and rejection specifications is a complex problem and cannot be solved in closed form.

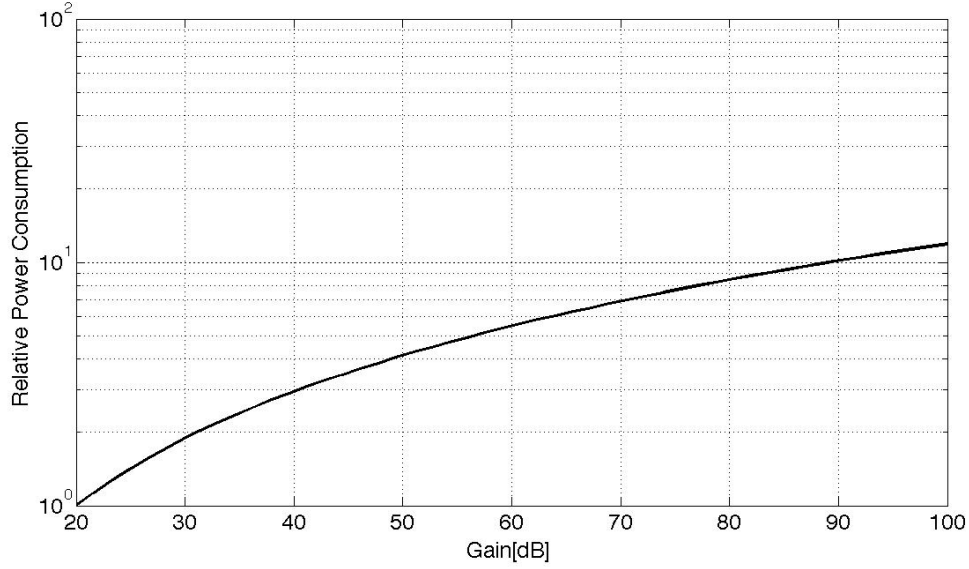


Figure 3.3. Power dissipation versus total baseband gain

When thermal noise is not a consideration however, the baseband stages can be constructed using small transistors and the power dissipation is only determined by the minimum capacitance allowed by the process, the signal bandwidth and the total amount of gain. In the simplest case we neglect selectivity and design a gain stage for a total amount of gain A_v and a bandwidth ω_{3dB} . Using an open loop cascaded gain chain with $N = 2 \log(A_v)$ equal stages results in minimum power dissipation, that equals (3.6)

$$P_d = V_{dd} 2e \log(A_v) \frac{\omega_{3dB} C_T n \frac{KT}{q}}{\sqrt{(2^{\log(A_v)} - 1)}} \quad (3.6)$$

Where it was assumed that the transconductor transistors operate in subthreshold, which is typically the case for scaled technologies and low-bandwidth standards. The dependence of power dissipation on gain value is essentially logarithmic (See 3.3). This implies that the power consumption of a gain-only-limited baseband chain will not be radically improved, even when the radio sensitivity (and hence the required baseband gain) is reduced.

VCO Power

The power consumption in a voltage controlled oscillator is determined by the requirement of short and long term stability. Long term stability is typically obtained using a phase-locked loop frequency synthesizer. For low-power applications, using a phase-locked loop often more than doubles power dissipation compared to the same free-running VCO. As a result, ultra-low power architectures have to avoid if possible a PLL, and employ open-loop frequency sources.

Short term stability specifications on the other hand originate from selectivity requirements and translate into specifications on the phase noise of the oscillator. Generally speaking, the phase noise of an oscillator is inversely proportional to the product of the square of the quality factor the tank Q and the signal power $P_{carrier}$, and proportional to the oscillator frequency itself:

$$L(\Delta\omega) = F\left(\frac{\omega_0}{\Delta\omega}\right)^2 \frac{k_B T}{Q^2 P_{carrier}} \quad (3.7)$$

Furthermore ([40]), an ultimate limit on oscillator power is dictated by the start-up conditions, when no phase noise constraint is considered.

$$gm \mid Z(\omega_0) \mid \geq 1 \quad (3.8)$$

For noise-limited VCOs, in order to reduce the power, we can either reduce the phase noise specification (by filtering prior to downconversion), increase the Q , or increase the fractional offset of the interference. In case the VCO is instead startup limited, maximizing impedance $Z(\omega_0)$ and minimizing the supply voltage are the main knobs.

In general, lower frequency VCOs tend to consume lower power than higher frequency ones, unless a better phase noise performance is required. In a system where the interferer characteristics over frequency were well known, one could choose the VCO frequency considering both startup and noise limitations, as well as channel loss, with the goal of minimizing power. This is however limited by FCC regulation and also results in a design which is highly application specific and might not work robustly in the field, when not-expected interferers are present.

Another technique to reduce VCO power dissipation is that of duty cycling the RF front-end, including the VCO itself, at a bit-by-bit level. This theoretically reduces the VCO power by a factor equal to the duty-cycle δ , while the extra complexity and power dissipation associated with implementing the duty-cycling clock. This approach is chosen in this thesis and will be analyzed in the next section.

3.1.2 Summary

This analysis has showed that shrinking the communication range to 5 centimeters results in only relatively small gains in the power dissipation of conventional radio architectures. Thermal noise limitations in the front-end do not limit the power due the small path loss, and the reduction in baseband gain and selectivity specifications can only bring a modest benefit, especially when compared to state-of-the-art sensor network radio designs (such as [8]), where similar considerations have already been exploited. One opportunity for power reduction lies in reducing the phase noise specifications of the local oscillator, as a result of the increased signal power available at the input. We have seen that the power consumption of a *startup limited* oscillator could still too high for our specifications. As a result, we introduce front-end duty cycling as a mean to reduce power consumption of the overhead components.




Figure 3.4. Transceiver Duty Cycling Concept: implemented by multiplying continuous waves with short windows (top) and implemented by power gating transmitter and receiver in a synchronized fashion (bottom)

3.2 Effect of duty-cycling on power dissipation

We now consider the effect of shrinking the duty-cycle of the transmitted signal, and similarly the duty cycle with which the receiver listen to the channel, in a synchronized fashion. This idea is shown in figure 3.4. Considering, for simplicity, OOK modulation, the transmitter baseband impulse is a square wave with period T_b and duty cycle δT_b . The bandwidth of this impulse (measured to the first null) is then $\frac{1}{\delta T_b}$. In order to proceed with the analysis, we define three more parameters :

- T_m , the time required for the receiver to power up and settle all internal bias point with sufficient accuracy to process the signal
- T_{start} , the transmitter start-up time. This time is defined symmetrically with respect to T_m , but is generally shorter (see below).
- The relative jitter of transmitter and receiver baseband clock, called here T_j .

The receiver, after acquiring the bit-boundaries, turns on only for a fraction of time $\delta T_b + T_m + 2T_j$ to downconvert, process the signal and decide whether a logic 0 or a logic 1 was transmitted. We are interested in understanding how the duty-cycling operation impacts the power dissipation of the transmitter and receiver blocks. For simplicity, we take the perspective of constructing a duty-cycled receiver by *scaling* an existing narrowband receiver. In other words, we assume that we have available a (still On-Off Keyed) radio design operating at bit-rate $R_b = \frac{1}{T_b}$. By decreasing the signal duty cycle δ , we increase inversely the signal bandwidth. We can then derive its power dissipation based on the following considerations:

1. The receiver duty-cycling is limited by the time necessary to startup-up the system, plus the pulse duration. As a result, the average power consumption P_{avg} (measured

over a bit period) is related to the instantaneous power consumption P_{inst} by $P_{avg} = P_{inst}(\delta + \frac{T_m}{T_b})$.

2. Bandwidth of RF stages is dictated by quality factor of the passives utilized, mainly inductors. Since typical quality factors are of the order of 10, the equivalent bandwidth of each RF stage will be of the order of 6GHz/10=600MHz. As a result, at least for low-data-rate systems, we can assume that even though duty-cycling increases signal bandwidth, we will not need to re-design the radio-frequency stages to account for this increase. Equivalently, we can say that due to technological limitations, RF gain stages are typically over-designed for bandwidth in narrowband systems. As a result, we find that the instantaneous RF power P_{RF} is constant as a function of the duty-cycling δ .
3. The bandwidth of baseband stages on the other hand is set by capacitance and resistor values only. Baseband stages are typically designed to accomodate a bandwidth similar to that of the channel being down-converted. As a result, these stages have to be re-designed for the increased bandwidth. This can be done by decreasing the impedance level of the circuit, simultaneously raising transconductance values to keep gain constant. We can identify two regions for this design. In the low-bandwidth region, transistors can be operated in the weak inversion region. Transconductance can here be increased linearly by increasing the bias current without changing the transistor capacitance, which is dominated by overlap. This leads to an increase in instantaneous power consumption that is linear with the increase in bandwidth, $P_{bb} \propto \frac{1}{\delta}$. This process however increases the transistor inversion level. Eventually, transistors are operated in the strong inversion region, where transconductance depends quadratically on bias current. In this region, power dissipation depends quadratically on the the bandwidth $P_{bb} \propto \frac{1}{\delta^2}$. In this work, transistors will be always biased in the weak-inversion region, so that $P_{bb} \propto \frac{1}{\delta}$.

Under these assumptions, we can model the power dissipation of these blocks with the following equations

$$P_{TX} = P_{TX}^{(ON)}(\delta + \frac{T_{start}}{T_b}) + P_{synch}(T_b, T_m) \quad (3.9)$$

$$P_{RX} = P_{LNA}(\delta + \frac{T_m}{T_b}) + P_{bb}(1 + \frac{T_m}{\delta T_b}) + P_{VCO}(\frac{T_m}{T_b} + \delta) + P_{synch}(T_b, T_m, \delta) \quad (3.10)$$

The dependence of P_{VCO} on duty cycling deserve further analysis, which will be the topic of a next section. Under the assumption that it can be kept constant as the signal duty-cycle is decreased, it decreases linearly with δ , towards the value of $P_{VCO} \frac{T_m}{T_b} \ll P_{VCO}$. Therefore duty-cycling can be an effective means to reduce overhead power consumption due the VCO. Notice also that the chosen duty-cycling strategy impacts receiver sensitivity requirements. As mentioned, we assume P_{TX}^{ON} is independent of the duty cycling δ , and so is P_{LNA} . Keeping P_{TX}^{ON} constant implies that the peak output power is constant, reducing the energy per-symbol E_b proportionally to δ . Keeping P_{LNA} constant implies (to first order) that the receiver noise factor is constant. These two effects combine in a reduction of the energy/bit at the slicer E_b/N_0 (N_0 , as usual, denotes the equivalent white noise power spectral density

at the slicer input), which now equals $\frac{\delta E_b}{N_0}$. Thus the sensitivity of the receiver is reduced, which is acceptable for our application as our preceding analysis of thermal noise shows that a narrowband scheme would result in a range greatly exceeding our desired 5 cm. For a transmitter peak output power of 3dBm, a receiver with 20dB noise figure and a slicer SNR of 20dB, the minimum usable duty-cycling factor for a bandwidth of 1MHz and a channel with 30dB of path loss is given by

$$\delta_{min} = 10^{-5.1} = 7.9 \times 10^{-6}$$

much smaller than the optimal duty-cycling factor derived below considering only minimization of power dissipation.

3.2.1 Overhead Costs in duty-cycled systems

Reducing the signal duty cycle can be quite effective in reducing the major overhead cost of a conventional narrowband transceiver (i.e. the VCO), while introducing a small penalty in the baseband power consumption. However, low-duty-cycle signaling requires the transmitter and receiver to achieve accurate synchronization. Such synchronization block represents the major overhead cost for the duty-cycled system.

In the simplest case, this block only needs to generate a power-gating signal of width $T_m + \delta T_b$ for the receiver every T_b seconds. The phase of this signal should be adjusted so that the energy of the pulse received during the window is maximized. There are several ways to implement this function, and a conceptually simple one (adopted in [11]) is to build an oscillator running at frequency $\frac{1}{2(T_m + \delta T_b)}$ (See 3.5), embedded in a phase-locked loop. Referring to figure 3.5, every N-th rising edge of this oscillator, a power gating signal is generated that last 1 oscillator clock cycle. A negative edge triggered flip-flop is then used to align the center of this interval to the incoming pulse, similar what is done in a CDR system.

If we expect the power consumption of this synchronizer to be dominated by the VCO and the divider, the power dissipation in this block is at least ¹ inversely proportional to $\frac{1}{T_m + \delta T_b}$.

$$P_{synch}(T_m, T_b, \delta) = \frac{k}{T_m + \delta T_b} \quad (3.11)$$

Even if we could make $T_m = 0$ then, an optimal duty-cycle value for power would exist and be determined by the increasing synchronizer power incurred at very large signal bandwidths.

3.2.2 Duty cycling optimization

We are now in a position to derive the optimal duty-cycling factor δ for a transmitter-receiver pair. Figure 3.6 shows the results for total power as a function of duty-cycling δ , for different values of T_m . The values of the parameters are here $P_{LNA} = 100\mu W$, $P_{VCO} = 100\mu W$, $P_{tx} = 2mW$, $k = .125pJ$ and $T_{start} = 1nS$. Somewhat unsurprisingly, we find that

¹Power reduction from scaling the digital supply voltage is not considered here

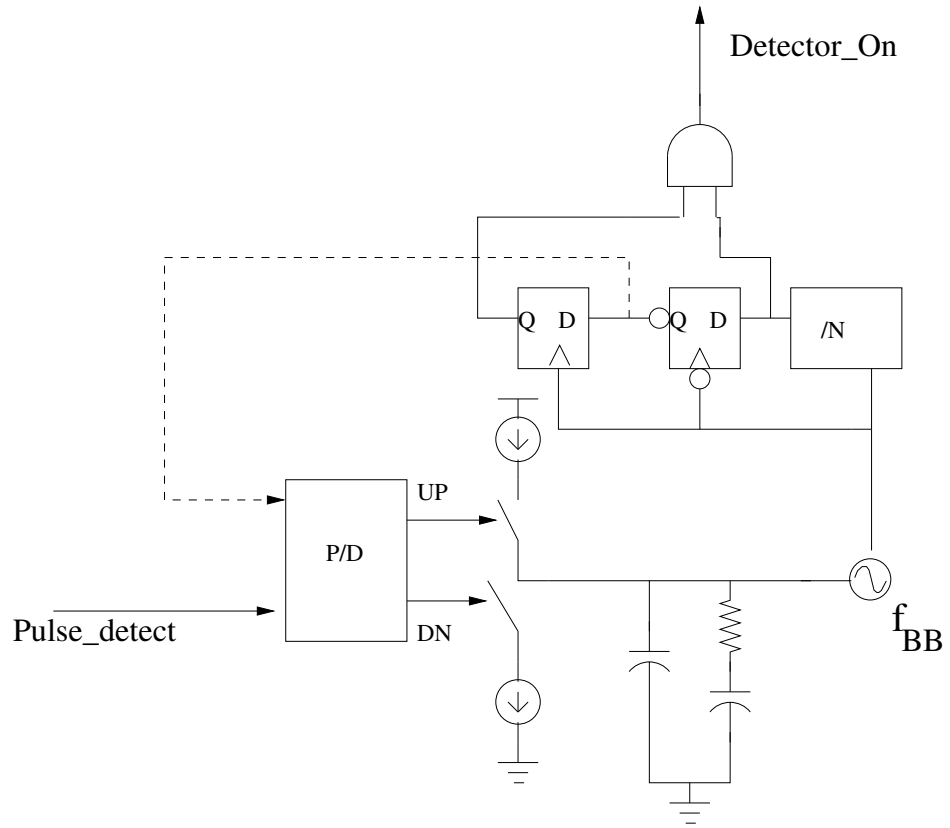


Figure 3.5. Analog Synchronizer

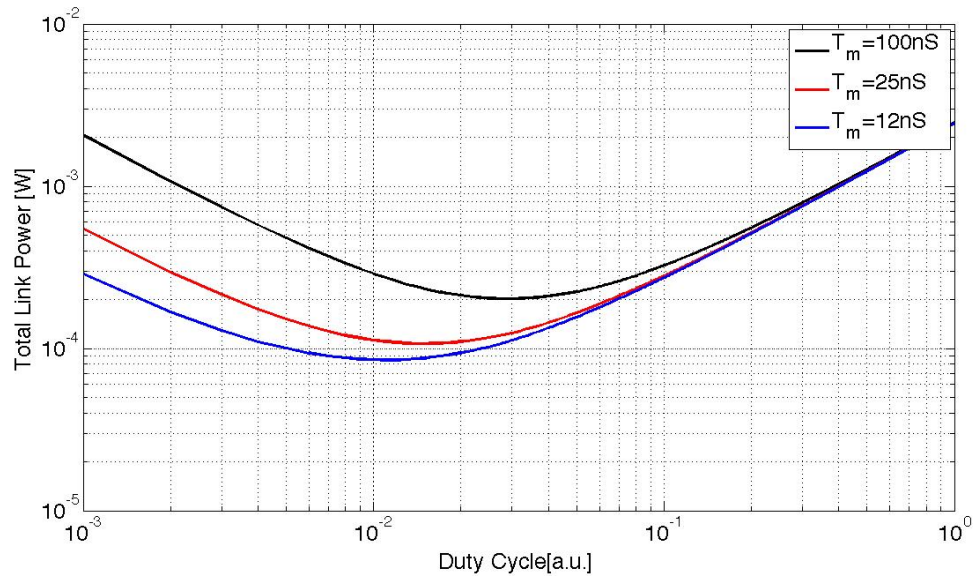


Figure 3.6. Power Dissipation of TX+RX as a function of duty-cycling




Figure 3.7. Minimum Power Dissipation as a function of receiver turn-on-time T_m

the optimal duty cycling is such that $\delta T_b \approx T_m$. This is intuitive as, for a fixed turn-on-time T_m , reducing the pulse duration below this point only causes an increase in the signal bandwidth and a degradation in sensitivity, without any power benefit. It is also interesting to investigate the dependence of the minimum power dissipation on T_m . A smaller T_m reduces the power in signal-path circuits by reducing the duty-cycle, but it simultaneously increases the jitter and speed requirements of the synchronizer. As a result, there is a point after which reducing T_m does not improve the total link power. This relationship is shown in figure 3.7. For the values used in this work, decreasing the receiver turn-on-time below 5 nS pays small returns in total power. More importantly, we see that duty cycling allows a reduction in power dissipation (at the expense of sensitivity) by a factor of 4 to 5 compared to a narrowband solution, making it a very attractive option for our application.

For what concerns this work, we chose not to design for the optimal transmission duty-cycle, as the optimization results in transmission bandwidths that are only allowed by FCC around certain carrier frequencies. Instead, we conform to UWB (Part 15) signaling regulations, which allow signals with bandwidth exceeding 500MHz and power spectral density below -41.6dBm/MHz to be radiated with arbitrary carrier frequency between 3.1 and 10.6GHz. This is further justified by the fact that the optimal duty-cycling value corresponds to a relatively shallow minimum, so that there is not a big penalty in operating at smaller duty-cycles. Complying with this regulation leaves freedom in the carrier frequency choice signal design, but puts an upper limit on the signal duty-cycle chosen.

3.2.3 Inherent Asymmetry of duty-cycled systems

One interesting property of duty cycled systems is their *Inherent Asymmetry*. This asymmetry arises from the fact that the transmitter does not need to estimate the position of the incoming pulse, and as a result does not need to use a wider turn-on period to account for jitter. This pushes naturally duty-cycled transceivers to be receiver limited, rather than transmitter limited. Also, since receiver circuits are typically class-A, low noise circuits, they rely on accurate bias points to achieve reliable operation. The settling of these bias lines contributes to make T_m larger than T_{start} . While achieving low-turn-on time is a challenge, this asymmetry represents an advantage of duty cycled systems, as designing an efficient, low-power transmitter is generally more difficult than designing a narrowband low-power receiver.

3.2.4 Power consumption of VCO and baseband: a more critical look

In the previous analysis, we assumed that VCO power P_{VCO} does not depend on duty-cycle. While these assumptions are to first-order true, they require further discussion. Since the main effect of decreasing the transmission duty-cycle is increasing the signal bandwidth, selectivity degradation is a concern. For what concerns the VCO, since the phase-noise spectrum is naturally rolling-off as $\frac{1}{f^2}$, the increased channel bandwidth results in looser specifications. To understand this fact, consider a system with a 1MHz bandwidth, and an interferer at 11MHz offset and 10dB higher peak power than the desired signal (Fig. 3.8). Using a Lorentzian model for the phase noise-corrupted voltage spectrum, with corner frequency f_c ([12]), the noise floor caused by reciprocal mixing and the signal power can be expressed by

$$P_n^{(rm)} = 2P_{INT}P_{LO}\left(\text{atan}\left(\frac{f_{INT} - f_{LO} + f_{BW}}{f_c}\right) - \text{atan}\left(\frac{f_{INT} - f_{LO} - f_{BW}}{f_c}\right)\right)$$

$$P_{sig} = P_{sig}P_{LO}$$

The resulting signal-to-noise ratio is

$$SNR_{rm} = \frac{P_{sig}}{2P_{INT}\left(\text{atan}\left(\frac{f_{INT} - f_{LO} + f_{BW}/2}{f_c}\right) - \text{atan}\left(\frac{f_{INT} - f_{LO} - f_{BW}/2}{f_c}\right)\right)} \quad (3.12)$$

Assuming $f_c/f_{LO} \approx 10^6$ ([12]) for a low quality oscillator, and $f_{LO} = 5GHz$, we obtain $SNR_{rm} = 30dB$. If the duty-cycling of the signal is decreased so that bandwidth is increased to 10MHz, and the interferer frequency offset from the carrier is set to 20MHz (So that it is still 10MHz away from the band edge), a ten-fold increase in thermal noise bandwidth is incurred. For the same interferer power and LO phase noise however, the reciprocal mixing contribution only increases by 30%. As a result, an LO with looser spectral purity can be used and LO power can scale more dramatically than with an inverse law when signal bandwidth is increased.




Figure 3.8. Desensitization due to reciprocal mixing

Different considerations hold for the base-band filters and converters. As the bandwidth is increased, more interferers are downconverted in the signal band, raising the baseband linearity requirements. In many standards (WIFI, WiMedia) ensuring reliable transmission requires OFDM-based processing and high-resolution A/D converters to ensure graceful degradation of the error-rate in the presence of blockers. These complex solutions require a power budget far exceeding that of the short-range applications we are targeting. While interference should be mitigated, this power budget ultimately will require a tradeoff between reliability and energy. This tradeoff will be eased for our applications by the large amplitude of incoming signal impinging the antenna.

3.3 Radio Architecture Design

Now that we identified duty-cycled transmission as a viable approach to realize our short-range radio, we can proceed to define the details of our system. First, we first set boundary conditions restating what are the technology and environment driven constraints we need to satisfy. We then proceed to compare existing radio architectures reported in literature, on the basis of power dissipation and interference robustness, showing that none of the existing solutions is suitable to communicating in the presence of interference without resorting to spreading. Based on this analysis, we conclude that an envelope-detector based receiver can reduce power consumption with respect to a receiver exploiting a local oscillator. We also show how to improve the interferer robustness over existing receivers by combining mixed-signal baseband processing with a reconfigurable RF front-end.

3.3.1 Technological Constraints

We repeat here the two specific characteristic of this radio system, as they impact much of the design choices that will be described in what follows. This system is characterized by two specific issues and opportunities

1. Low channel loss (high SNR) with line-of-sight propagation
2. Single chip implementation

The presence of a line-of-sight channel with minimal multipath dictates that the a single tap-approximation to the channel response is sufficient. This makes a complex, rake-based digital back-end un-necessary. Also, low channel loss reduces the receiver dynamic range, making receiver jamming by a large blocker less likely. We make use of these facts to speed-up signal acquisition and simplify the digital back-end.

The absence of any off-chip component, and specifically a frequency reference, makes the most process-insensitive frequency reference an on-chip LC tank, which has a total variability of the order of $\pm 5\%$. This limits from below the usable signal bandwidth, and also demands for the baseband clock generation stages a relatively acquisition range of the order of $\pm 10\%$.

3.3.2 Transmitter section and link budget

We have seen above that due to the timing estimation problem, low-rate ultra-wideband transmitters are inherently more energy efficient than the companion receivers. Since the system under design is peer-to-peer, we take advantage of this fact by increasing the transmitter peak output power as much as possible. Spectral mask limitations impose a power spectral density of -41.6dBm/MHz over a 500MHz bandwidth, corresponding (for brickwall filtering) to an average power of -14.6dBm . For a 1Mbps data rate and 2nS pulse, the maximum output power imposed by this FCC limit is $+14\text{dBm}$. A lower peak output power of 6dBm is chosen in this work to allow for some margin, as well as to avoid requiring a complex PA architecture. Since 6dBm correspond to a voltage swing of 600mV on a $50\ \Omega$ load, this value can be realized without impedance transformation and using only thin-oxide transistors in the chosen 65nm CMOS technology, which has a nominal supply voltage V_{dd} of 1.2V . Under the expected worst-case channel loss of 30dB for a radiative channel, this corresponds to a receiver input voltage of approximately 20mV 0-pk, which will define the desired receiver sensitivity as -24dBm when referred to peak power. The resulting link margin calculation is summarized in table 3.1 A receiver with 44dB of noise figure will guarantee reception. We now try to understand how to build a receiver with such a high noise figure in the most power efficient manner.

Item	Value	Running Total
TX Output Power	6 dBm	+6dB
Path Loss	30 dB	-24 dB
Demodulation SNR(2-PPM)	16 dB	-40dB
Thermal Noise (2GHz BW)	-84 dBm	+44dB
Receiver Noise Figure	44 dB	0

Table 3.1. Receiver Link Margin Calculations

3.3.3 Receiver section

In recent years, the space of impulse-radio UWB receivers has been populated by several entries ([16],[27],[9],[37]). These receivers differ in the details at many levels, but can be broadly categorized in four categories based on the type of radio-frequency and baseband signal processing implemented. At radio frequency, we distinguish between heterodyning receivers ([16],[36]) which utilize a (free-running) local oscillator to downconvert the desired signal to baseband and self-mixing ([9],[27]) receivers in which the RF signal experiences amplification followed by self-synchronous rectification. At baseband, the distinction lies between analog ([27],[9] and [16]) and digital ([42]) correlation receivers.

Power efficient front-end selection

We now proceed to determine the most power efficient RF architecture for our system. In addition to the aforementioned self-mixing and heterodyning architectures (Fig.3.10, (a) and (b)), we exploit the high-SNR channel and explore an architecture that cannot be employed in radios detecting μV level signals. This architecture is based on direct AM-detection (without any pre-amplification, Fig.3.10,c)) and could be feasible in this case due to the small channel loss. In an overhead-limited regime, such architecture seems advantageous as it does not require any active device to have a transit frequency comparable to the carrier frequency, and furthermore, it does not require an RF oscillator. We start by comparing choices (a) and (b). The choice between these two designs options is straightforward. Consider in fact a receiver based on a single-stage, differential LC-loaded gain stage, feeding a passive-mixer based envelope detector (See left side figure 3.9), with tank tuned at frequency ω_0 . Rewiring the transistors, we can obtain the circuit to the right of figure 3.9, which consists of a differential oscillator directly driving the passive mixer gate terminals. Since in both cases, the output of the passive mixer is proportional to $V_g(t)V_s(t)$, we find that called A_v the gain of the differential amplifier, in the case of self-mixing we find $V_{out} = \alpha A_v^2 V_{rf}^2$, while in the case of heterodyning we have $V_{out} = \alpha V_{LO} V_{rf}$, so that as long as $V_{LO} \geq V_{rf} A_v^2$, the comparison is in favor of the heterodyning approach. Since for a current-limited design, we also have $V_{LO} = \frac{A_v}{gm_{eff}} \geq A_v n V_T$ ($gm_{eff} = \frac{gm}{I_d}$), we can rearrange terms to obtain:

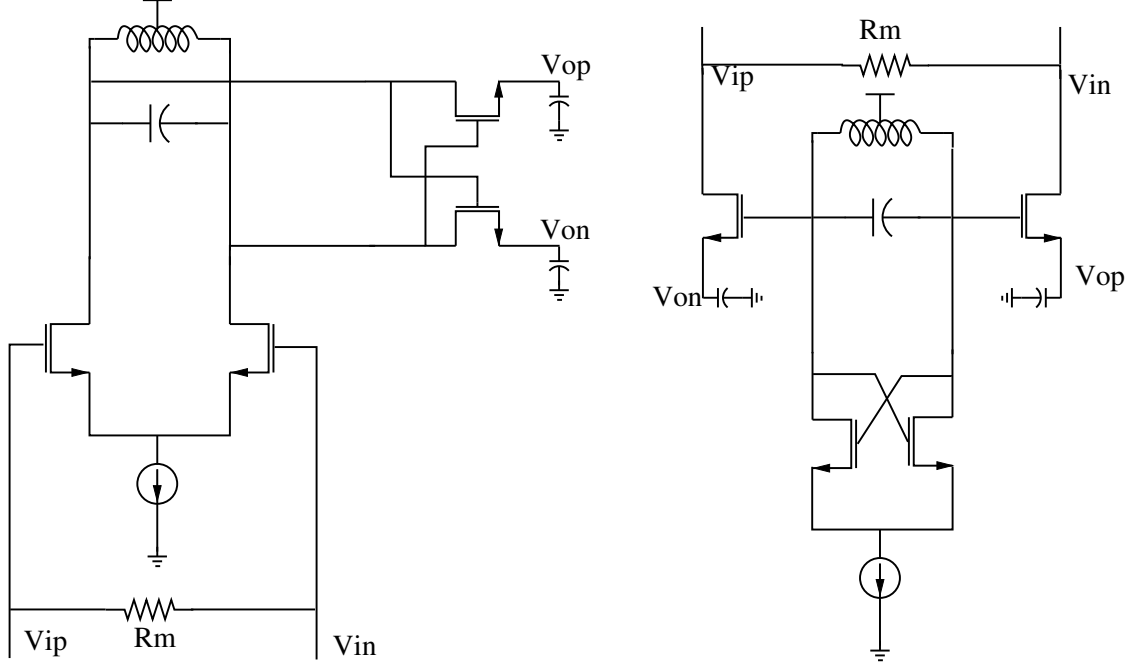


Figure 3.9. Comparison between a receiver comprising of LNA+AM detection (left), and lo power oscillator+passive mixer(right)

$$\frac{V_{out}^{SelfMixer}}{V_{out}^{ActiveLO}} \leq \frac{A_v V_{rf}}{n V_T} \quad (3.13)$$

Given that in most radios scenarios the ratio in 3.13 will be much smaller than 1, heterodyning receivers will almost invariably have higher gain for a given power consumption in the RF path ².

Determining if instead an architecture employing direct AM detection is advantageous over one using heterodyning is more difficult and requires a deeper understanding of noise effects in self-mixing receivers. As a first step, we review the noise-figure performance of a squaring detector. In a squaring detector ([56]), the downconversion process occurs as a result of the incoming RF signal being rectified by active device nonlinearity. A common model of the AM downconverter is that of a squaring nonlinearity followed by a low-pass filter (3.11). More accurate models can also be used ([52]), but this simple representation is often sufficient for gain calculations. In this context, the input of the AM detector will be composed of an amplitude modulated sinusoidal signal $b_k \sin(2\pi f_0 t)$, plus a noise process $n(t)$ which we assume has known autocorrelation function $n(t)$. The output of the AM detector under these conditions is given by

$$V_o(t) = g_2 \left(\frac{b_k^2}{2} + LPF[2 \sin(2\pi f_0 t) n(t) + n^2(t)] \right) \approx g_2 \left(\frac{b_k^2}{2} + N_{BB,Q}(t) b_k + N_{env}(t)^2 \right) \quad (3.14)$$

²The reader might be concerned that we neglected thermal noise in this analysis. However, since both receiver topologies exhibit loss, their noise figure is really set by the thermal noise of the following baseband stage and by the total input-output gain. As a result, the architecture with higher gain will also have lower noise factor.

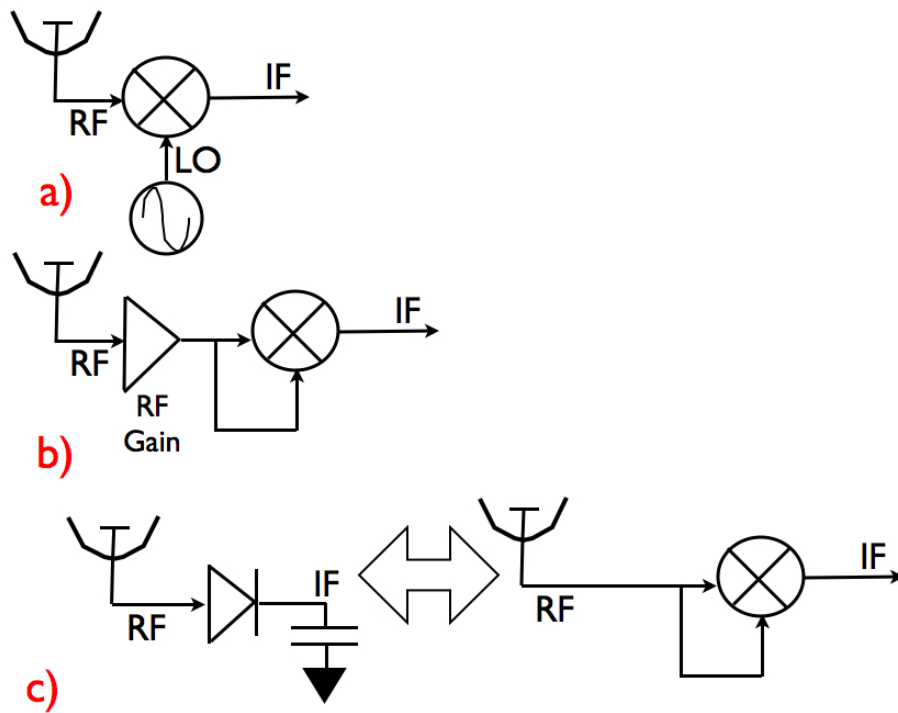


Figure 3.10. Low Power Radio Architectures: a) Heterodyning b) RF Gain plus AM-Detection c) Direct AM Detection




Figure 3.11. Behavioral model of AM detector




Figure 3.12. Operation of a single transistor MOS AM Detector

Where g_2 is the gain of the memoryless nonlinearity $y = g_2 x^2$, and $\text{LPF}[\cdot]$ denotes low-pass filtering operation. Furthermore, we take the incoming signal as reference in the constellation plane, and write $n(t) = N_{bb}^Q \sin(2\pi f_0 t) + N_{bb}^I \cos(2\pi f_0 t) = N_{env}(t) \cos(2\pi f_0 t + \phi)$. Most often (but not always, see [56]), the term in N_{env}^2 is small and can be neglected. In this case, the output noise process is still gaussian, and the the output SNR can be written as

$$SNR_o = \frac{\frac{b_k^4}{4}}{\frac{N_0 b_k^2 BW}{2}} = \frac{b_k^2}{2N_0 BW} = \frac{SNR_{in}}{2}$$

So that the ideal noise factor of an AM-detector is 3dB. In a real implementation however, several other phenomena combine to further degrade performance. Most notably, the active device in the AM detector will run a DC bias current I_{dc} , and consequently generate noise. For an MOS detector biased in the active region, the rectified variable is usually a current, so that current noise is the quantity of interest. As a result, the detector adds a noise $i_n^2 = 4KTn_F g_m$ per root Hertz. In several cases ([39]) this is the dominant source of noise at the slicer input. In the presence of this additional contribution, the output SNR is

$$SNR_o = \frac{g_2^2 b_k^4}{4\left(\frac{g_2^2 N_0 b_k^2}{2} + 4KTg_m n_F\right)BW} \quad (3.15)$$

For the relevant case of a transistor biased in sub-threshold, $g_2 = \frac{g_m}{2nV_T}$, and equation 3.15 becomes:

$$SNR_o = \frac{g_m^2 b_k^4}{2BW(g_m^2 N_0 b_k^2 + 32KTg_m n_F n^2 V_T^2)} = \frac{g_m b_k^4}{2BW(g_m N_0 b_k^2 + 32KTn_F n^2 V_T^2)}$$

Put $N_0 = \frac{4KTn_{IN}}{gm}$, the final result

$$SNR_o = \frac{g_m^{Diode} b_k^4}{8KTBW(n_{IN}b_k^2 + 8n_F n V_T^2)} \quad (3.16)$$

For a current commutating mixer, considering as output the downconverted current, we find

$$SNR_o^{Mixer} = \frac{b_k^2 g_m^{Mixer}}{2KTBW(n_F + n_{IN})} \quad (3.17)$$

In order for the output SNR to be the same, we find

$$\frac{gm_{Diode}}{gm_{Mixer}} = \frac{4(b_k^2 n_{IN} + 8n_F n^2 V_T^2)}{b_k^2 (n_F + n_{IN})}$$

Since typically $b_k \leq nV_T$, this equation indicates that for an equal bias current (equal g_m) the noise performance of the AM detector will be worse than that of the mixer counterpart. Alternatively, to achieve the same noise performance, the mixer transconductor will be biased at much lower current. How can the AM detector ever be power efficient then? Since a mixer needs an LO generator and drivers, we need to account for this power consumption as well in our comparison. By doing so, and called P_{LO} a constant summarizing the power consumption of LO generation and distribution, we can determine that the value of the input voltage V_{in}^* such that for $V_{in} \geq V_{in}^*$, $P_{diode} \leq P_{mixer}$ equals

$$V_{in}^* = \sqrt{\frac{\gamma(3n_{IN} - n_F) + \sqrt{\gamma^2(n_F - 3n_{IN})^2 + 128n_F\gamma P_{LO}(nV_T)^2}}{2P_{LO}}} \quad (3.18)$$

where $\gamma = \frac{2KT^2}{q} BW \cdot SNR \cdot V_{DD}$. Values of equation 3.18 as a function of bandwidth for different values of P_{LO} , are shown in figure 3.13. If we assume that the cost of local oscillator and drivers is $100\mu W$, the AM-detector begins to be advantageous in term of power dissipation for input voltages larger than 14mV, which lines up with our expected voltage range and confirms that the AM-detector radio is a good idea in the context of this application.

3.4 Conclusions

The analysis presented in this chapter led to the conclusion that circuit sizing must be combined with architectural modifications in order to exploit the reduction in path loss resulting from the reduced range and obtain a significant reduction in power. We have shown that by using impulse radio signaling and transmitter and receiver duty-cycling, the overhead costs due to the RF VCO could be reduced by more than ten times.

However, our analysis also showed that for the expected signal amplitude, a duty-cycled *Direct-AM Detection* receiver, where rectification is preceded by no active RF gain, can be even more energy-efficient than a heterodyning one. Adopting this architecture poses




Figure 3.13. Minimum Input Amplitude s.t. AM-Detector is more power efficient solution

significant serious challenges, caused by its inherent interference sensitivity. In the next chapter, we analyze the effect of interference in self-mixing receivers, and devise techniques to reduce the impact of narrow-band blockers on system performance without requiring RF filtering. By combining these techniques with the chosen AM-detector, we obtain the robust, ultra-low-power transceiver we are looking for.

Chapter 4

Baseband and Timer Subsystem Design

In the previous chapter, we focused on techniques to reduce power dissipation. We found that a tradeoff exists between interference robustness and power: to minimize energy, we intend to use a ultra-wideband communication system and a direct AM-detection receiver. Both these techniques are known to increase the chance that system performance is degraded by unwanted signals. As a result, we devote the bulk of this chapter to exploring techniques to mitigate interference. We propose a system that changes its RF front-end topology from being AM-detection to heterodyning when close in-interference is present. The system operates most of the time in diode mode to save power, and switches to heterodyning by monitoring the bit-error rate (B.E.R.) in real time. We also describe how the B.E.R. monitoring function is embedded into the timing recovery subsystem (also described in this chapter), and discuss what are the implications of this choice on the overall specifications of the baseband subsystem.

4.1 Baseband architectures for UWB systems

Ultra-wideband receivers have been the subject of relatively intense academic research in the last ten years. While it is generally recognized [42] that for these systems, interference is a bigger concern than white noise, the debate on baseband architectures mostly focuses correlation receivers([37]) , which are only optimal from the standpoint of thermal noise, and the discussion is centered on whether to implement the correlation function in the analog or the digital domain. For example, for rectangular pulses the correlating baseband consists of an integrate and dump matched filter, and because it is easy to implement, it is proposed




Figure 4.1. Receiver architecture with linear downconversion and analog correlation

also in situations in which it is only an approximation of a matched filter.

The main drawback of this approach is, as described below, its lack of robustness to interference. Since the system under design emphasizes the relative importance of interference versus thermal noise, after assessing the limitations of the integrating baseband we proceed to search alternative baseband techniques that are better able to reject interference in a wideband system.

Interference robustness of analog-correlating architectures

Consider the receiver topology shown in figure 4.1, that realizes an analog correlating coherent receiver for OOK/PPM. The output of a mixer is fed to an integrator with integrating window of duration T_p . Since the integrator transfer function is $T_I \frac{\sin(\pi f T_I)}{\pi f T_I}$, the equivalent response of the receiver reported to RF is shown in figure 4.2. While higher order baseband filtering can improve the resiliency to interfeferers in far our-band, signals lying in a bandwidth of at least $\pm \frac{1}{T_p}$ of the desired carrier frequency are harmful. Increasing the integration time T_p beyond the duration of the desired signal pulse only makes the performance for close-in disturbances (4.2) worse, and it increases the noise floor as well. The situation worsens significantly if the linear mixer is replaced by a self-mixer. In this case, in the absence of RF prefiltering, any interfering signal is downconverted to DC and integrated, resulting in a false alarms. This leads to the conclusion that in order to improve the interference performance of UWB systems, simple integration at baseband has to be abandoned, and other signal processing techniques must be investigated.




Figure 4.2. Conversion gain for receivers with different baseband integration times

4.2 Improving Interferer Performance in Self-Mixing Radios

We start by analyzing the *Selectivity* of an AM radio. An intuitive argument suggests that since the detector is only concerned with energy in the channel, interferer rejection is only determined by the RF filter preceding it. In reality, this is correct only if the signal-processing performed at baseband is integrating, in which case, as shown in the previous section, false-alarms will limit the performance. Different signal processing strategies at baseband lead however to different results. In order to analyze this problem, we need two introductory remarks. The first regards the nature of errors made by the detector, assuming ideal signal processing blocks follow the diode. The possible hypotheses are presence or absence of a desired signal. The error corresponding to the event that a pulse is present, but is not detected, is here referred to as *False Negative*. The dual error, occurring when our decision procedure declares the presence of a pulse even when no pulse is present at the input of the receiver, will be called a *False Positive*.

We also introduce an intuitive model of the envelope detector in the phasor domain. According to this model (See Fig.4.3), if we represent the inputs to the detector itself in the phasor domain, the detector output is proportional to the length of the vector sum of the inputs. This model will prove useful in understanding the analysis in an intuitive way.




Figure 4.3. AM Detection as phasor domain length-measurement

False Positive Performance

Consider now the case of interferer-induced false-positives. In this circumstance, a large undesired narrowband signal is present at the input of the receiver. By our analogy, this interferer can be represented with a phasor rotating in the complex plane at a speed ω_I , and with length $V_I(t)$. The output of the detector is therefore related to the quantity $V_I(t)$. In particular, reverting to a squaring nonlinearity, it will be $\frac{g_2 V_I^2}{2}$. Now, if we assume the signal processing following the detector consists only of linear gain followed by optimal thresholding based on the signal statistics (i.e. decision value V_{th} is midway between "1" and "0"), this output voltage will be classified as a 1, as soon as $g_2 V_I^2 \geq \frac{g_2 V_{sig}^2}{2} \approx V_I \geq \frac{V_{sig}}{\sqrt{2}}$ leading to false positives. Under this assumption, the error rate of the detector is given by the red line in figure 4.4. Due to the slowly-varying nature of $V_I(t)$ however, these errors could be greatly reduced by high-pass filtering the baseband signal. Since this results in energy-loss in the signal itself as well, a tradeoff is incurred. For example, for a simple first order filter with

$$H(\omega) = \frac{j\omega}{j\omega + \omega_p}$$

the signal-to-interferer power ratio at the output of the filter can be expressed as in equation 4.1

$$SIR_o = \frac{P_S BW_I^{(N)} (BW_s^{(N)} - \text{atan}(BW_s^{(N)}))}{P_I BW_s^{(N)} (BW_I^{(N)} - \text{atan}(BW_I^{(N)}))} \quad (4.1)$$

Where $BW_I^{(N)}$, $BW_s^{(N)}$ are the bandwidth of the signal and the interferer squared envelopes normalized to ω_p . The situation is favorable when $BW_s^{(N)} \gg 1 \geq BW_I^{(N)}$ which explains why this strategy is only really advantageous in a wideband system. For instance, if $BW_s = 100\text{KHz}$, choosing $BW_s^{(N)} = 10$, results in an effective suppression only for interferers with at most 10KHz bandwidth¹. In a UWB system with 1GHz BW on the other hand, the same conditions correspond to an attenuation of interferers with a bandwidth as high as

¹Not accounting spectral regrowth due to squaring




Figure 4.4. False Positives Rate for system employing high pass filter versus conventional case

100MHz on the envelope, which encompasses most existing wireless systems. This is shown in figure 4.4, where the computed False Positive rate is reported for $\frac{P_s}{P_I} = 1$ and varying envelope bandwidths, assuming a signal-to-thermal noise ratio of 16dB. So even without any prefiltering, in a ultra-wideband system interference from most standards can be rejected through high-pass filtering at baseband. This encouraging result represents unfortunately only half of the picture, as we'll see next.

False Negative Performance

We now consider the case of simultaneous presence of desired signal and a narrowband interferer and try to analyze the probability that the detector rules for a *No signal* condition. Let us start our analysis using an intuitive argument and our phasor-domain analogy: we represent the interferer and the signal as phasors rotating in the complex plane. The output of the AM detector is proportional to the squared length of the vector sum of the two phasor and depends on the relative phase of the two vectors. If we can treat the signal as an infinitely-narrow pulse (or a delta function), the AM detector output will consist of an equally narrow delta function, superimposed to a baseline voltage due to the slowly varying length of the interferer. If we employ high-pass filtering at baseband as discussed above, the baseline voltage will be rejected, and we can focus on the remaining contribution due to the pulse. The variable of interest is therefore

$$V_o = \frac{g_2 V_{sig}^2}{2} \left(1 + 2 \frac{V_{int}}{V_{sig}} \cos(\Theta) \right)$$

Where Θ is the unknown phase offset between pulse and interferer carriers. In the absence of noise, whenever $\Theta = \theta^* - \text{acos}(\frac{V_{sig}}{2V_{int}})$, $V_o = 0$ so that a false negative is obtained. A graphical interpretation of this special angle value in the phasor domain is shown in figure 4.5. Since we are looking for the initial phase such the length of the vector sum between signal and the interference is the same as that of the interferer alone, this can be found by looking for the intercept between a circle of radius $|V_{pulse}|$ and center \vec{I} , indicating the possible locations of the vector sum of desired signal and interferer and a circle of center 0 and radius $|V_{int}|$, which is spanned by the interferer phasor in the absence of any signal.

In reality, false-negative rate will degrade even when this condition does not hold exactly as a result of the decreased SINR. To calculate the receiver B.E.R. , we can then use Bayes' Formula and average over Θ the conditional probability of error for a given realization of θ . Since this interference mechanism dependence on a phase offset variable, which is uniformly distributed on $[-\pi, \pi]$, error rate is however dominated by the relatively frequent worst case conditions, i.e. values of Θ in a neighborhood of θ^* . In fact, the effect of interference in this case maybe thought of as multiplicative noise, similar to the case of a fading channel, which explains the large SNR penalty.

While this analysis captures the essence of the false-negative mechanism, the assumption made in modeling the baseband pulse as a delta function introduces errors, as it suggests that the false-negative rate is independent of the relative frequency of desired signal and interference.

This is incorrect, as if interferer and signal have different frequencies and the baseband pulse has finite duration, the relative phase of desired signal and interferer will drift over time, resulting in a different amplitude for the vector sum and hence a non-zero output. The rate of change of the vector sum depends on the frequency offset between interferer and signal. To model this mechanism correctly, we need to make an assumption about the behavior of the receiver stages following the envelope detector. To seek an approximate expression, we ignore finite sampling time at the detector output and consider a receiver that continuously looks at the output of the front-end, and declares a pulse received whenever the voltage threshold V_{th} is exceeded. We also approximate the real B.E.R. by the probability of error obtained for the worst case initial phase θ^* . Called ΔF the frequency offset between interferer and desired signal carrier, and T_p the pulse duration we find that

$$\begin{aligned}\theta^* &= -\text{acos}(\frac{V_{sig}}{2V_{int}}) \pm \pi \mid \Delta F T_p \mid \text{ if } \frac{V_{sig}}{2V_{int}} \leq 1 \\ \theta^* &= \pi \pm \pi \mid \Delta F T_p \mid \text{ otherwise}\end{aligned}$$

This choice of initial phase is least favorable as, for the given ΔF , it forces the output of the envelope detector to be small for the longest amount of time. In this case, the error-rate performance will be determined by the probability of detecting either of the two peaks in the output of the the envelope detector. The low frequency component of the envelope detector output (i.e. ignoring the second harmonic terms) is shown in the botton right of figure 4.5. In a real system, the abrupt transitions at the beginning and end of the pulse would be smoothed by the time constant at the output node of the detector. For each peak,

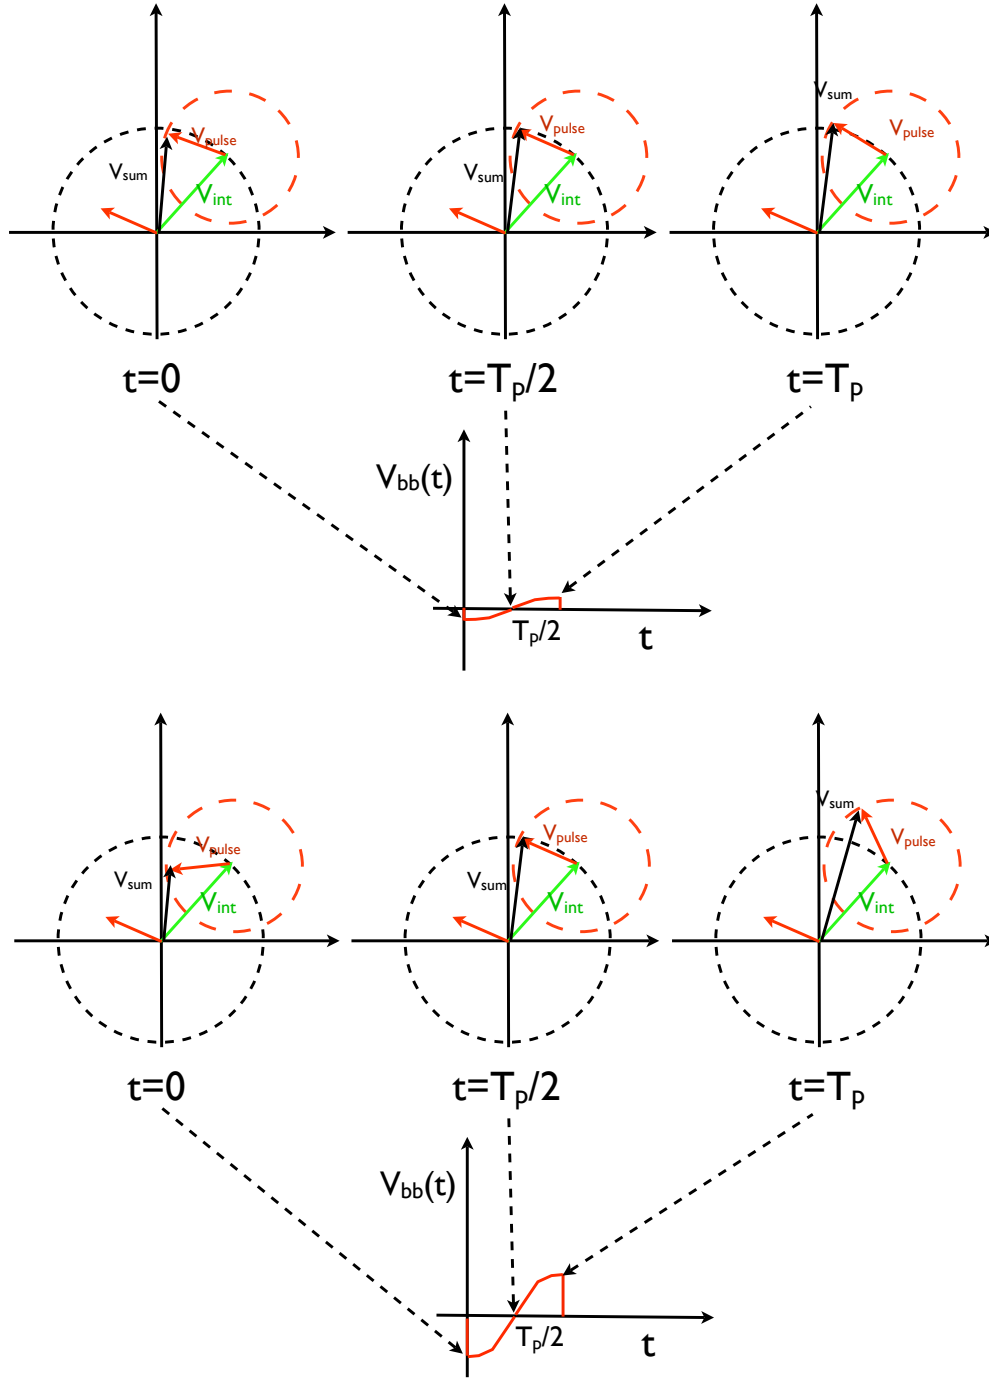


Figure 4.5. Graphical interpretation of the worst case angle θ in the case of $\Delta F \ll \frac{1}{T_p}$ (top), and of baseband pulse of finite length T_p (Bottom). The baseband output is shown for $\Delta f T_p = \frac{1}{2}$




Figure 4.6. Complete Baseband architecture for false-negative rate calculation




Figure 4.7. Bit-Error Rate as a function of power for different frequency offset values

the probability of missed-detection evaluates to

$$P_{md} = Q\left(\frac{g_2 \mid (V_{sig}^2 - 2V_I V_{sig} \sin(\theta^* \pm (\pi \Delta F T_p))) \mid -V_{th}}{\sigma}\right)$$

. The dependence of P_{md} on the frequency offset ΔF and the interferer level V_I is shown in figure 4.7. Consider first the case of small ΔF (i.e. $\Delta F T_p \ll 1$). In this case, when the interference amplitude is increased from 0, the performance degrades constantly until $V_I = \frac{V_{sig}}{2}$ (Figure 4.7, top), and then remains constant.

When $\Delta F T_p \approx 1$ on the other hand, the receiver false negative rate first increases with increasing interferer amplitude, and then it decreases sharply again (See Fig. 4.7). The worst case (maximum error rate) occurs when $P_I = \frac{P_s}{2}$.

This surprising phenomenon can be explained by modeling the operation of the squaring detector as multiplying the incoming signal by itself. In the presence of two tones, this operation produces a DC shift term for every tone, as well as an intermodulation product. When only the desired signal is present, only the DC shift term due to it is detected by the following circuitry.

An interferer at large frequency offset and smaller than the signal creates an intermodulation product smaller than the DC shift itself ($V_{sig} \geq V_{int}$) that effectively reduces the noise margin on the DC shift detection. If the interferer is much larger than the signal on the other hand, the intermodulation product is larger than the DC bias shift, and becomes the quantity detected by the baseband circuitry. Clearly, the interferer needs to be at a sufficiently large frequency offset that its arbitrary phase relationship with the signal carrier does not statistically reduce the signal gain. In this circumstance, the diode downconverter is behaving as a *Single Gate Mixer* [54], where the large interferer is acting as (the wirelessly coupled) LO signal.

We can under these conditions characterize the false negative performance of this diode receiver with a *Worst Case* BER versus Interferer offset curve, corresponding to the case of $V_{sig} = 2V_I$. The calculated worst-case interference sensitivity curve is shown in figure 4.8 for a pulse duration value $T_p = 1.6nS$. The resulting interferer-sensitive bandwidth(Defined as the offset from the signal such that the BER is larger than 10^{-3})for this model is $\pm 350MHz$, corresponding to $\Delta F T_p = .42$.

Figure 4.9 shows the results obtained when repeating the same experiment in a time-domain *Simulink* behavioral model of the system, which are in good agreement with theory, considered the many approximations that were made in deriving the results in Fig.4.8. We therefore see that the combination of an envelope-detector and a high-pass baseband filter displays *Selectivity* even when not preceded by an RF prefilter. *Far-out Interferers* are rejected by the receivers, while *Close-In* interferers produce false-negative errors. As a result, any interference mitigation technique we develop for this application should focus on improving false-negative errors introduced by narrowband jammers .

Coherent Radio Performance

The false-negative performance of the AM-detector radio directly depends on the phase between the incoming signal carrier and the interferer carrier. These two are unknown at the




Figure 4.8. Bit-Error Rate as a function of interferer offset of SIR=6dB




Figure 4.9. Simulated error of the receiver in diode mode for pulse durations of 2nS (blue) and 8nS (red). SIR=6dB

receiver side and mutually independent; furthermore the arrival of the desired signal in time is also not known. These effects combine to make it difficult to improve interferer resiliency while using an AM detector architecture. This situation can be improved instead by using a downconverter that employs a local oscillator (See 4.10). Notice that we are not interested in improving the interference performance of an AM radio for arbitrary relative location of signal and interference, but only for interfering signals that lie within $\Delta F = \pm \frac{1}{T_p}$ of the desired pulse carrier in the frequency domain. For a hard-switching mixer with perfectly linear RF port and phase aligned to that of the desired signal, the baseband output will be

$$V_{if}(t) = A_v \frac{2}{\pi} (V_{sig}(t) + V_I(t) \cos(2\pi\Delta F t))$$

Where we assume $\Delta F T_p \ll 1$. Compared to the case of AM-Detection, the intermodulation term has disappeared thanks to the presence of the LO signal, which enables to decouple the RF port linearity from the conversion gain. We can now observe that under the hypotheses made, the term $V_I(t) \cos(2\pi\Delta F t)$ is slowly varying compared to the signal term, and as a result it can be removed by high-pass filtering without impacting the receiver SNR. In particular, using a discrete-time first order high-pass filter with transfer function $1 - z^{-1}$, the baseband discrete time process become

$$\begin{aligned} V_{if}[k] &\approx A_v \frac{2}{\pi} ((V_{sig}[k] - V_{sig}[k-1]) + V_I(\cos(2\pi\Delta F T_s k) - \cos(2\pi\Delta F T_s (k-1))) \\ &= A_v \frac{2}{\pi} ((V_{sig}[k] - V_{sig}[k-1]) + \frac{V_I}{2} \sin(\pi\Delta F T_s) \sin(\pi\Delta F T_s (2K-1))) \end{aligned} \quad (4.2)$$

The interferer is suppressed by a factor $\sin(\pi\Delta F T_s)$, which can be made small by reducing T_s . For $T_s = 400pS$ and $|\Delta F| \leq 350MHz$, the worst case SNR improvement is 3dB. To further maximize the performance in this condition, we can choose the baseband transmitted waveform to match the interferer-rejecting $1 - z^{-1}$ filter, so that this acts as a matched filter. This choice of baseband shape does not impact diode-mode reception, as in this case the phase modulation is rectified and does not impact the output. Since this receiver topology is now linear, interference shows up as added noise, and it has the same effect on both false-positive and false negatives so that both these errors will increase for out-of-band interference, and be small for in band interference as shown in figure 4.11. Notice that the dependence of errors on frequency offset appears to be complementary to that of the AM-detector (i.e. far-out interferers are more detrimental to the performance), a fact that we will use to our advantage.

Proposed Dual Mode Radio Receiver

It has already been mentioned that the AM-detector and the linear mixer perform in a complementary way with respect to interference. This is shown in figure 4.12, where the overlaid false-negative error characteristics of the two receivers are reported. Imagine now a signature of the presence and location of interference was available. We could then configure the receiver as an AM detector in default conditions, and switch to a linear downconverter when close-in-interferers are detected (See Fig.4.12,top). As long close-in interferers are only




Figure 4.10. Receiver with linear down converter and high-pass filtering at baseband




Figure 4.11. Bit-error rate performance in a linear downconverter as a function of interferer frequency offset

rarely present, and the energy and complexity costs of the interference-monitoring function are small, this guarantees robust operation with a small power price.





Figure 4.12. Proposed dual-mode receiver with real-time quality of service control(top) and overlaid false negative rates response of linear and AM-detecting down-converters(bottom)

We claim that the receiver bit-error-rate in diode mode is a good statistical signature of the location and strength of an interferer. In fact, we have seen that the receiver false-negative rate will increase only in the presence of a close-in interferer, while the false-positive

rate will not change. In reality, a sudden increase in the receiver bit-error rate can indicate either a decrease in the signal level, or an increase in the noise and interference level. However typically, switching from an AM detector architecture to one based on a linear downconverter will both improve the linearity and increase the gain of signal chain, this solution will improve both circumstance. We therefore choose to utilize this indicator as a signature of interference.

4.3 Synchronizer and modulation scheme design

We have seen at the beginning of the section that the overhead power dissipation of a UWB system is largely determined by its synchronizer/timing generation subsystem which should therefore have ultra-low power consumption. In several UWB receivers ([16],[27],[42]) this subsystem is either off-chip or realized through a DLL slaving a low-frequency local oscillator to a quartz crystal. Due to our constraint on not using a crystal, this approach is not viable for us and we need to seek a different strategy. Similarly to [11] we chose to recover the clock from the incoming data stream, by co-designing modulation scheme and timing subsystem. Several choices for the timing-recovery scheme are available, largely based on mostly-analog bang-bang or linear phase or delay locked loops. All feedback clock recovery loops exhibit a similar tradeoff between jitter tolerance(demanding a small loop bandwidth) and acquisition speed (demanding a large loop bandwidth), as well as between transition density and synchronizer frequency accuracy. PLL systems are more robust to missed detection of incoming pulses, but suffer from long acquisition times due to the small loop bandwidth, and often consume a sizable amount of power due to the many analog components. For example in [11], the PLL consumes $72\mu A$ when running at 250Kbps, of which almost half is spent in the low-frequency loop filter. In addition, a linear PLL can be too sensitive to input jitter, and as a result, bang-bang-operation could be preferred, leading to even lower bandwidth and longer acquisition. DLL systems on the other hand achieve broader loop bandwidth and hence faster acquisition, but allow less filtering of the incoming reference edge.

A clock recovery technique that mitigates both these tradeoff was introduced in [6], and is based on an oversampling clock-recovery architecture. The oversampling clock-recovery system continuously runs, and then XORs consecutive decisions to determine the position of the transitions. This architecture lends itself to low-power operation because it can be implemented using exclusively digital circuits, and relies on transistor speed which is more readily available than analog accuracy in deeply scaled technologies([50]).

Adopted Clock Recovery Technique and modulation scheme

The architecture used in this work is inspired by the oversampling clock recovery system, as well as by classical communication systems employing a numerically controlled oscillator. Numerically controlled oscillators (NCOs) have the advantage of ultra-wide tuning range with frequency resolution and low output jitter [15]. They are completely digital circuits,

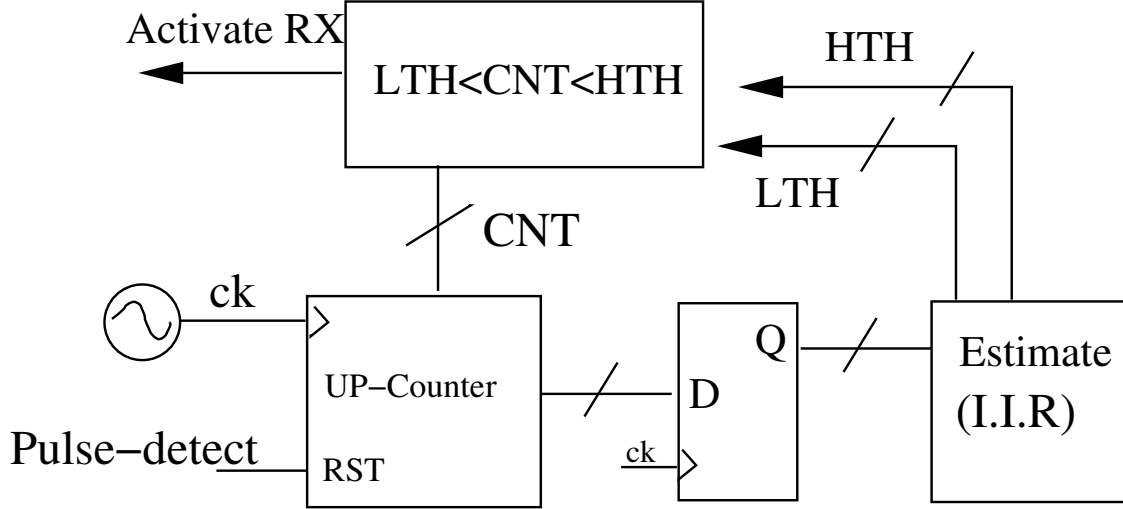


Figure 4.13. Conceptual Synchronizer Diagram

and as a result their power efficiency has improved steadily over the past 20 years thanks to Moore's law. The NCO proposed in [15] for example, consumed 500mW while running at 700MHz. We will show in this work that this figure can be reduced by an astonishing 4 orders of magnitude, for essentially the same clock rate. A conceptual system architecture is shown in figure 4.13 . A high-speed free-running oscillator sets the system time base and feeds an up-counter that generates a phase-domain ramp. The same signal is also used to as clock input to the analog receiver's slicers. When a pulse is detected by the radio front-end, the count value is sampled and the counter reset. Count values after the first are therefore a measurement of the TX clock frequency normalized to the RX time base. These counts are averaged by an IIR filter forming an estimate of the TX period \hat{T} , and control a duty-cycle controller, which turns on and off the analog receiver when the count respectively reaches $\hat{T} - \delta$ and $\hat{T} + \delta$. The synchronizer relies on two feedforward mechanisms to align the phase of the recovered clock to that of the desired signal: first, the reset operation forces the edge of the recovered clock to be aligned to that of the incoming signal. Second the averaging loop enforces the frequency to be the same to within a oscillator clock cycle T_{ck} .

The speed at which the local oscillator runs and the bandwidth of the loop filter determine the ultimate resolution achieved by the timer in producing the estimate \hat{T} . Defined $\epsilon = \hat{T} - T$, ϵ and the duty-cycling window width 2δ are related by the maximum run length RL of the input without transitions: $+T_m + T_{pulse} + \epsilon RL \geq 2\delta$. A tradeoff therefore exists: a short code run length translates into a short duty cycling window, enabling a coarse measurement and ultimately reducing power consumption in the synchronizer but likely increasing the power in the transmission. Conversely, a smaller ϵ allows both a small duty-cycling interval and a longer run-length to be tolerated, reducing the transmission power at the expense of the synchronizer power. Since at low pulse-repetition rates power dissipation is limited by overhead, it makes sense to trade-off power consumption in the data-detecting receiver in exchange for a simplification of the synchronizer.

The modulation scheme chosen in this work is therefore differential PPM with a time deviation equal to half a bit-period. The full modem diagram is shown in figure 4.14 .The

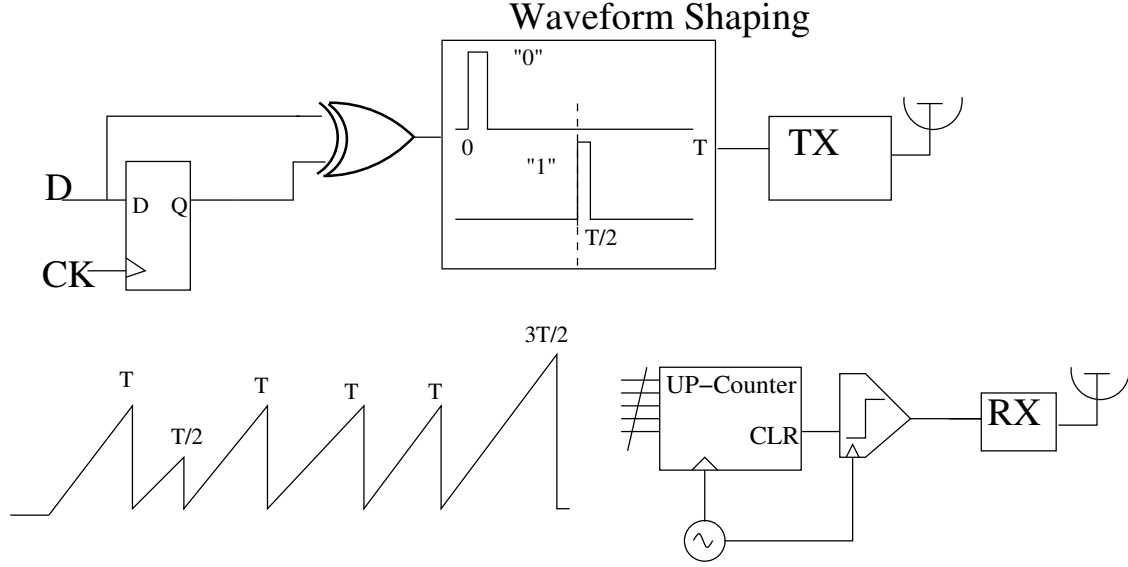


Figure 4.14. Block Diagram of Modem employing differential-PPM modulation and merged synchronizer/demodulator

bitstream is first digitally differentiated, and then encoded by the waveforms shown in figure 4.14 and sent over the air. This scheme requires only non-coherent RF detection, which is compatible with an energy-detection front-end. It also allows to merge the demodulator and the synchronizer in a single structure, saving power consumption and area . This occurs because the synchronizer phase ramp is reset at every pulse detection, and hence its value before reset is a measure of the time distance between consecutive received pulses. This value can be used to differentially demodulate the data, as shown in the modulation trellis diagram shown in Fig. 4.15. The maximum code run length can also be inferred from the trellis. Since the modulation scheme has a maximum run length of $1.5T_b$, the minimum turn-on time width becomes $2\delta = T_m + T_{pulse} + 1.5T_{ck}$. Since typically $T_m \geq T_{pulse} \geq T_{ck}$, this choice of modulation scheme allows the use of a short reception window, while keeping the synchronizer completely digital and hence maximally simple. Finally, this scheme is very tolerant of baseband clock imperfections. By using differential demodulation, ([41]) the detection process is insensitive to long-term oscillator drifts. While the effect of cycle-to-cycle jitter is still present, the large large time deviation used in the modulation scheme itself make it unimportant, enabling a low-quality, low power ring oscillator to be used to synthesize the baseband clock.

Jitter transfer and bandwidth analysis

We now compare the jitter-transfer acquisition bandwidth trade-off of this architecture with that of a clock recovery phase-locked-loop, showing that the two architectures serve different purposes, and suggesting why the chosen clocking scheme is more suitable to this




Figure 4.15. Trellis diagram for differentially demodulated PPM

system . In the PLL case, the jitter transfer function is given by

$$JT(s) = \frac{K_{vco,eff}H(s)}{s + K_{vco,eff}H(s)} \quad (4.3)$$

Where $K_{vco,eff}$ is the VCO effective gain (including the divider), and $H(s)$ is the loop filter transfer function. $JT(s)$ is a third order low-pass transfer function with a low-frequency zero and poles located around the PLL loop bandwidth [28]. The cycle-to-cycle jitter can be calculated by integrating $JT(s)$ times the input phase noise spectrum and taking into account the sinc() [35] mask used for cycle to cycle jitter measurements. The result is that jitter from the reference is largely suppressed, as it is suppressed by $JT(s)$ past its crossover frequency, and by the sinc() mask in band. This indicates that from a noise point of view, we could use a large loop bandwidth, which is desirable for fast acquisition. However, a loop bandwidth larger than $\frac{1}{10}$ of the reference frequency itself results in PLL loop instability, and so cannot be used in practice. The problem with the phase-locked loop approach is the double integration loop that enforces exact matching of the recovered clock phase and frequency with incoming data-stream. While in an applications using long-run codes and high-data rate this is vital, the modulation scheme chosen is tolerant of long-term clock stability, making the optimal clock solution one that can deliver moderate short term stability, but high acquisition speed. This is achieved by proposed structure which similarly to a DLL, relies on feed-forward phase correction enhancing the stability of the loop. This enables extending the loop bandwidth beyond $\frac{1}{10}$ of the reference frequency, improving convergence speed. In fact, the jitter transfer function of the proposed architecture is readily derived by observing that the loop generates an edge occurring at time $E(n)$ that is located $(1 - \alpha) \sum_k \alpha^k (E_i(n - k) - E_i(n - k - 1))$ from




Figure 4.16. Effect of jitter on digital timing measurement

the last received one. As a result

$$JT(Z) = z^{-1} \left(1 + \frac{(1 - \alpha)(z - 1)}{1 - \alpha z} \right) = z^{-1} \frac{(2 - \alpha - z^{-1})}{1 - \alpha z^{-1}} \quad (4.4)$$

Which is the same transfer function of a delay-locked loop. A more interesting parameter in this system is the variance of the estimation error, i.e. the difference between the estimated arrival time of the next edge and the last one. This quantity is related to the input by $H(z) = JT(z) - 1$ and equals

$$H(z) = \frac{(1 - \alpha)(1 - z^{-1})}{1 - \alpha z^{-1}} \quad (4.5)$$

. Which shows that the loop tracks the low-frequency content of the input clock variations. The jitter peaking exhibited by Eqns. 4.3 and 4.4 is a concern when many nodes are cascaded and forward clock from one to the next[28],[64]. Every node amplifies the incoming jitter and adds its own. This leads to unbounded increase in jitter along the chain, and synchronization failures. This issue can be mitigated by locally regenerating the timing signal prior to forwarding, similar to the regeneration-forwarding strategy of communication networks. Regeneration of timing is obtained in this system by performing a *Digital* timing measurement. Consider the scenario in figure 4.16, where the incoming reference edge is shown shaded by the jitter pdf. The quantized signal will be the same as long as the jitter amplitude is smaller than the distance of the nominal to the decision edge, and the input jitter is suppressed. Calculations in appendix B show that as long as the $\frac{J_{in}}{T_{ck}} \leq 0.25$, the output jitter variance increases by less than 10% over the baseline given by quantization. This jitter-suppression property enables the use of this architecture in a dense-network environment where distributed synchronization is required.

Operation under presumed-error free decoding and data-dependent jitter removal

Similarly to a DLL the proposed clock-recovery technique relies on an incoming edge generated by a detected pulse to synchronize. However, even if the receiver is operating error-free, edges will not occur during *every receiver on-interval*. By considering the trellis in Fig.4.15, we see that if the receiver is in state "1", values of inter-pulse arrival times of either T_b or $T_b/2$ are allowed. Similarly if the receiver is in the "0" state, values of inter-arrival time of T_b or $\frac{3T_b}{2}$ can occur. Since the incoming data is not known at the receiver, this has to turn on and try to detect a pulse in both intervals. On average, a successful pulse detection will only occur for half the receiver "wake-ups". The randomness in the edge arrival time can be seen as a large amount of data-dependent jitter. In fact, it is easy to prove that for independent, equally likely, binary input data and jitter-free transmitter clock, the variance of the sampled counter output equals $\sigma_J^2 = \frac{T_b^2}{8}$. This data-dependent jitter limits the width of the duty-cycling window, resulting in increased power dissipation, and must be mitigated. A straightforward approach to reduce this jitter is to decrease the loop bandwidth, i.e. increase the amount of averaging performed. This is however impractical at low data rates, as it would lead to average a number of samples proportional to $(\frac{T_b}{T_p})^2 \approx 25000$, completely jeopardizing the tracking ability of the loop and increasing unacceptably the acquisition time. The idea pursued in this work is that of using *decision-directed averaging*. The idea behind decision-directed averaging is that since the amount of data-dependent jitter is much larger than random jitter due to thermal noise in the TX and RX clocks, the distribution of the measured inter-pulse intervals at the receiver will consist of three well defined peaks at $\frac{T_b}{2}; T_b$ and $\frac{3T_b}{2}$. Therefore, data-dependent jitter can be removed by multiplying the measured time interval with a constant that depends on the current decision. Called CNT_n the n_{th} measurement and X_n the input to the filter at time n , the update rule becomes

$$X_n = CNT_n \leftrightarrow \frac{3\hat{T}}{4} \leq CNT \leq \frac{5\hat{T}}{4} \quad (4.6)$$

$$X_n = 2CNT_n \leftrightarrow \frac{1\hat{T}}{4} \leq CNT \leq \frac{3\hat{T}}{4} \quad (4.7)$$

$$X_n = \frac{2CNT_n}{3} \leftrightarrow \frac{5\hat{T}}{4} \leq CNT \leq \frac{7\hat{T}}{4} \quad (4.8)$$

The effect of this update rule on the jitter of the recovered clock is shown in figure 4.17. A simpler approximation to this technique, which is actually implemented in this work, is that of using a *conditionally recirculating DLL*. Recirculating DLLs are already widely known in literature ([61]) as clock multipliers. The idea behind the conditionally recirculating DLL is to update the frequency estimate only when a valid clock-edge occurs. In other words, if we believe that our estimate is unbiased, and the receiver receives no-pulse upon turn-on, the synchronizer resets the counter at $T_{rst} = \hat{T} + \delta$, and keeps the estimate \hat{T} unchanged. In reality, this is not sufficient, as the turn-on-time of the receiver has now been shifted to the right by δ due to the missed pulse, leading to an enhanced probability of missing the next pulse. This shift can be corrected for by simply shifting the receiver turn-on thresholds to $\hat{T} - 2\delta, \hat{T}$ (See Fig.4.18). This also holds for the case in which an RF pulse was present

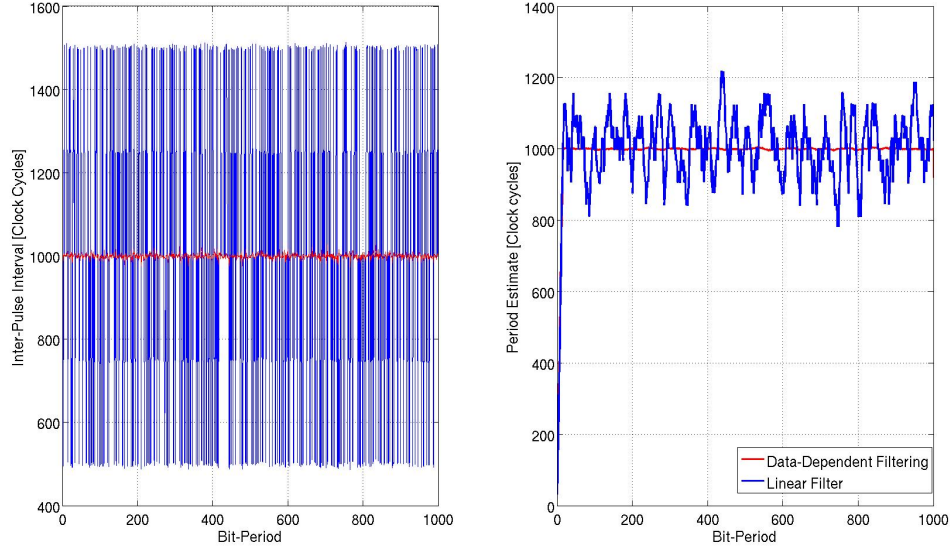


Figure 4.17. Period Estimate in the synchronizer with (red) and without (blue) data-dependent rescaling: input to the filter (left) and period estimate(right)

but went undetected by the analog front-end. Had N pulses gone undetected, the receiver should use $\hat{T} - (N + 1)\delta$, $\hat{T} - (N - 1)\delta$ as the $N+1$ th turn on window. However, the receiver should not miss any more than 2 pulses, as long as faithful data reception is occurring. If a large number of consecutive pulse reception intervals result in no reception, a synchronization failure is declared, and the system has to be restarted in synchronization mode. We therefore assume that this condition will be handled in higher-layers of the protocol stack, and ignore it in the hardware design.

Acquisition impact on analog receiver specifications

An important implication of recovering the clock from the decisions of the analog front-end is the interaction of circuit noise on the synchronization process. This is particularly important during acquisition, as in this time the receiver is constantly listening to the channel and the effect of noise is maximum. We will now derive the analog-receiver error rates required for synchronization to occur successfully. In order to obtain this result, we model the input-output of the analog-receiver as a binary channel, characterized by a probability matrix P .

$$P = \begin{pmatrix} 1 - p_{fa} & p_{fa} \\ p_{md} & 1 - p_{md} \end{pmatrix}$$

Here p_{fa} represents the probability of false-alarm, and p_{md} the probability of missed-detection of the receiver. Consider now that the receiver just successfully detected a pulse, and we are interested in determining the probability density for the arrival time of the next detected pulse, knowing that for an error-free receiver, this would occur at $t = MT_{ck}$. Since as soon




Figure 4.18. Waveforms in the synchronizer in the event of a missing pulse

as a pulse is detected, the receiver restarts its time measurement, the random variable of interest is the *first crossing* of the output of the analog receiver front-end. This occurs at discrete time k if at all previous discrete times the receiver output was 0, and it becomes 1 exactly at this instant. Therefore, the probability density function of this crossing is

$$P(k) = K_1 \begin{pmatrix} (1 - p_{fa})^{k-1} p_{fa} & k \leq M - 1 \\ (1 - p_{fa})^{M-1} (1 - p_{md}) & k = M \\ (1 - p_{fa})^{k-2} p_{md} p_{fa} & k \geq M + 1 \end{pmatrix} \quad (4.9)$$

Where K_1 is chosen such that $\sum_k P(k) = 1$. This probability density function is shown in figure 4.19 for $M=1000$, $P_{md} = 10^{-3}$ and different values of P_{fa} . This graph shows that during synchronization, keeping false-alarm rate minimal is more important than reducing the probability of missed-detection. This occurs as during synchronization, many more decisions for "0" have to be correctly made than decisions for "1" (In this case 1000 versus 1). In fact, approximately the same statistics are achieved with missed-detection rates varying between 10% and 1%. We can use this information to determine the SNR penalty that we incur by synchronizing our system using the data, as opposed to a low jitter clock. The scenario is different if the system synchronizes using a preamble or not. In case a preamble is used, the data bits are known, and P_{md} can be reduced to decrease P_{fa} while keeping SNR constant. For example, at 13dB SNR the pair $P_{fa} = 2.2e-7$; $P_{md} = .1$ can be achieved, which results in an input jitter of 5 clock cycles only. During demodulation, the slicer threshold can be changed to obtain $P_{fa} = 10^{-3}$; $P_{md} = 10^{-3}$ performance. If a preamble is however not used, and data needs to be correctly demodulated, P_{md} cannot be reduced below an arbitrary rate. The case $P_{md} = 10^{-3}$, $P_{fa} = 10^{-7}$ can be achieved at a slicer SNR of 15.4dB, giving a 2.4dB penalty compared to the case of sole data-demodulation. Once the SNR is specified, the bandwidth requirement of the loop, as well as the acquisition period length in number of symbols can be determined. For a cycle-to-cycle jitter target after filtering of $T_j^{(o)} = 2T_{ck}$




Figure 4.19. Probability density function(top) and cumulative distribution of input transition

, and $T_j^I = 8T_{ck}$ input jitter from the data during synchronization $N_{pulse} = (\frac{T_j^{(i)}}{T_j^{(o)}})^2 = 16$ is obtained. These are the values used in our design. Notice that the situation is quite different *after duty cycling is enabled*. In this case, the number of correct "0" decisions to be made is drastically lower and the P_{fa} spec can be relaxed.

Bit-Error rate Estimation

One interesting opportunity given by the chosen modem is that of real-time BER estimation. As we mentioned in the previous section, we utilize error-rate estimation to adapt the radio architecture to the interference scenario. Consider the trellis diagram in figure 4.15 and assume that a pulse has just been detected. If the demodulator, based on past decisions, identifies the present state as "1", it will power up the analog receiver at instants $\frac{\hat{T}}{2}$ and \hat{T} from now to search for pulses. If instead the present state is identified as "0", the receiver will be powered up at \hat{T} and $\frac{3\hat{T}}{2}$. In doing so, we ignore the reception interval around $\frac{\hat{T}}{2}$, as our knowledge of source statistics excludes the possibility of the transmitter purposefully transmitting data in such window. However, if the receiver were to power up and successfully detect a pulse in that window, this would be evidence that its conviction to be in state "1" was incorrect, and as a result that an error occurred in some other detection. This can be extended to a more general situation by considering the forbidden sequences in the trellis in Fig.4.15. It can be seen that all sequences of the kind $\{\frac{T}{2}, T, \dots, T, \frac{T}{2}\}$ correspond to transitions $0 \rightarrow 1, 1 \rightarrow 1, \dots, 0 \rightarrow 1$ which are illegal. Similarly, $\{\frac{3T}{2}, T, \dots, T, \frac{3T}{2}\}$ correspond to transitions $1 \rightarrow 0, 0 \rightarrow 0, \dots, 1 \rightarrow 0$ which are also illegal. It therefore seems

that by adopting a slightly different wake-up schedule, the receiver can estimate the current error-rate in real-time. This is possible as long as the increased observation window does not increase the probability of false-alarms in the receiver above the existing error rate. This is true in the case of this work because as we have seen, a low probability of false alarm is already required by the synchronization process. It can be proven that the average on-time of a system that does not try to estimate error rate is $K_{on} = \frac{3\delta}{2T_b}$, where δ is the duty-cycling window width. For a system employing B.E.R. estimation, this ratio is increased to $\frac{2\delta}{T_b}$. As a result a roughly 33% power penalty in the active power (meaning excluding synchronization) of the receiver is incurred by this choice.

Synchronization simulations

Since the synchronization subsystem is a crucial part of the transceiver, it was also verified individually at the behavioral level. Crossing instances were generated using the distribution 4.9, and fed to a state-machine simulating the synchronizer behavior. The synchronizer starts up in a *Listening State*, in which power gating is inactive. In this phase, incoming edges are detected from the synchronizer and frequency and phase locking is sought. The length of this listening time in terms of pulses successfully received can be varied in simulations. The simulations allowed to investigate the locking behavior of the synchronizer and the impact of error propagation in the synchronization. Ultimately, since the effects of imperfect synchronization can be mitigated by extending the duty-cycling window of the receiver, an optimal listening time exists for the receiver as a function of packet length. A longer packet length demands more accurate synchronization, as the impact of a narrower duty-cycling window can be amortized over a larger number of detections. In a short packet on the hand enabling the duty-cycling mode earlier enables larger power savings. Considering 256 bits long packets, it was found that at a false-alarm rate of 10^{-5} , a minimum duty-cycling width of 35 clock cycles can be employed guaranteeing an error-rate of less than 10^{-3} on all packets when the bit transmission period equals 1000 clock cycles. This duty-cycling window width can be reduced to 25 clock cycles if a false-alarm rate of 10^{-8} is achieved by the analog front-end. Figure 4.20 shows a plot of estimated and simulated B.E.R. as obtained from a *Simulink* model of the complete system. As clearly shown in the graph, the constructed B.E.R. estimate is an extremely good approximation of the real B.E.R., and can be therefore used reliably to adapt the system architecture.

Amplitude Estimation

In addition to timing acquisition, another important task that the receiver has to perform at startup is amplitude estimation to optimally place the data-slicing threshold. It is well known that this optimal threshold (for symmetric data input probabilities and error-costs), is situated halfway through the received eye. Performing this operation requires a training sequence to be transmitted, as well as either a multi-bit ADC or a variable threshold DAC. For short packet lengths, the overhead associated with the increase in ADC resolution and




Figure 4.20. Simulated BER Estimate versus real BER

the training sequence transmission can be significant. To reduce the header length and avoid a multi-bit ADC, we propose to utilize an asymmetric decision scheme, where the threshold is set by fixing a desired false-positive error rate, independent of the signal amplitude. This is advantageous, as since the receiver noise can be assumed independent of the signal level, the decision level can be estimated in between packets, and a shorter header can then be sued. As a result, a simple threshold tuning algorithm can be designed, where first the slicer threshold is initially set to the minimum allowed value, and the false-positive rate measured. The threshold is increased until the false-positive rate is lower than the desired target. Since this system already requires asymmetric thresholds during synchronization, this strategy is particularly suitable. The result of amplitude-independent threshold setting is a degradation in the receiver amplitude-error rate characteristic for both low and high SNR, as shown in curve 4.21. At low signal-levels, errors are dominated by the false negatives, while the floor due to false positive shows up at high-SNR. This degradation is not particularly problematic in neither case. At low SNR in fact, while error rate is degraded, the target 10^{-3} B.E.R would not be achieved anyway. At high SNR, quality of service increase is limited, which can also be tolerated.

4.4 Summary

We presented a complete analysis of the interferer performance of self-mixing receivers. We showed that by replacing baseband integration with high-pass filtering, the false-alarm rate of these receivers can be drastically improved, so that their interference performance




Figure 4.21. Comparison of waterfall curves with conventional and proposed amplitude estimation algorithm

is only limited by false negatives generated by interferers with carriers closer to that of the desired signal than $\frac{1}{T_p}$. In order to remove this limitation, we proposed to *adapt* the front-end architecture depending on the channel state, using a linear downconverter when close in interferers are present. We have also shown how interference detection can be realized through B.E.R. monitoring by using a differential PPM modulation with differential detection, so that minimal extra hardware is required in addition to the low-power timer needed by the duty-cycling function.

Given that our communication system is now completely specified, we can now describe how to implement the RF and baseband functions that we described so far in a power efficient manner. This is the topic of chapter 5.

Chapter 5

Circuit Design

5.1 Introduction

We now describe the transistor-level design of the cm-range transceiver. This system was designed in a 6M1P 65nm Low Power CMOS process from ST Microelectronics, which has a nominal power supply of 1.2V and a FO4 delay of 35pS. A higher performance process (lower FO4 delay) would help reduce the power dissipation of the analog and RF subsystems; but it was not utilized in order to demonstrate the ideas developed in the rest of the thesis in a industry-standard technological platform.

While the major goal of the receiver was to achieve low-operating power, an emphasis was placed also on minimization of area-consuming integrated passives, especially inductors. A single inductor is used in the whole transceiver as part of the input matching network/ duplexer. Thanks to the low required RF frequency accuracy, the resonance of the matching network can also be used as a frequency reference for the RF oscillator, resulting in a compact implementation that is desirable for low-cost systems.

A high level description of the system architecture is given in figure 5.1. Transmitter and receiver share the antenna through a single on-chip matching network. The LO generation block is also shared. While the receiver front-end employs the dual-mode diode/mixer mentioned in the previous chapter, circuit considerations induced some changes in the baseband stages. In particular, running the synchronizer out a single-phase 2.5GHz clock results in excessive power dissipation. A 5-phase ring oscillator is therefore employed to generate 400pS spaced signals while running the synchronizer at 500MHz. As a result, a 5-way interleaved baseband is employed, and the inter-pulse intervals are measured using a combination of counting and phase picking. All the RF circuits on the chip share a 1.2V supply voltage, while the baseband and clock recovery operate off a 0.75V V_{dd} . While a finer-grained V_{dd} allocation strategy could have resulted in a lower power dissipation, reducing the number of

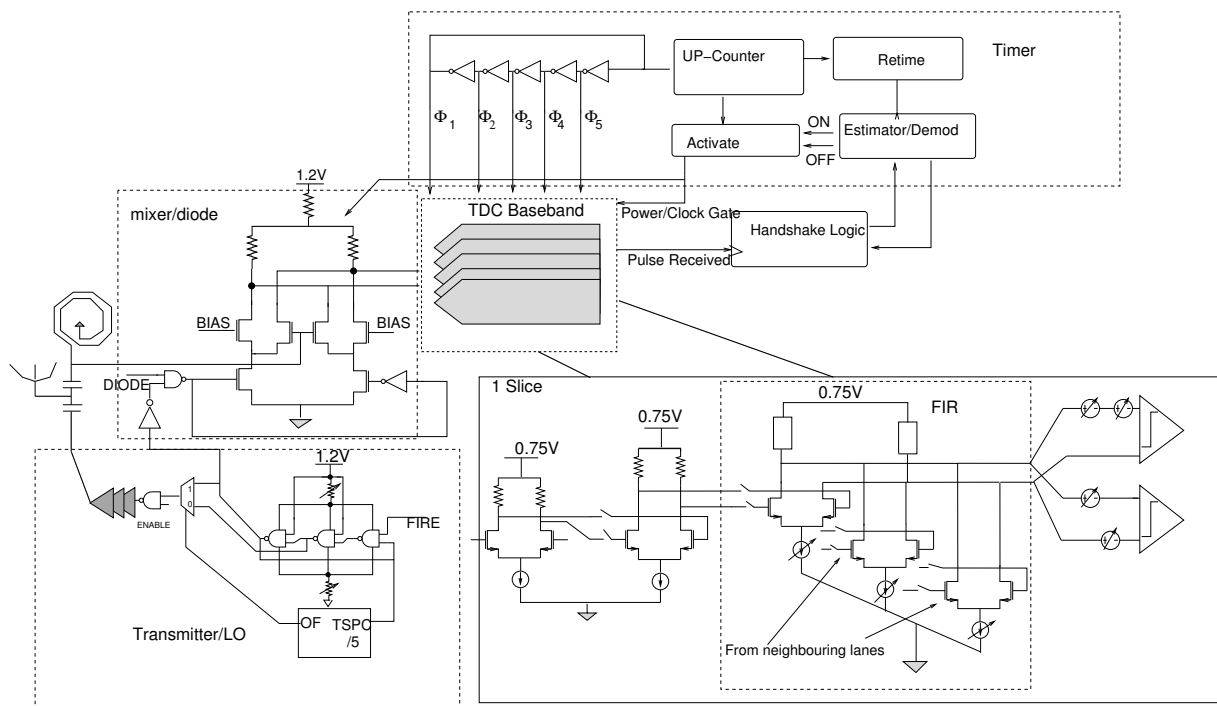


Figure 5.1. Block diagram of transceiver architecture

supplies is preferred to avoid further losses and complexity in the (not implemented) supply generation and distribution subsystem.

5.2 Transmitter Circuits

5.2.1 Architecture selection

Compared to many general-purpose UWB pulsed, the transmitter in this system necessitates only a relative small operating frequency range (6GHz plus process variations), while the pulse rate is fixed to a value of 1MBps. Differently from most of the literature, this transmitter also requires a high-speed antipodal modulator to realize the $1 - z^{-1}$ baseband pulse shape that we chose to increase interference robustness. Thanks to the low-data rate, spectral mask is easily met at the chosen 6dBm peak output power, and complex pulse shaping circuits are not required, increasing the power efficiency.

To understand the tradeoffs in the implementation of this block, a literature search in the space of pulse-based transmitters was first conducted, and we identified three main architectures for UWB pulse generation in the 3-10GHz band. The main distinction lies in whether

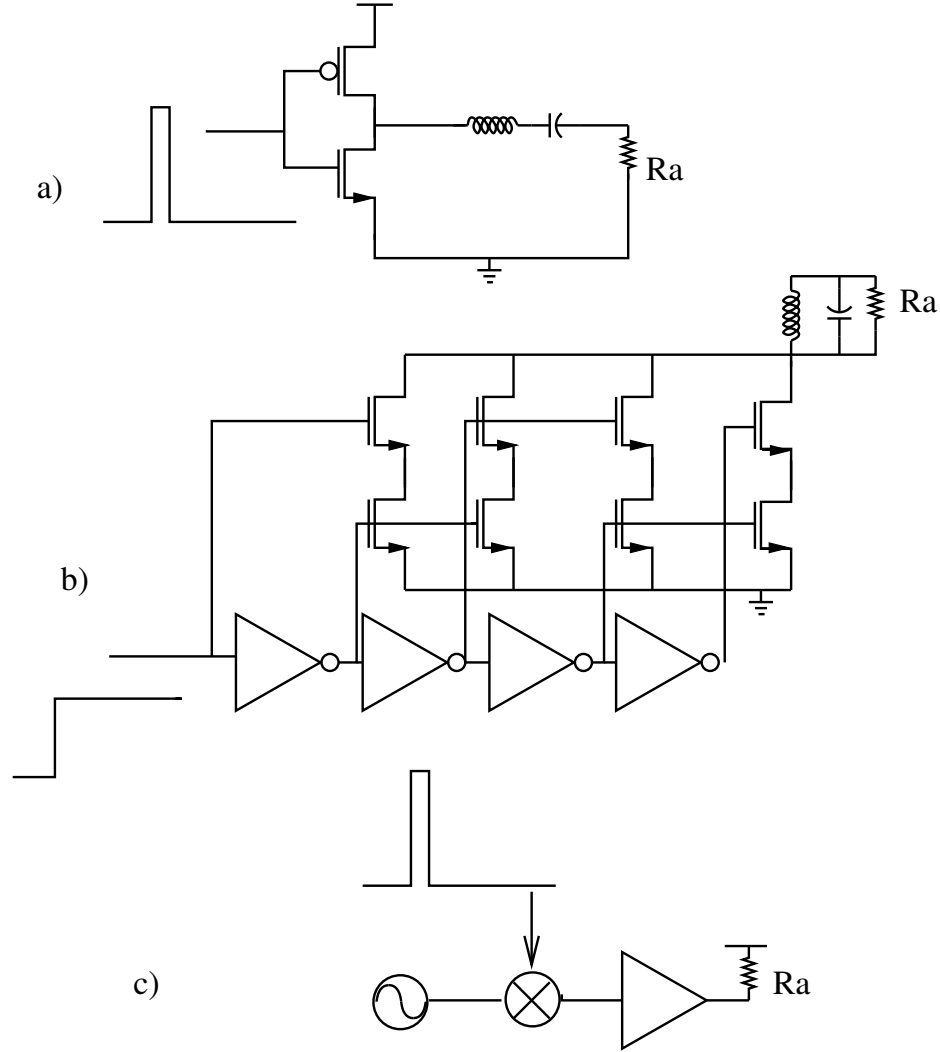


Figure 5.2. Bandpass Pulse Generation Architectures: a) filtering; b) Edge Combining; c) Gated Oscillator

an architecture employing a local oscillator[16],[47]¹ or not [13],[31],[1] should be used. In general, architectures that do not employ a local oscillator claim the advantage of simplicity, but compromise output power and robustness. In [31] for example, the pulse is generated at baseband, and filtered by a passive, LC- based RF band-pass filter. This method suffers from poor efficiency, as much of the energy in the baseband pulse is dissipated in filter instead of radiated, and has either limited tuning range on the RF carrier or the pulse bandwidth, as both are controlled by the same LC filter. [13] and [1] on the other hand, generates a pulse by edge combining from a delay line. While this method is potentially attractive, the high-parasitic capacitance associated with the edge combiner make it difficult to obtain a rail-to-rail swing at the output of the pulse generator, and hence either results in low output power, or in a challenging and potentially inefficient driver. The transmitter in [53] employs a local oscillator, modulated by a triangular wave generated through a separate on

¹ Architectures employing an LO are referred to as *Gated Oscillator* in figure 5.2

chip-circuit. This transmitter achieves higher output power, and low power consumption; however it is susceptible to variations in the triangular wave generator, which are independent from those in the local oscillator and require to be calibrated. A better approach is that showed in [53],[24] and later in [47], where a local oscillator is used and the baseband pulse shape is derived from its period through high-speed synchronous division. This approach does not require calibrations, but requires high-speed mixed-signal circuits to perform the pulse shaping, which are readily available in modern technologies.

The architecture in [53] was chosen for this transmitter, and to maximize simplicity and power efficiency, a single-ended chain has been chosen. In order to realize an efficient transmitter, a low power pulse shaping subsystem, composed of an oscillator, a frequency divider and a PSK modulator, is necessary. In addition, a low power antenna driver is also required. These building blocks are described in the following sections.

5.2.2 Oscillator

We have seen that in ultra-wideband systems, local oscillator phase noise performance is relaxed due to the large signal bandwidth. Therefore, we can choose between LC and ring oscillators; in order to minimize power consumption and obtain the desired tuning range. We have shown in [40], that the impedance level of capacitors is superior to that of a resonant tank at 2GHz in 90nm CMOS, and therefore when phase noise is not a concern, ring oscillators should be used as local oscillators. At 6GHz, in the chosen 65nm CMOS process, the impedance of a $.5\mu/.06\mu$ device is $-j44K\Omega$, which is still at least an order of magnitude larger than what can be obtained with an LC tank. We therefore employ a ring oscillator as the LO. Another advantage of using a ring oscillator lies in the fast startup, which is due to the low-quality factor. This helps minimize the transmitter startup time and reduce overhead power. Startup time can be further minimized by using a NAND-based oscillator (Fig.5.3). When EN is low, the p-mos devices in the NAND gate pull the output to V_{dd} , and the loop cannot oscillate. When EN is raised to 1, the output of the first NAND to the left is pulled to ground, and oscillation is started. Notice that the oscillator startup is repeatable, and coherent with the EN signal. Digital tuning capability with $200MHz/step$ resolution is required to deal with process variations. To minimize analog voltages, this is obtained by using a digitally programmable resistive string, as used in [40], to change the virtual V_{dd} of the oscillator. Because

$$V_{dd}^{(v)} = V_{dd} - RI_{dd} = V_{dd} - RC_{osc}V_{dd}f_{osc} \quad (5.1)$$

$$f_{osc} \approx f_{osc}^0 + \alpha(V_{dd}^v - V_{dd}) \quad (5.2)$$

, we find that in order to obtain a linear code-frequency characteristics, the resistance should be made linearly related to the code. Using straightforward parallel MOS switches results instead in a conductance that is linear with the code, which makes the tuning curve highly non-linear and ultimately demands an increase in number of bits for a given resolution. We therefore used a hybrid series-parallel configuration to increase linearity, which is shown in figure 5.6. The simulated oscillator frequency versus tuning code characteristic is shown in 5.4. An specific issue arises when resistive supply tuning is coupled with oscillator duty-




Figure 5.3. Digitally Controlled Oscillator Schematic




Figure 5.4. Simulated tuning range




Figure 5.5. Impedance network model

cycling. When the oscillator is off, $I_{dd} = 0$ and hence $V_{dd}^{(v)} = V_{dd}$. As a result, the first edges propagating through the ring travel at increased speed when compared to those traveling when the supply network has settled. This effect introduces a trade-off between frequency accuracy and startup-time. To surpass this tradeoff, the positive and negative virtual supplies are decoupled on chip, and the tuning resistors are all turned off when the oscillator is deactivated, trapping charge stored on the decoupling capacitor. The value of the decoupling capacitor must be chosen carefully. A too small decoupling capacitor cannot prevent the charge from leaking away due to the switch off-currents. A too large capacitor on the other hand does not allow the supply network to settle quickly enough. For this design $R = 1k\Omega$, result in a choice of $C_{dec} = 1pF$ to obtain a fast settling of the supply combined with ripple smaller than 10mV. A complete schematic of the VCO and the symmetric root clock drivers is shown in figure 5.3. In post-layout simulations, the VCO core consumes 140uW at 6GHz from a 1.2V supply (operating on a virtual supply of 1V).

5.2.3 Divider

The divider chain used to generate the baseband waveform needs to provide maximum speed at low power. For this reason, TSPC logic is used in the divider chain, and a Johnson architecture is chosen for its small logic depth and low power (See figure 5.7;[16],[58]). Since the TSPC divider uses dynamic storage, the main challenge in the design lies in simultaneously guaranteeing high speed during the pulse transmission and long data retention in between pulses. In particular, called C_{node} the capacitance at a generic node in the signal path, the logic delay t_d is proportional to $\frac{C_{node}}{I_{ON}}$, while the leakage rate is $\frac{C_{node}}{I_{off}}$. Therefore, operating this divider correctly requires a $\frac{I_{on}}{I_{off}} \geq \frac{T_b}{T_{delay}} = 25000$. In this process it is impossible to ensure this I_{on}/I_{off} in the worst temperature and process conditions (i.e. I_{on} measured on slow, cold, while I_{off} on fast, hot). As a result, keepers were added to the basic TSPC

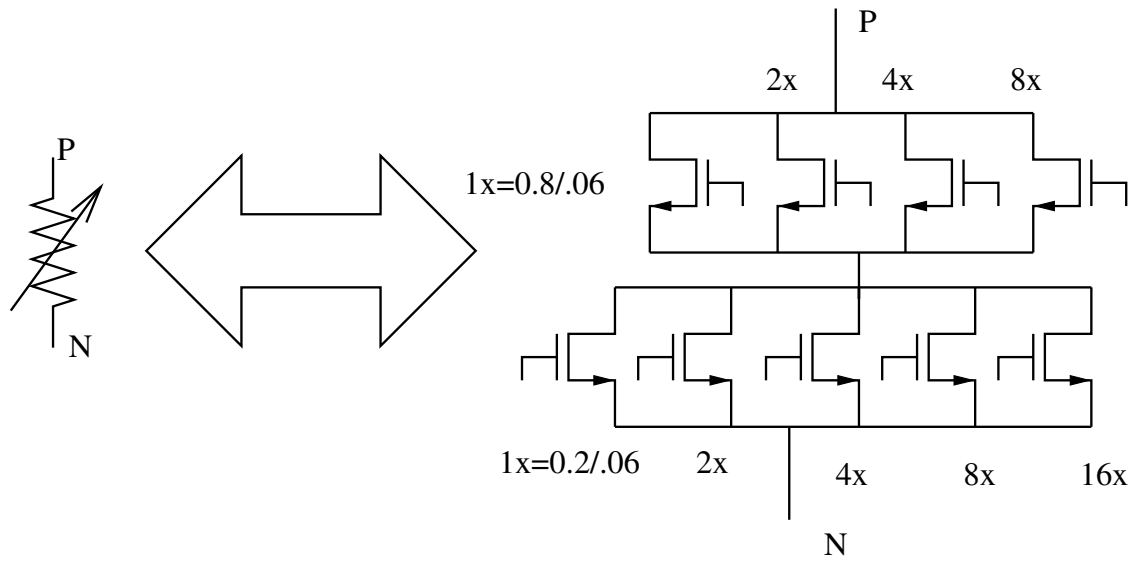


Figure 5.6. Tuning control shown for the bottom rail. Top rail is symmetric but uses PMOS transistor with 2x width




Figure 5.7. Johnson divider architecture

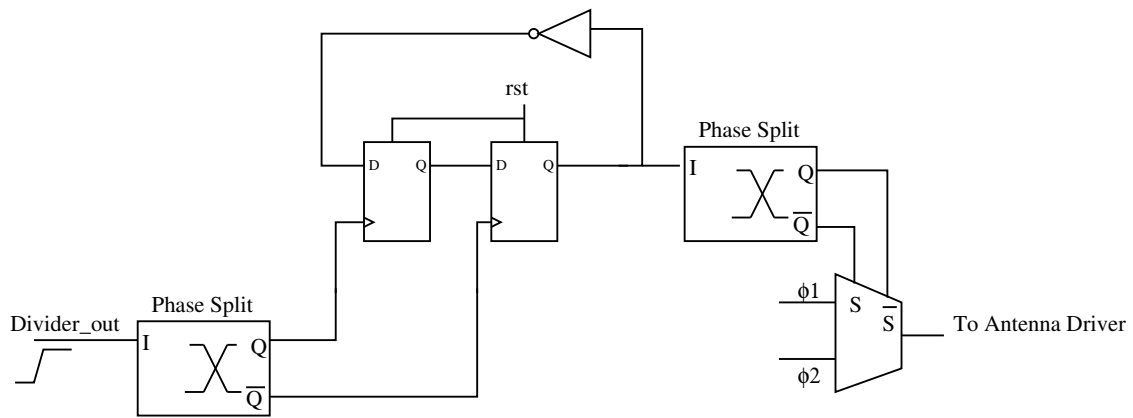


Figure 5.9. PSK Modulator/multiplexer and control circuits




Figure 5.10. Simulated Pulser Waveforms




Figure 5.11. Effect of delayed inputs on glitching at the output of the multiplexer

of $\{1, -1\}$ waveform. A small delay compared to the 160pS carrier period is required in switching the phase of the carrier to avoid glitching. In order to minimize this effect while still generating the phase inversion using a single-ended pass-transistor based multiplexer, we use two techniques. First, the switching of the multiplexer is resynchronized to the rising edge of one of its inputs ϕ_1 . Since $\phi_1(t)$ is a square wave, the switching occurs then during the time the square wave is at logic "1", minimizing the effect of the delay. Second, the inputs of the multiplexer are connected to two different phases of the ring oscillator ϕ_1 and ϕ_2 , which are in the relationship $\phi_2(t) = \phi_1(t + \frac{T_c}{6})$. This extra delay compensates for the multiplexer delay and further reduces the glitching at the output, as shown in figure 5.11.

The Multiplexer directly feeds the antenna driver/PA, which is realized with a static CMOS inverter chain. The last stage of this chain is sized to obtain an output resistance of 10Ω in the typical corner, corresponding to NMOS and PMOS widths of $32\mu m$ and $64\mu m$ respectively. To ensure proper functionality across corners, the chain is sized with a fan-out equal to 2; so that the combined effect of the finite output impedance and the fan-out limits the driver efficiency to $\eta_{Driver} = \frac{5}{6} \frac{1}{\sum_0^\infty 2^{-k}} = \frac{5}{12} \approx 40\%$. The average simulated power consumption of the transmitter, when running at 1Mbps, is $28\mu W$ from a 1.2V supply, or 28pJ/pulse. The overall power breakdown for the pulser is shown in Table 5.1. The energy efficiency is substantially limited by the antenna driver, which by itself has an efficiency of 20.1%. This efficiency is much lower than the theoretical value of 40% due to the capacitive losses in driving the gates of the CMOS transistors. Leakage currents account for 2.4pJ per pulse, and are therefore not a limiting factor at this frequency. A total efficiency of 18%

(excluding matching network losses) is simulated, which is superior to all the published work in general purpose UWB transmitters.

Energy Delivered to the antenna	5pJ
Driver	18.9pJ
Oscillator	0.282 pJ
Divider and Modulator	1.55pJ
LO distribution	1.1pJ
Total	27.5pJ

Table 5.1. Power consumption breakdown in pulse generator

5.3 Receiver Circuits

We now describe the receiver circuits, from the antenna to the comparators. The receiver should be able to detect pulses with a sensitivity of -24dBm peak power, corresponding to -52dBm average power or 15mV 0-peak impulse amplitude, consuming minimum power. The receiver adopts a single-ended matching network and antenna interface, given that this solution, as shown in [40] improves the sensitivity in diode mode.

5.3.1 Matching network/duplexer

Since transmitter and receiver share a single antenna, a strategy to share this resource is necessary. In impulse-UWB systems, the short signal duration makes it natural to opt for a Time-Division-Duplex (TDD) solution, where pulse transmission and reception do not occur at the same time. In addition, the low-output power of the transmitter makes it unnecessary to isolate transmitter and receiver with a switch to ensure reliability. For these reasons, transmitter and receiver are connected to two different ports of a single matching network. We seek a matching network design that allows us to exploit the inherent asymmetry of UWB systems, i.e. to shift the burden from receiver to transmitter as much as possible. Consider now a generic two-port matching network, shown in figure 5.12. This network transform the antenna impedance at the transmitter port to a value R_{tx} , while the resistance at the receiver side is R_{rx} . Assume now that we are interested in maximizing the *Voltage* induced by the transmitter at the receiver input, considering the transmitter has a maximum output swing V_{dd} . For a channel power loss G_{ch} , we obtain

$$V_{rx} = V_{dd} \sqrt{G_{ch}} \sqrt{\frac{R_{rx}}{R_{tx}}} \quad (5.3)$$

We therefore desire a matching network that maximizes R_{rx} and minimizes R_{tx} . An example of such matching network is shown in figure 5.13 ([8]). This matching network is the same

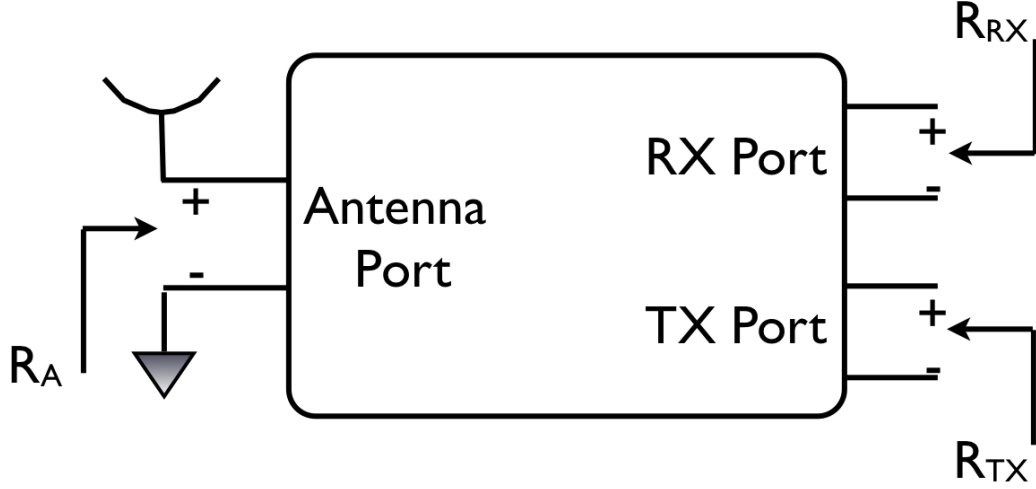


Figure 5.12. High-level matching network view

used in [8], but transmitter and receiver are connected to two different ports. Reasoning from the receiver side at the parallel resonance of the tank, the parallel impedance of the tank R_p is transformed to the antenna impedance R_A by the capacitive transformer C_1, C_2 . This provides a voltage gain $G = \frac{1}{2} \sqrt{\frac{R_p}{R_A}} = \frac{1}{2} (\frac{C_2}{C_1} + 1)$. From the transmitter port, the real part of the impedance looking into the matching network at the same frequency can be calculated (see Appendix) to be

$$Z_{tx} = R_A \frac{(\frac{C_2}{C_1})^2 + \omega_p^2 C_1^2 R_s (R_A + R_s)}{(\frac{C_2}{C_1})^2 + \omega_p^2 C_1^2 (R_A + R_s)^2}$$

. Where $R_s = \frac{\omega_0 L}{Q}$ is the inductor series resistance. For high inductor quality factor (high-Q) $\frac{C_2}{C_1} \rightarrow \infty$ and hence $Z_{tx} \rightarrow R_A$. In addition, a small imaginary part exists in series with Z_{tx} . This imaginary part can however be neglected for high-Q matching networks.

Taking a component sizing perspective, since $R_p = \omega_0 L_m Q_s$, to maximize the gain the value of both quality factor and inductance should be maximized. However, an increase in quality factor increases the group delay associated with the transfer function from the antenna to the RX input, which is undesirable. Similarly, the inductance L_m cannot be increased without limit due to self-resonance limitations. We chose an inductor value of 1.5nH, with an unloaded quality factor of approximately 10 at 6GHz. The layout of the inductor was generated by an inductor-selection tool integrated in the design kit, resulting in a 3.25 turn layout with 10 μ m metal width 158 μ m diameter. Since we are operating with small inductance and capacitance values, a design methodology to include the parasitics associated with the bonding pad and the associated wire inductance had to be developed. An estimate of the pad capacitance and the bondwire inductance is assumed known. In this case, the effective antenna impedance as seen by the chip can be estimated by Thevenin

equivalence. For a 50Ω antenna with 1nH series inductance and 120fF pad capacitance, the parallel equivalent circuit of the source at 6GHz is given by $Y_A = \frac{1}{78} - j\frac{1}{200}$, equivalent to a 50Ω resistor in parallel with a -137fF capacitor. Values of C_1, C_2 are found that match the 681Ω inductor parallel resistance to 78Ω , while providing a resonance of 6GHz

$$\sqrt{\frac{C_1 + C_2}{LC_1C_2}} \frac{1}{2\pi} = 6\text{GHz} \quad (5.4)$$

$$\frac{C_2}{C_1} = \sqrt{681/78} - 1 = 1.95 \quad (5.5)$$

This results in values $C_1 = 700\text{fF}$, $C_2 = 1.4\text{pF}$. C_2 is then increased to compensate for the negative 137fF of input capacitance, giving $C_2 = 1.54\text{pF}$ and has digitally controlled value ensure robust operation. The simulated value of S_{11} for different values of C_2 is shown in figure 5.14. The impedance from the transmitter side is instead shown in figure 5.15 and at 6GHz is $23 - j19\Omega$. This degrades the transmitter output power by less than 1dB . Furthermore the matching network provides a passive gain of 5dB from the antenna to the receiver input, boosting the input voltage of the diode to 30mV 0-peak and relaxing its noise specifications.

5.3.2 Reconfigurable mixer/diode

Several topologies of reconfigurable mixer-diode were considered to find one that guarantees low power consumption and high sensitivity. Since the receiver is expected to operate in diode-mode most of the time, the focus was given to optimizing the performance in these conditions. We first considered the classic envelope detector described in [52] and shown in figure 5.16. This topology is well understood and robust; and it can be operated as a mixer by periodically switching the biasing current-source M2 on and off. However, this circuit has two limitations when used in ultra-wideband applications. The first issue is that the topology is based on a source-follower, where rectification is obtained thanks to the V-I characteristic of M1. However, the transconductance of M1 sets both the value of the rectifier gain as well as the output impedance, which is fixed to a $\frac{1}{gm}$ value. As a result, the nonlinear voltage gain of this topology is limited to

$$\frac{V_o}{V_{in}} = \frac{V_{in}}{4(nV_t)}$$

for a transistor biased in subthreshold.

Second, due the circuit structure, the nonlinearity is not excited by the input voltage directly, but rather by the transistor V_{gs} , which is a high-pass version of the input. Using basic circuit analysis, the transfer function from input to V_{gs} can be calculated to be the value in 5.6.

$$\frac{V_{gs}}{V_{in}} = \frac{\frac{C_L s}{gm}}{1 + \frac{C_1 + C_L}{gm} s + \frac{C_1 C_L s^2 R_s}{gm}} \quad (5.6)$$

To obtain high-bandwidth and maximize conversion gain, device M1 should be biased in the sub-threshold region. However, this increases the device capacitance C_1 and hence reduces




Figure 5.13. Matching Network Used in The Design




Figure 5.14. Simulated input matching network for different values of digitally controlled capacitor C_2




Figure 5.15. Simulated impedance at the transmitter port




Figure 5.16. Envelope Detector topology of [52]

the overall gain again. For example, to obtain a bandwidth of 1GHz on a 100fF load, at a transconductance efficiency of 24, a transistor width of $40\mu m$ and a length of $.06\mu$ is necessary, with $25\mu A$ bias current. This device has an intrinsic capacitance of the order of 26fF, which results in a voltage loss of $2dB$ in the $V_{in}-V_{gs}$ conversion, and hence a $4dB$ loss in the overall conversion gain.

In the common-source (CS) topology (See 5.18,[1]) instead, the transistor non-linearity rectifies the input voltage translating it into a current, which is then converted into a voltage by the load resistor R_L . Choosing $R_L \geq \frac{1}{gm}$ voltage gain is obtained at the expense of bandwidth. Alternatively, the bandwidth can be maintained constant while increasing the power dissipation and input capacitance, thus decreasing the input-referred noise of the baseband stages and hence improving the receiver sensitivity. In order to compare the two topologies, AM-detectors were designed with both topologies using identical core-transistor size and bias conditions. The simulated voltage conversion gain, output noise, and signal-to-noise ratios are reproduced in figure 5.17, showing that the CS topology provides higher gain and signal-to-noise ratio over the input voltage range of interest.

Based on this discussion, we chose to utilize a common-source AM detector in this work. Several modifications were introduced in the basic structure of 5.18 to improve performance and enable dual-mode operation. First, a cascode transistor was added at the drain of the transconducting device M1. The cascode device improves output impedance and hence gain with a small noise penalty; additionally, it reduces the swing on the drain of M1. This is important as drain voltage and current are related in a submicron device by Drain Induced Barrier Lowering (DIBL). Due to DIBL, a decrease in drain-source voltage increases the transistor threshold voltage, and hence it decreases the current. Since the common-source amplifier is inverting, drain oscillations are in anti-phase with gate oscillations, meaning the nonlinearity due to DIBL will act to cancel the nonlinearity due to the I_d, V_{gs} characteristic. Transistor sizes were chosen to optimize conversion gain in diode-mode operation. In particular, choosing the width of transistor M1 is critical as it entails a trade-off between the conversion-gain of the diode device, and the capacitive load imposed on the matching network, which limits the matching network gain. In principle, it is beneficial to decrease the size of transistor M1 if this allows increasing the gain of the matching network, as this




Figure 5.17. Output SNR and Amplitude level comparison for CS and CD diodes

has a quadratic effect on the overall conversion gain. Unfortunately however, exploring this trade-off requires a transistor non-linearity model that maintains accuracy also in the moderate inversion region. To avoid the risk of model inaccuracies, transistor M1 was sized $32\mu\text{m}/.06\mu\text{m}$ and its bias current was made programmable with a 1:7 range from $80\mu\text{A}$ to $560\mu\text{A}$ through an on-chip programmable resistor. In the nominal operating point, the bias current is $220\mu\text{A}$, corresponding to a transconductance efficiency of 18. $2.7K\Omega$ un-salicided polysilicon resistors R_L guarantee a 600MHz bandwidth on the nominal 100fF load.

In order to reconfigure the device into a mixer with minimal overhead, we decided to introduce switching at the source of the trans-conducting devices and reconfigure the circuit into a switched-transconductor mixer ([14]). Compared to a Gilbert-cell mixer, a switched transconductor mixer can achieve high-performance even when the switching devices are operated in the triode region, and hence has larger headroom potential. In addition, the LO drive can be effected through a simple static inverter chain, without a concern on limiting the swing to avoid large glitches on the transconducting devices drain terminals. The complete schematic of the reconfigurable topology is shown in figure 5.19. A double-balanced, pseudo differential structure is chosen for its superior LO-IF leakage performance. Configuration is accomplished by driving the LO ports appropriately. When LO is a static logic 1 and $\overline{\text{LO}}$ is a logic 0, one side of the mixer is shut-off and the circuit operates as a pseudo-differential common-source amplifier. Otherwise, LO and $\overline{\text{LO}}$ can be connected to complementary square waves, in which case the circuit operates in mixer mode.

An important consideration in mixer mode operation is that since large devices and large transconductance values are used in diode mode, the LO buffer chain consumes a significant amount of switching power. This is particularly true as the on-resistance of the clock-driver

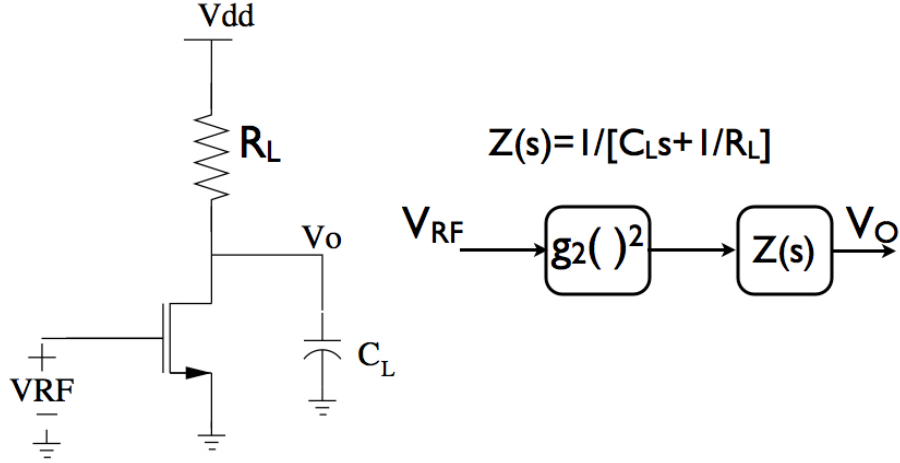


Figure 5.18. Common-source diode

acts as a degeneration resistor to the common-source amplifier in 5.18, and hence should be minimized, requiring wide buffer transistors. To minimize the power consumption of the clock drivers, a static switch M2 with low R_{on} is inserted in parallel with the clock buffer final stage, decoupling its width from the performance in diode mode. In addition transistor M1 and its clock driver are realized as the parallel combinations of four individually enabled unit cells containing an 8/.06 transistor and its clock buffer. A graceful power-gain tradeoff can then be realized in mixer mode by enabling a subset of unit cells at the time. A last issue occurring in mixer mode, is that a common-mode current is injected into the load by the LO drivers due their non-50% switching point. This current results in a variation in the stage output common-mode between mixer and diode mode, and could drive the succeeding stage outside of the correct operating region . The programmable common-mode resistor array R_{cm} , realized with an array of digitally controlled transistors, is introduced to compensate this problem. The simulated performance of this block is summarized in Tab.5.2, while the overall CAD layout of the RF front-end of the chip is included in figure 5.20.

Diode Mode Gain @ 20mV input	0.7
Diode Mode Output Noise (rms)	1mV
Diode Mode Power Consumption	$200\mu W \rightarrow 720\mu W$
Mixer Mode Conversion Gain	$-4.5dB \rightarrow 6.5dB$
Mixer Mode Output Referred Noise(rms)	$386\mu V \rightarrow 525\mu V$
Mixer Mode Power Consumption	$550\mu W \rightarrow 2mW$
3dB Bandwidth	590MHz

Table 5.2. Mixer-Diode Performance

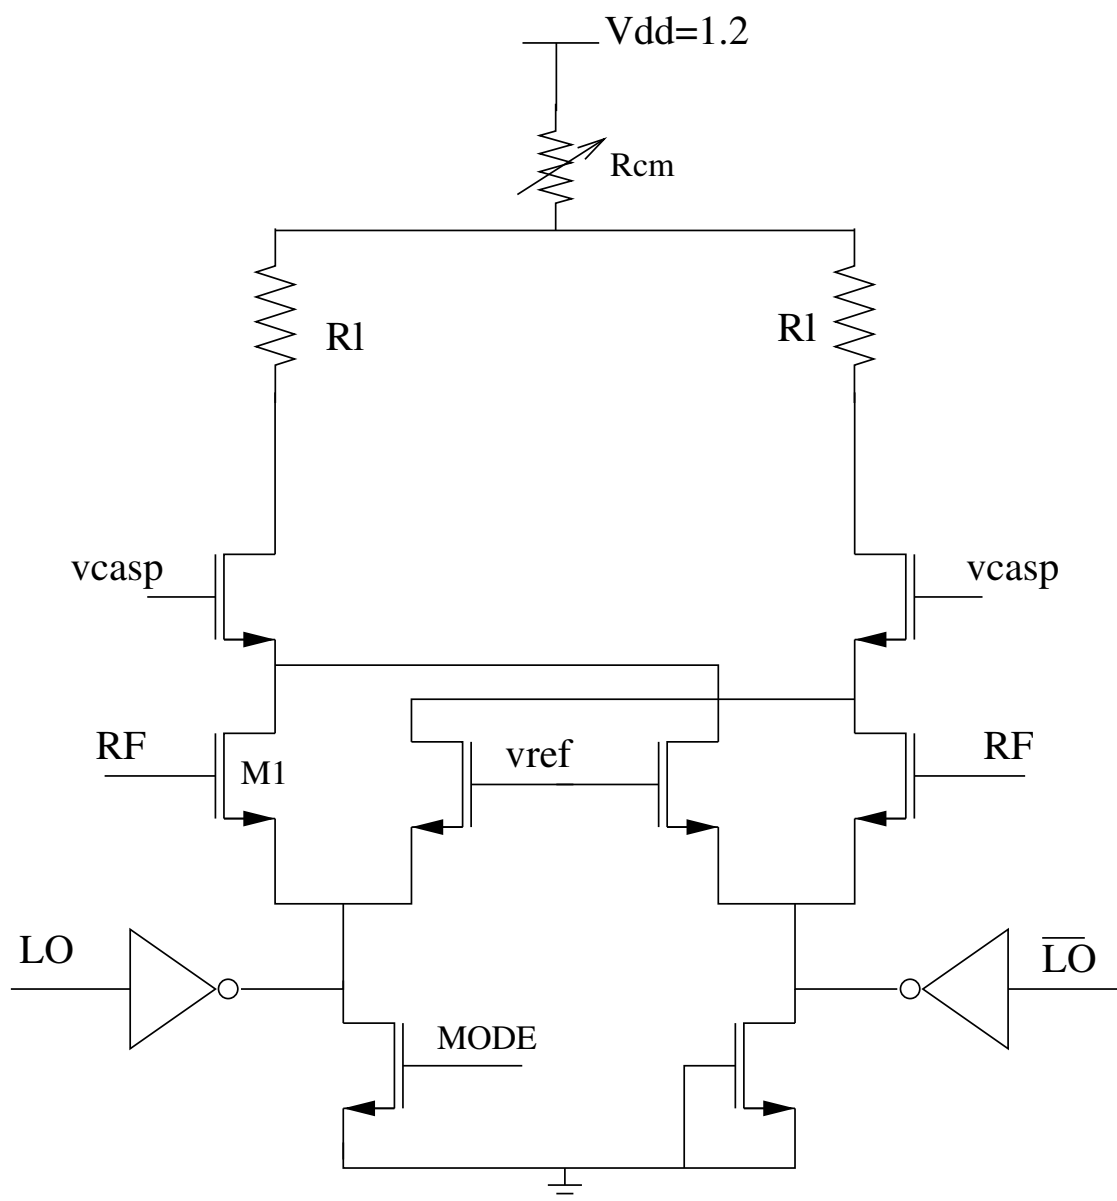


Figure 5.19. Implemented Mixer-Diode stage




Figure 5.20. CAD Layout of RF Front-end

5.3.3 Baseband Filter

The interference rejection properties of the system rely on the performance of the baseband filter. While this block was described as a simple differentiator in the previous chapter, more careful considerations show that this is not feasible. In fact, we mentioned in the previous chapter that in mixer-mode the receiver would operate as a direct-conversion radio. However, notice that the RF carriers of TX and RX are not enforced to be at the same frequency, and neither are the baseband clocks. As a result, the phase of TX and RX carriers can be modeled as uniformly distributed over the $[-\pi; \pi]$ range. A direct conversion architecture cannot therefore be used, as when the phase wanders around $\pm\pi/2$ it would generate no baseband output, raising the error floor. This issue was solved in [25] by using a quadrature receiver and nonlinearly combining the receiver outputs after correlation and digitization. This solution pays the price of having quadrature mixers, LO distribution as well as baseband circuits, nearly doubling the power consumption of the receiver. An alternative solution is to perform the energy detection and signal conditioning at IF (As in[40]), by enforcing a frequency shift between transmitter and receiver that is large enough that

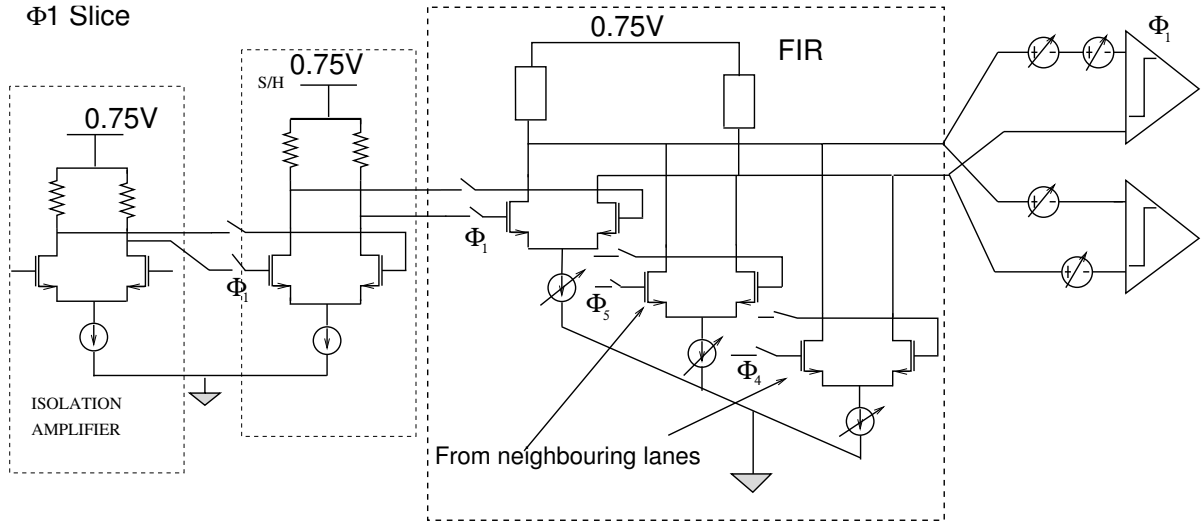


Figure 5.21. Circuit diagram for lane of interleaved baseband

even signals starting with a $\pm\pi/2$ initial phase would be detected correctly. There is a obviously a tradeoff between the chosen IF value and power dissipation. We found that by using absolute value slicing(see below) a minimum shift of 200MHz is required to detect correctly the transmitted 1.6nS pulses. As a result this value was chosen as nominal distance between RX and TX carriers. This choice places the image-frequency within the down-converted channel resulting in relaxed image rejection specifications.

Programmable Filter

Operating the circuits at IF converts the high-pass filter into a notch filter centered at 200MHz, increasing the required FIR filter order from 2 to three taps. The filtering operation is obtained by sampling the incoming signal on the gates of transconducting transistors, which convert their differential input voltage into a current proportional to their transconductance. Currents from different differential amplifiers are then summed into a resistive load. The filter weights are then the transistor transconductances, which can be varied by varying the tail current of each differential pair with a current-mode- DAC. To embed the filter in the interleaved baseband, 5 separate groups of 3 transconductors each are used (Total of 15 transconductors), with sampling phases connected to consecutive phases of the ring oscillator(Fig. 5.21) and each filter driving a slicer operating on different clock phase. For instance, the filter driving the ϕ_1 slicer samples inputs on phases ϕ_1, ϕ_5, ϕ_4 , while the filter slice which feeds the ϕ_2 comparator combines samples taken on ϕ_2, ϕ_1, ϕ_5 .

As mentioned above the filter transfer function is given by

$$H(z) = R_L(gm_1 + gm_2z^{-1} + gm_3z^{-2}) = gmR_L(1 + \alpha_1z^{-1} + \alpha_2z^{-2})$$




Figure 5.22. Loads with high common-mode compliance: diode (left); and diode with level shift(right)

where for $F_s = 2.5GHz$ and a 200MHz notch, $\alpha_2 = 1, \alpha_1 = -0.87$. The ratio of the filter DC gain to the total DC current in this case is $\frac{1-|\alpha_1|+|\alpha_2|}{1+|\alpha_1|+|\alpha_2|}$. Since this ratio is lower than 1, a large DC drop is required on the resistors if a DC gain larger than 1 is desired. To overcome this issue, the filter load is constructed using a low-voltage active load with low common-mode resistance. This load is shown in figure 5.22. It is derived from the classic PMOS diode load shown to the left of the same figure, but with an added NMOS source follower to drive the gates of the PMOS devices. This biases the PMOS devices with $V_{gd} = |V_{thP}| - V_{thN} \approx 0$. To guarantee operation across process variations, the NMOS is then implemented with a standard V_{th} device, and the PMOS with a *low* $-V_{th}$ device, increasing the nominal V_{gd} of the PMOS to $\approx 30mV$ and leaving margin for mismatches. The complete schematic of filter slice is shown in figure 5.23. Butterfly switches implemented in the cascodes allow sign-selection on each of the paths, and combined with the current-DACs setting each transconductor bias enable control of the transfer function. The load resistor of $12.8K\Omega$ give a settling time constant of 500pS on the 40fF load coming from the parasitics and the comparator input. Each stage in the filter consumes $20\mu A$ from the 0.75V supply, leading to a total power dissipation of $300\mu A$ or $225\mu W$. The input capacitance of each trans-conducting stage is 30fF, leading to a 90fF average load on the S/H stages.

The design of the bias DACs also presents some challenge as current-source transistors need to be duty-cycled between pulses, leading to a trade-off between turn-on speed, matching and power dissipation. We found that using a 4 bit DAC cell with fully-thermometer architecture, to obtain an $\sigma(INL) \leq 0.5LSB$ a moderate unit current matching of 25% is required. Assuming a $5mV/\mu m$ V_{th} mismatch coefficient, a $\frac{2\mu m}{.06\mu m}$ device biased at $gm_{eff} = 15$ achieves these results with $800nA$ bias and is then utilized. The complete DAC unit cell is shown in figure 5.24, and a plot of INL/DNL variance in figure 5.25.

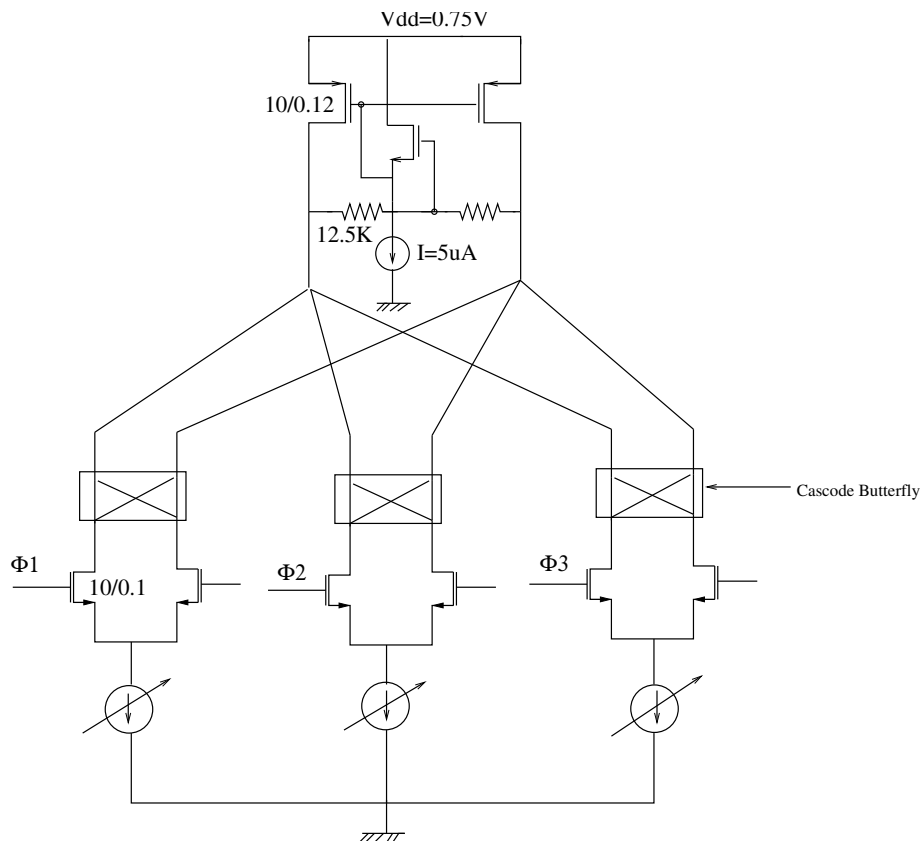


Figure 5.23. Distributed filter schematic (1 lane)

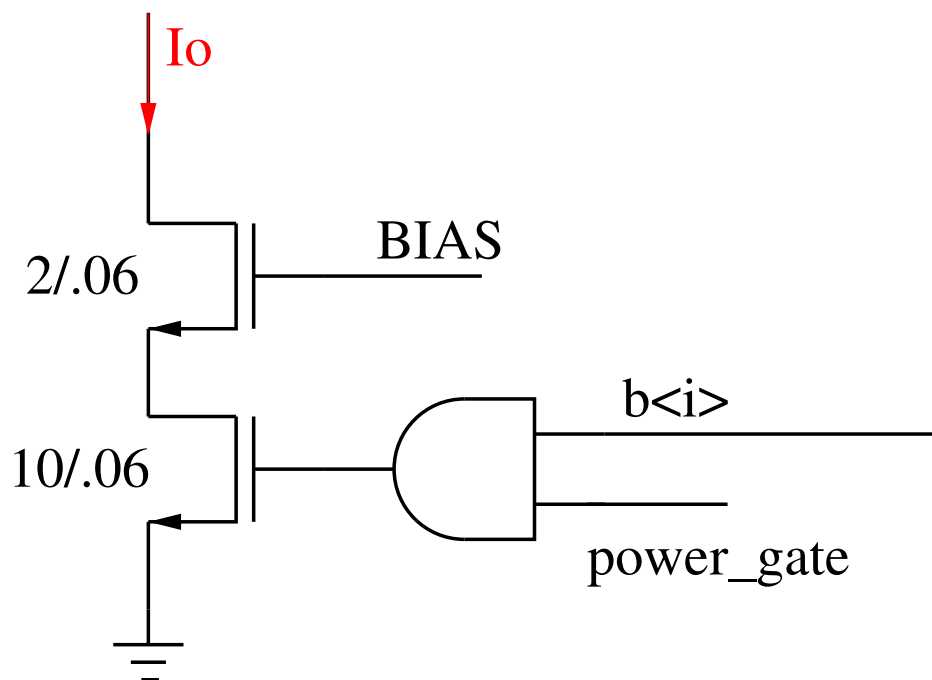


Figure 5.24. Dac Unit Cell Schematic. Nominal LSB current is $1.5\mu A$




Figure 5.25. INL and DNL standard deviation for one the FIR filter bias-DACs

Preamplifiers and Sample and Hold

The filter is preceded by an active track and-hold that relaxes its bandwidth requirements and reduces glitching errors. Sampling is performed open loop onto the gates of NMOS transistor MN1, MN2 (See figure 5.26). PMOS transistor switches with common mode close to V_{dd} are employed, and dummy switches are added to compensate for clock feedthrough and charge injection errors. Under nominal conditions, each of the 5 S/H circuits consumes $15\mu W$, and achieves a 3dB bandwidth of $500MHz$ and a voltage gain of $2.4dB$. The switch capacitance associated with multi-phase sampling is further isolated from the sensitive mixer-output by a continuous-time voltage amplification stage with a nominal gain of 8dB (See figures 5.1 and 5.27). Since this stage operates across the 1.2V domain and the 0.75V domain, it has a high common-mode input voltage, which prevents the use of cascoding in its transconducting core. A tradeoff exists in the design of this isolation amplifier. A single isolation amplifier loaded by the 5-way interleaved samplers could be used. Since at most 3 track and holds are in the transparent mode at any point in time, the circuit can be designed for a load capacitance which is only three times the input capacitance of a T/H transconductor. However, as track-hold unit i switches from transparent to hold state, and unit j undergoes the opposite state transition, a voltage error is injected onto the the isolation amplifier output node, which must be settled before the arrival of the next sampling instant, i.e. within $\frac{T_{clock}}{5}$. Since the speed requirement is increased more than the capacitive load is reduced, it is advantageous to design a single isolation amplifier for each track-hold phase. The current consumption of each isolation amplifier is $60\mu W$ to provide 8dB of gain with a 3dB bandwidth of 6.1GHz.

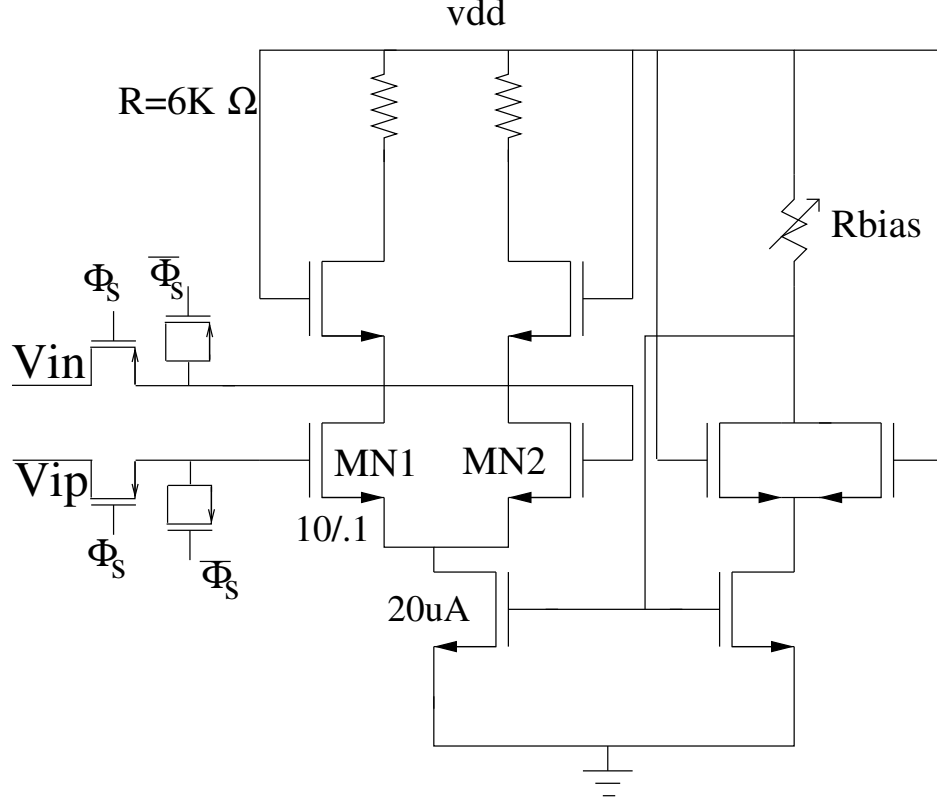


Figure 5.26. Track and Hold Amplifier

5.3.4 Comparators

The comparison chain is shown in figure 5.29. It comprises a regenerative latch, followed a metastability-preventing static SR latch and a retiming flip flop. Comparator outputs are retimed with the ϕ_1 clock, which gives every regenerative latch a time of $T_{ck}/5$ to regenerate. The regenerative latches were designed as sense-amplifier based latches ([65]). They were designed to resolve a 5mV input over a variable decision threshold within a $\frac{T_{ck}}{5} = 500pS$ timing interval. This results in a total voltage gain requirement of 150, or 100pS time constant. Since achieving such time constant over process variations was not possible in the given process, a preamplifier was added in front of the comparator to relax its sensitivity requirement and reduce kickback.

An important architectural feature of this design is that each lane is equipped with two separate comparators, which have decision thresholds situated symmetrically above and below ground. The decisions of the two comparators are combined by an XOR gate in the digital domain. During diode-mode front-end operation, the polarity of the incoming pulse is known a-priori and as a result, one of the comparator chains is clock gated to reduce power consumption. During mixer-mode operation however, the polarity of the incoming signal is unknown. Since the implemented symmetric comparison around 0 is then equivalent to performing an ideal absolute value rectification followed by thresholding, it reduces the IF bandwidth requirement in this mode by $\frac{1}{2T_p} = 320MHz$ (See figure 5.32 for an explanation). The comparator core is shown in figure 5.30. To minimize power dissipation, relatively small




Figure 5.27. Isolation Amplifier Schematic




Figure 5.28. Simulated FIR Filter frequency response for different settings

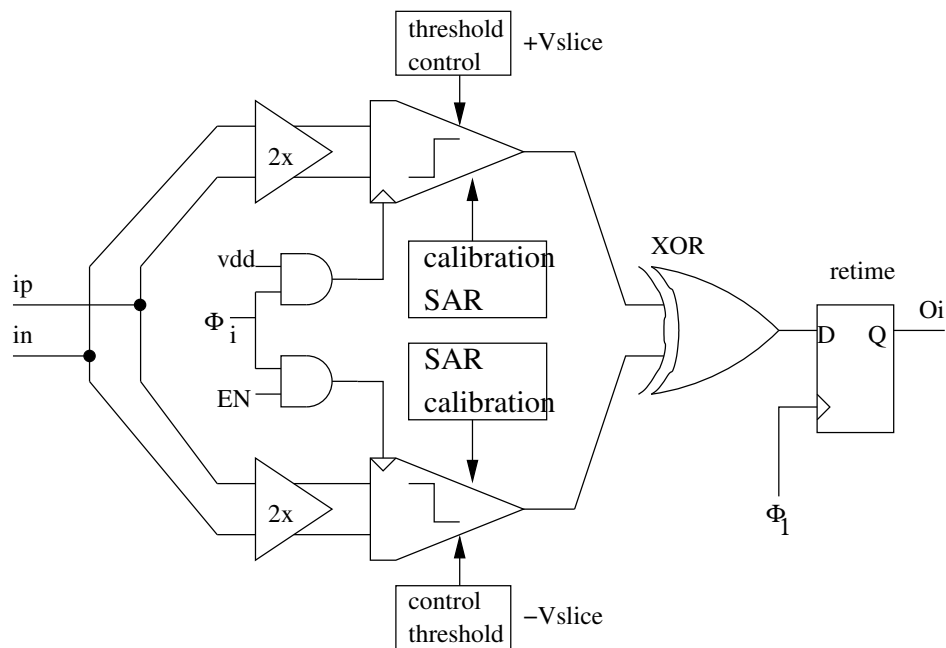


Figure 5.29. Block diagram of slicing chain

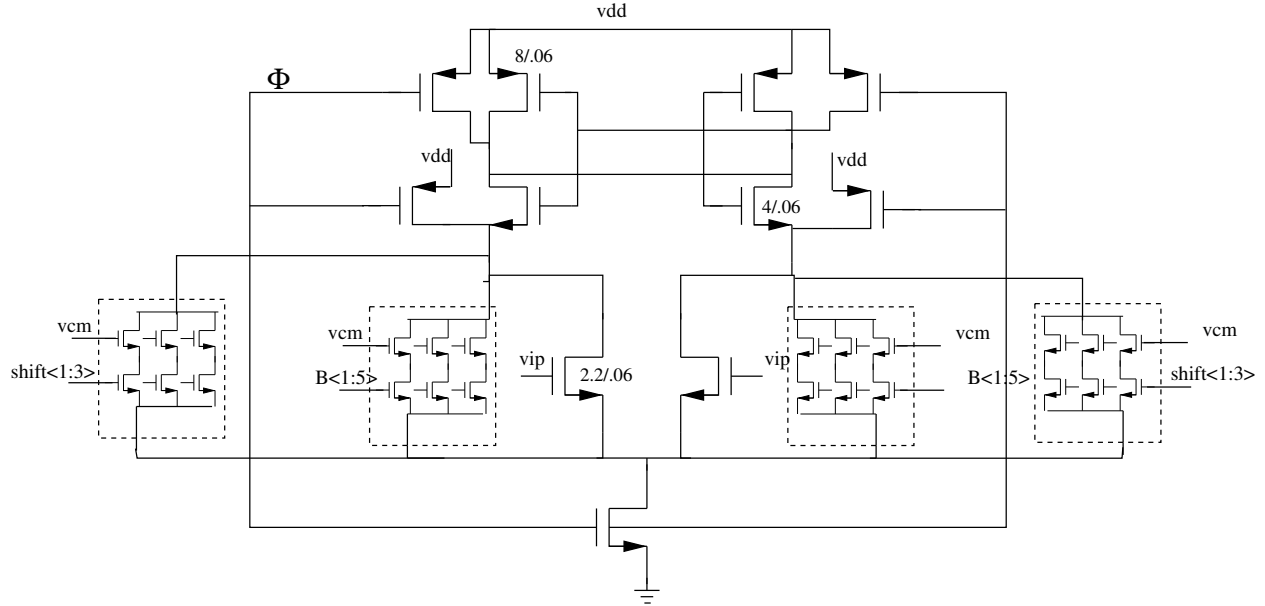


Figure 5.30. Annotated Comparator Schematic

devices were employed for all the transistors, requiring the addition a digital offset calibration loop on chip.

The decision threshold adjustment and the offset calibration are both performed by adding shunt resistors to either side of the sense-amplifier inputs. Note that the gates of the skewing devices are not connected to the comparator inputs, but rather to the supply voltage V_{dd} to prevent unbalanced loading of the preceding stage. Since the offset-correction devices are small and feedback monotonicity is necessary for calibration convergence, these devices are thermometer coded. However, the binary-to-thermometer decoder increases the loop delay above 1nS and the calibration is performed with a locally generated slow clock running at 125MHz ². After calibration, the comparator achieves an offset voltage smaller than 5mV . The distribution of the comparator offset, as well as the simulated relation between applied threshold setting code and the comparator offset, is shown in figure 5.31. The total comparator noise, referred to the preamplifier input, is simulated through periodic noise analysis to be 2mV_{rms} . When running at 500MHz , each preamplifier/comparator pair dissipates only $38\mu\text{W}$. This leads to a total power consumption in the comparators of $190\mu\text{W}$ for a 2.5GS/s sample rate in diode mode ($380\mu\text{W}$ in mixer mode), which maps into a competitive energy consumption of 76fJ/decision .

²The divider generating such calibration clock is gated when the calibration is not active by a calibration controller that is shared across all channels. The SAR control machine and state register are instead replicated for each comparator.




Figure 5.31. Comparator simulations: decision threshold versus digital code(left) and uncalibrated input-referred offset distribution (right)




Figure 5.32. Example of the effect of dual-thresholding and equivalence between dual thresholding and absolute value thresholding: the blue waveform, corresponding to $\delta F = 100 \text{ MHz}$ and initial phase of π , never crosses the positive threshold, and goes undetected by single-comparator scheme.

5.3.5 Layout and summary

The CAD layout of the baseband portion of the receiver is shown in figure 5.35, and a breakdown of its power consumption in table 5.4.

Output noise [1KHz-2GHz](V)	1.9mV
BasebandDC Gain/Peak Gain	6dB/13dB
Input Referred Noise	1 mV
Comparator Offset	$\leq 5mV$

Table 5.3. Baseband circuits performance breakdown

Block	Power Consumption
Isolation Amplifier	$300\mu W$
Track/Holds	$75\mu W$
FIR Filters	$100\mu W$
Comparators	$190/380\mu W$
Bias	$120\mu W$
Total	$775/965\mu W$

Table 5.4. Baseband circuits power consumption breakdown

5.3.6 Bias Circuits and power gating strategy

A unique challenge to UWB system concerns the implications of receiver power gating on the bias circuits. Bias circuits can either not be duty-cycled and then heavily decoupled-on chip, or they are duty-cycled with the receiver. Since our expected duty-cycling factor is of the order of 3% at 1MHz, a bias current mirror ratio of the order of 1:1000 is needed to make the bias circuit power consumption negligible when operating at 1MBps. This current mirror ratio is realizable, but it leads to inaccurate bias point setting. We chose therefore to duty-cycle the bias together with the circuits. Each circuit block is equipped with a local bias generator, realized with a current mirror with a gain of 4 that gets power gated. The bottleneck for bias settling is constituted by the bias node of the mixer-diode front-end, which is shown in figure 5.33. To prevent leakage from the RF port to the bias, an RC filter is inserted between the RF input and the current mirror. A blocking capacitor is further introduced between one port of the matching inductor and the diode gate to allow separate bias points. A low RC filter bandwidth is necessary for good isolation, while a large bandwidth is desired for fast bias settling. In addition, since the filter resistor R appears in parallel with the 700Ω inductor parallel resistance at resonance, and therefore must be much larger than this value to avoid degrading input match and voltage gain. Due to the presence of the RC filter, settling of the offset voltage on this stage exhibits a two-time constant

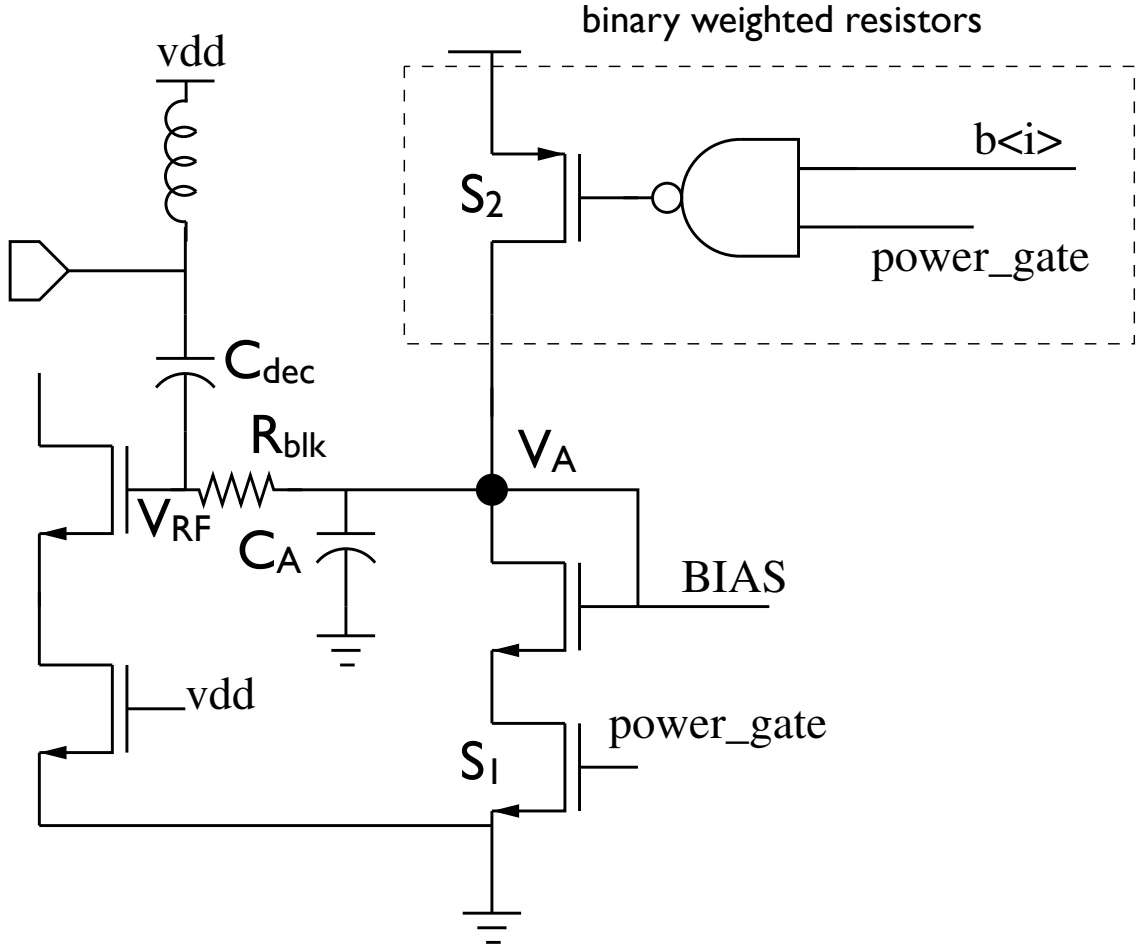


Figure 5.33. Schematic of bias circuit for RF front-end

behavior. At first, the bias resistor R_{blk} charges the bias capacitor C_A . Subsequently, this voltage propagates to the RF port of the mixer-diode through the delay set by the resistor R_{blk} and the capacitor C_{dec} . The values for R_{blk} and C_{blk} were optimized through simulations and set to respectively $10K\Omega$ and $500fF$, leading to a $5nS$ time constant for the filter settling. Settling of node V_A is then sped up by introducing a $500fF$ decoupling capacitor C_A , and power-gating switches S_1 and S_2 . During the receiver off-state, S_1 and S_2 are open and the voltage of node C_A is preserved, so that node C_A does not need to be recharged.

The switching transient is this way reduced to the duration of $13nS$, for a residual settling error of $2mV$ (See figure 5.34). Throughout the reset of the receiver, tail-current based biasing with source switching is adopted (See Fig.5.24)). This scheme has the advantage of low turn on time and isolation from the signal path. To minimize leakage, all the receiver circuits combine this technique with an additional header switch. LO drivers and TX circuits however are only clock gated to maximize their speed, and therefore contribute the bulk of the leakage




Figure 5.34. Simulation of bias settling:mixer differential output (left) and bias lines shown in Fig.5.33




Figure 5.35. CAD Layout of analog baseband

power.

To minimize complexity, all the bias circuits are enabled with a single control wire, EN_{Bias} , which is controlled by the timer described in the next sections. The complete power up circuit and sequence is shown in figure 5.36. Bias circuits and clock tree are enabled first. Since the impulse response of the FIR filter is 2 clock cycles long, the first two samples after the bias settling occurs should be disregarded. Therefore, the comparator outputs are not

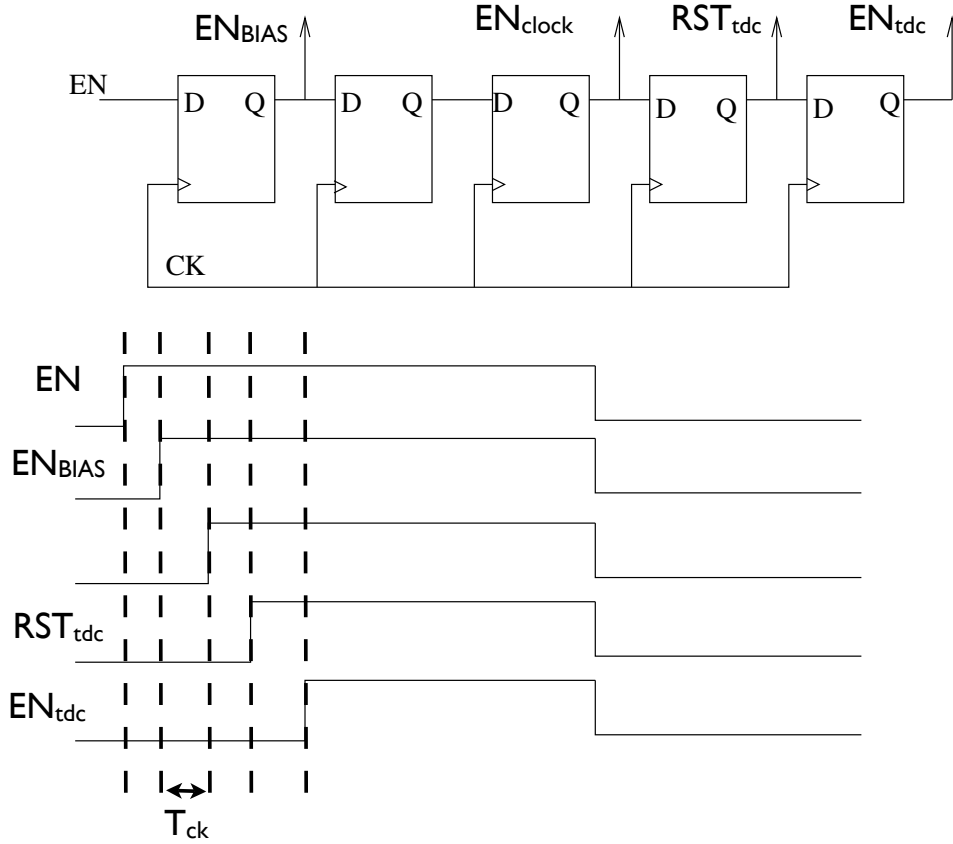


Figure 5.36. Power-on sequence

armed until 2 cycles after the FIR filters, which is turn is enabled 4 clock cycles after the bias circuitry.

5.4 Timing Circuits

The timing subsystem needs to generate the 400pS-spacing sampling clock as well as to track and align the receiver duty-cycling signal to the incoming transmitter data-stream. The architecture of this timer is shown in figure 5.37. It is comprised of a high-speed core, in turn containing the baseband DCO, counters and digital comparator; and a low-speed digital controller that processes the pulse inter-arrival times to produce an estimate of the transmitter clock frequency, demodulate the data and update the BER estimate.

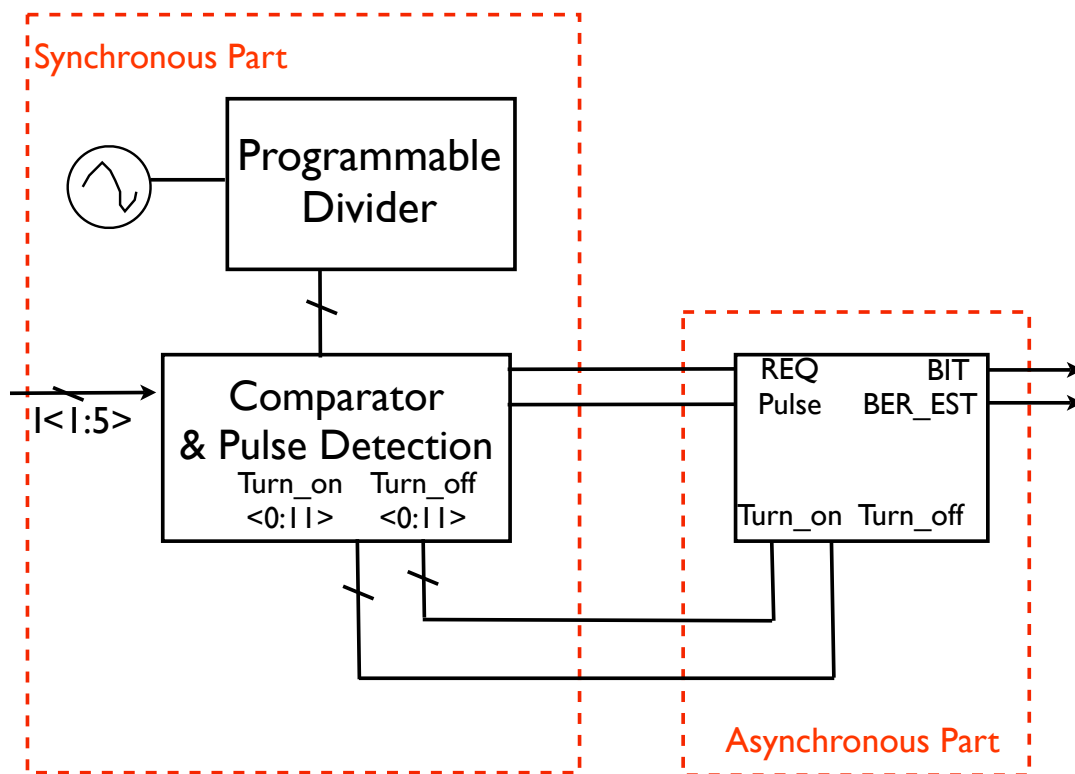


Figure 5.37. Synchronizer Architecture and Synchronous-Asynchronous Partition

5.4.1 Baseband oscillator

The system clock is generated from a 500MHz, 5 stages ring oscillator. We rely on positive clock-edges only to obtain the 400pS spacing required. The choice of the interleaving factor $N=5$ stems from a trade-off between design complexity and divider power consumption. A single-phase solution $N=1$ allows the simplest design, and demands no matching on the oscillator internal phases. However, it maximizes the speed of the divider and hence its power dissipation. When N is increased, the divider and oscillator speed requirements is decreased inversely, but matching between the ring oscillator delays starts to be required, leading to larger devices in the ring oscillator itself. At low supply voltage, when the sensitivity of delay to process variations is very high, designs operate close to the weak-inversion region and the range of N resulting in power savings becomes small. The chosen design point of 5-way interleaving allows the oscillator to run on an internal supply of 550mV, which even for small devices (See Fig. 5.38) results in a mean-to-standard deviation for the delay of about 10, and guarantees over 5x power savings in the divider. The oscillator frequency is digitally controlled through symmetric resistive degeneration on both supply rails, and delay cells can be individually adjusted by a 3 bit digital control word to further improve phase matching (Again, see Fig.5.38). The nominal power supply of the oscillator is set to 750mV, as required by the divider chain and the comparator. Under these conditions, the oscillator core power dissipation is $4.5\mu W$.

In order to minimize system power dissipation, clock gating is used aggressively in the whole digital baseband. To insert the clock gates as close as possible to the root, all of the oscillator output phases are symmetrically loaded. For phases 1, 2, 4, 5 this load consists of a NAND gate used for clock gating, and a second NAND gate with a grounded input. This second NAND gate is used to equalize the load to that of phase 3, which feeds the divider through an identical gate with an input tied to V_{dd} .

Switching the clock phase drivers on and off changes the capacitive load and hence the clock frequency. Suppose now that the oscillator frequency when the clock tree is disabled is $f_{osc} + \Delta F$, while it equals f_{osc} when the clock tree is enabled; and the estimated pulse reception window equals $\{\frac{N}{f_{osc}}; \frac{N+K}{f_{osc}}\}$. Since during the acquisition phase, the clock-tree is always active, the first pulse reception periods after power gating is enabled will have an inaccurate positioning. In fact, the first incoming pulse after duty-cycling is enabled will occur $\frac{N}{f_{osc} + \Delta F}$, or it will be early (assuming ΔF is positive) by $\frac{N\Delta F}{f_{osc}}$ clock cycles. However, if the receiver duty-cycling window is wide enough that the pulses get received however, the estimation loop tracks the new value.

As a result, the minimum pulse-reception window width should be increased by $N\frac{\Delta F}{f_{osc}}$ at the beginning of the duty-cycling period to accomodate this change. To minimize this effect, the oscillator output phases are connected to the bottom transistor of the N-MOS NAND stack. This way, capacitance variation is only due to variations in C_{gd} modulation. This limits $\frac{\Delta F}{f_{osc}}$ to less than 1% and the increase in duty-cycling width to 5 clock cycles.

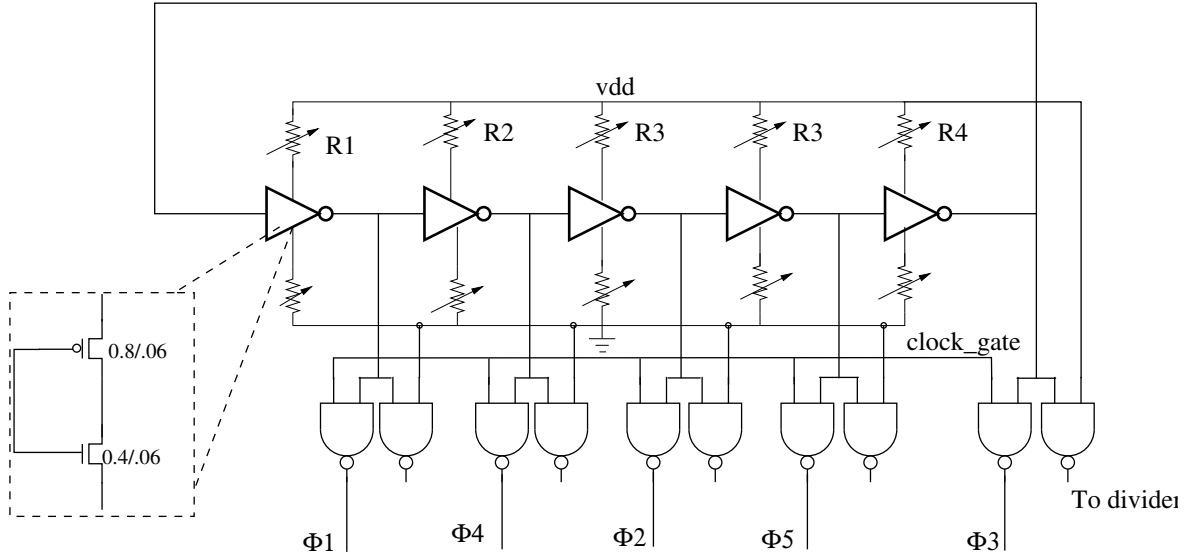


Figure 5.38. Baseband DCO and driver schematic

5.4.2 Baseband divider and comparator

Divider speed is critical as to ensure correct start and stop time detection, the divider output codeword and the comparator have to settle within a clock cycle ($2nS$) (Fig.5.39). To minimize divider latency, a fully synchronous counter would be desirable. However, this leads to excessive clock load and power dissipation. The adopted solution is shown in figure 5.39. A TSPC divide by 2 prescaler feeds a synchronous divide by 2 TSPC counter. The output of this counter is the clock input to a 9-bit fully synchronous divider, which is realized using static CMOS. The complete 12-bit counter output is fed to a static digital comparator together with the start and stop times of the duty cycling window. The complete settling time of the counter is 800pS, leaving 1.2nS available for the comparator to resolve; while its average power dissipation is $28\mu W$ at 0.75V after parasitic extraction.

The divider reset logic is shown in figure 5.40. The external reset signal is first retimed to the 500MHz ring clock. Then, flip-flop DF1 and logic form a 1-clock cycle wide pulse that is used to sample the output of the counter in a parallel fashion. In the subsequent clock-cycles, the counter is synchronously reset. The latency introduced in the reset process introduces bias in the period estimate is compensated in the digital controller block.

The digital comparator is built as a pair of tree-based 12-way XOR (See figure 5.41). When the counter output becomes equal to the lower turn-on threshold, the output of the first XOR tree goes high the set input of SR latch SR1, transforming this transition into a level that is held until the reset input of SR1 is asserted by the second tree. Because the inputs from the counter settle at different times through the clock period, the output of the comparator must be retimed to avoid glitching. This is accomplished by flip-flops

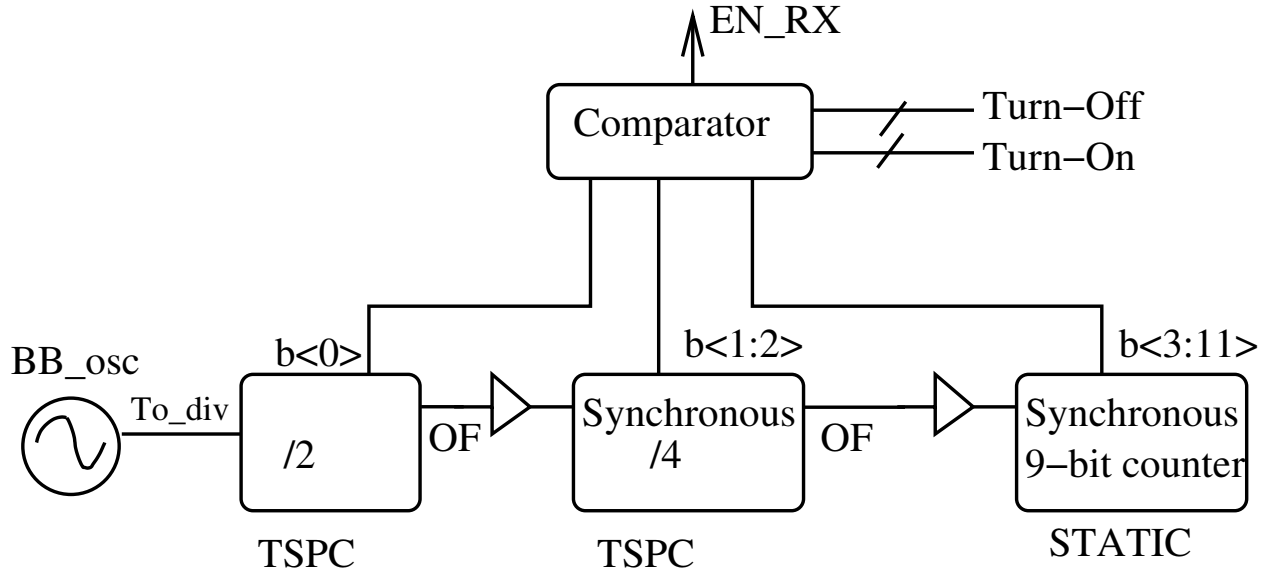


Figure 5.39. Baseband divider architecture

FD1 and FD2. Since the comparator speed is also critical, more registers could be added to pipeline its operation and relax the speed constraint. However, since we are operating under the assumption of a constant 0.75 V supply voltage, such pipelining registers would also increase power dissipation by increasing the clock load and are therefore omitted. The extracted comparator critical path delay, including the t_{su} of FD1 and FD2, is 1nS on the typical process corner.

5.4.3 Pulse detection and retiming logic

A drawback of the chosen architecture is that since the divider can be reset in any state, only the 500MHz ring output can be used as a system clock. This increases the power dissipation of a synchronous architecture considerably, since all the communication registers between the controller and the high-speed core (counter+comparator) need to be clocked at this rate. This power consumption can be reduced by implementing an asynchronous interface between the high-speed core and the controller. The operation of this interface is based on a three wire interface between the high-speed core and the controller. When a counter reset-event is generated, a request for service is issued to the controller block, along with an active-high bit indicating whether a pulse has been detected or not. This request for service acts as a clock signal for the signal processing blocks in the controller, and is also fed to a replica-delay line. The output of this delay line acts as an ACK signal for the high-speed core, which is then re-armed and can issue a new request.

The schematic of the interface blocks is shown in figure 5.42. The request to service can be issued either because a valid pulse was received (ones of the comparator outputs goes low), or because the upper limit of the duty-cycling window has been reached without a valid detection. The request to service event is converted to a level by SR latch SR2, and passed

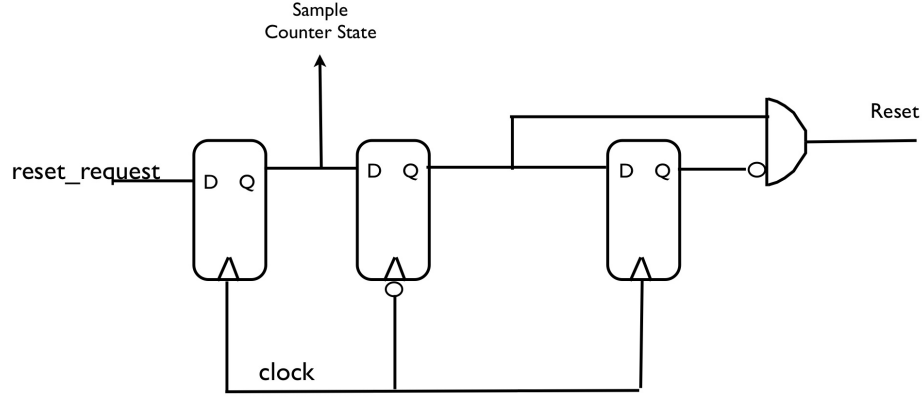


Figure 5.40. Counter Reset Logic

to the controller, together with the output of the C-element, that will be high if a pulse has been received. Both these sequential elements are reset by the ACK signal mentioned above. Finally, figure 5.43 shows the digital logic implementing the pulse detection function. The output of the slicers in each lane are ORed together through 2 stages of logic to generate an event detection. When the event detection goes high, the bank of D-FF DF1:DF5 freeze the state of the comparators, which is used by the controller to calculate the fractional part of the inter-pulse interval. Next, the output of the C-element mentioned in the previous paragraph is raised, initiating the request to service process.

5.4.4 Controller

The digital controller , performing IIR filtering for frequency estimation, data demodulation and B.E.R. calculation has been coded to VHDL and mapped onto the standard cell library by John Crossley of U.C. Berkeley, starting from the simulink description used in behavioral simulations. To accomodate asynchronous operation, the design is forced to complete all its calculations within 1 clock cycle(as only one request to service signal is generated by the high-speed core)and synthesized as a synchronous block. The request to service signal is connected to then connected to the clock-port of this block. To guarantee timing closure of the synthesized block at $0.75V$, the timing constraints of the synthesis were tightened by a factor equal to the ratio on an inverter delay at $V_{dd} = 0.75V$, to the delay of the same inverter at $V_{dd} = 1.2V$, which for this process is about 4. This digital block It occupies an area of $100\mu m \times 100\mu m$ and consumes approximately $5\mu W$ from the $0.75V$ supply when running at 1Mbps.




Figure 5.41. Binary comparator cell

5.4.5 Block Layout and summary

The timer layout is shown in figure 5.44., and the power consumption breakdown is given in Table 5.5. In order to minimize routing parasitics, blocks were placed as close as possible. Yet, extracted simulations show over 50% increase in dissipation compared to simulation, largely due to the non-optimized layouts of the static gates used. The power consumption of the whole synchronizer remains however over a factor of 2 lower than that of the PLL reported in [11].

Block	Power Consumption
Oscillator	$6\mu W$
Divider	$28\mu W$
Digital Logic	$5\mu W$
Total	$39\mu W$

Table 5.5. Synchronizer power consumption breakdown

5.5 System Layout and summary

The entire transceiver layout is shown in figure 5.45. The system occupies an area of $600\mu m \times 600\mu m$ and is comprised of 32500 transistors .

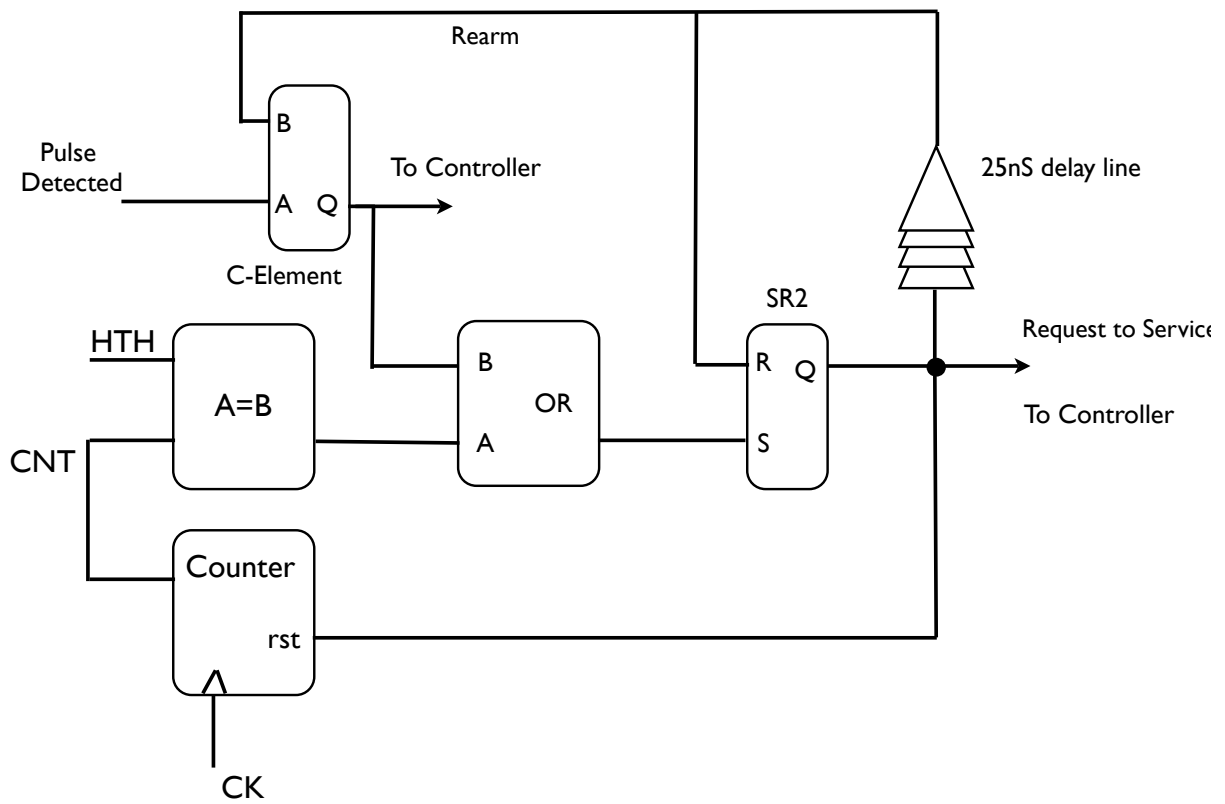


Figure 5.42. Asynchronous interface building blocks

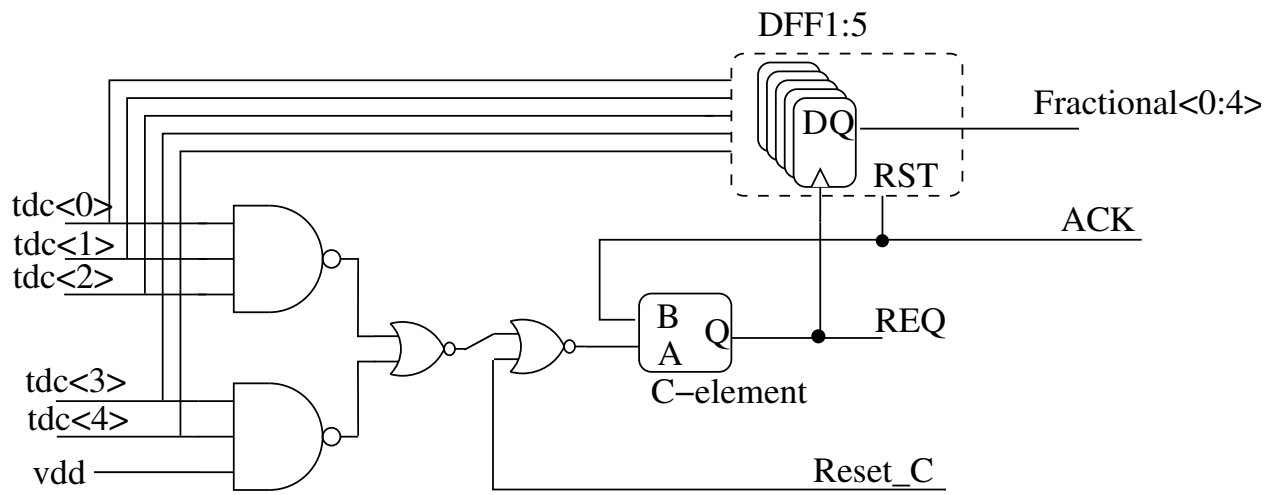


Figure 5.43. Pulse Detection Logic

Tuning Range	240KHz-16MHz
Output jitter [pk-pk]	$\leq 2nS$
Period resolution	$2nS$
Acquisition time	$\leq 32updates$

Table 5.6. Synchronizer performance summary

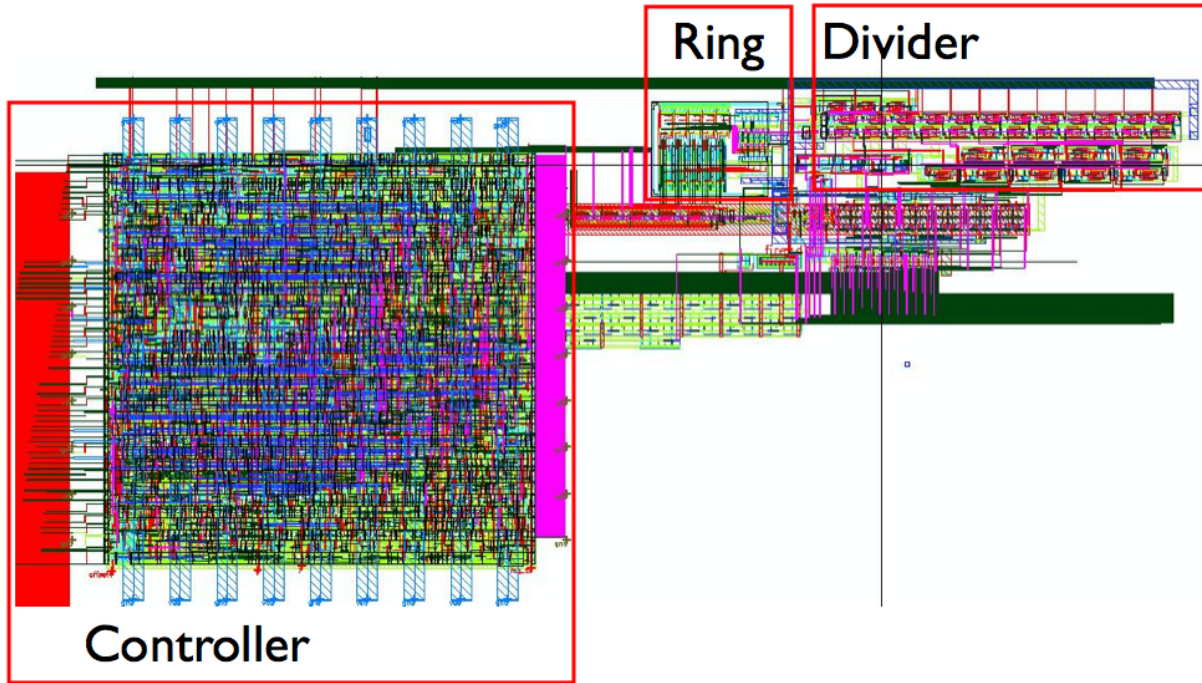


Figure 5.44. CAD Layout of timer




Figure 5.45. CAD Layout of transceiver

The matching network inductor, and the transmitter decoupling capacitance are the largest area consumers. This large decap is needed on the transmitter V_{dd} due to the high peak current flowing during transmission. This decoupling capacitor is built using gate capacitance from 2.5V devices to prevent gate-leakage. The expected performance and the power dissipation breakdown for this chip are reported in table 5.7.

Sensitivity	-24dBm
Output power	5dBm
Maximum interference power	-30dBm
RF Power Consumption (on state)	600 μ W
Baseband Power Consumption (on state)	1mW
Transmitter energy consumption	28 pJ/bit
Timer power consumption	40 μ W
Expected receiver duty-cycling window width	28nS
Expected duty cycled RX power consumption	90 μ W

Table 5.7. Transceiver performance summary

Chapter 6

Measurement Results

The transceiver was implemented in a 65nm 6M1P CMOS process with LP option and dual-threshold devices. The die photographs is shown in figure 6.1. The chip has an area of $1.4 \times 1.7mm$ limited by IOs. In addition to the main transceiver (top left), standalone blocks are also implemented for characterization. A replica of the timing recovery loop is implemented in the bottom left corner of the chip, while the top right corner contains a standalone baseband subsystem. A replica of the RF front-end and matching network is laid out in the bottom right corner of the chip. The chip is controlled by an on-chip 512-bit wide SPI controller, which is interfaced to the outside world via a 3-wire interface and controlled by a PC through a Python script and an *AARDVARKTM* adapter. All the testing was performed with the chip is directly attached mounted on an FR4 board using a COB technique in order to reduce parasitics and facilitate testing.

6.1 FIB Issue

Due to a layout error, the secondary port of matching network inductor $L1$ was not connected to ground (See figure 6.2). This port was therefore connected to ground through a focused-ion-beam (FIB) tungsten deposition. The expected effect of the FIB is to add 10Ω of series resistance to the integrated inductor reducing its Q value to 3.76. The quality factor reduction degrades input matching, transmitter efficiency and sensitivity. According to simulations, the S_{11} minimum should increase to $-8dB$, while the voltage gain should be reduced to $1.1dB$ from the initial $5dB$. Furthermore, the impedance seen from the transmitter port is simulated at $23 - j36\Omega$, degrading the matching network efficiency to 47%. The simulated reflection coefficient however does not match the measured values (See Fig.6.3), and appear downshifted in frequency by about 600MHz. Using gating to eliminate the reflections at the connector results in better S_{11} ,but does not shift the center frequency.




Figure 6.1. Die Photo

This seems to indicate that the center frequency shift is due to on-chip component variations, and not to board parasitics. However, since the exact influence of the FIB is not known, it is hard to draw conclusions.

6.2 Transmitter measurements

Transmitter measurements were taken after FIB on the complete transceiver. Figure 6.4 show the time-domain TX output waveform and power spectrum captured by a 20GS/s Agilent Infiniium real-time scope. The oscilloscope bandwidth is 6GHz, and as a result the spectral content of the pulse above this frequency is strongly attenuated. The measurement shows that the peak in the transmitted pulse spectral density is much lower than the FCC limit of -41.6dBm/MHz (-61.2dBm/MHz-). The spurious tones that violate the mask in the GPS band are not due to the transmitter, but originate from clocks inside the oscilloscope itself, as shown in the inset showing the noise floor. For the same sample, the transmitter power dissipation and energy/pulse are shown in figure 6.5. At 1Mbps, the transmitter energy consumption is 26J/bit. Fitting a linear relationship $P_{tx} = P_{leak} + E_{act}f_{rep}$ to the measured data results in $P_{leak} = 1.7\mu W$, $E_{act} = 25.4pJ$ ¹. The output energy per pulse, de-embedding 3dB of measured cable loss, is 2.880pJ, resulting in a transmitter efficiency

¹Leakage power increases at lower pulse-repetition rates as a result of dynamic gates starting to loose their states

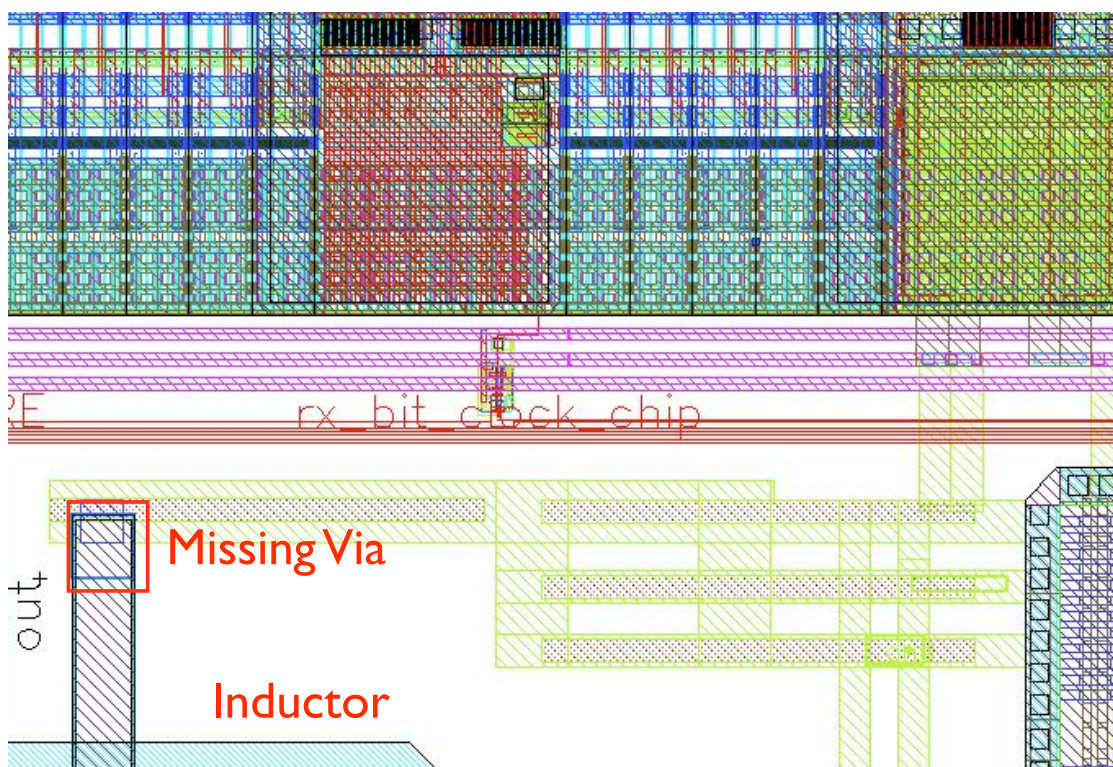


Figure 6.2. Location of missing connection




Figure 6.3. Measured input reflection varying digital tuning of matching network capacitor C_2




Figure 6.4. Measured Transmitter Output Waveform and Spectrum




Figure 6.5. Fibbed Transmitter Power dissipation (top) and energy/pulse (bottom)

of 11.6%. The simulated transmitter output power matches this measurement within 10% once the 10Ω extra loss in the matching network inductor is taken into account.

The tuning curve of the RF DCO is shown finally in figure 6.6 along with simulated data. Out of the 256 possible digital codes, only a fraction could be measured due to frequency limitations in the antenna driver, which fails to operate above 6.5GHz. The measured data shows good correlation with the post-layout simulations of the oscillator structure, except for the points around tuning code 257, which have a much higher oscillation frequency in simulation.

6.3 Receiver Measurements

6.3.1 Sensitivity Measurements

The receiver sensitivity measurements were also conducted on a FIB-bed part. First, the performance was characterized by a wired test. An Agilent 81134A Pulse-Patter generator was used to generate 2nS wide pulses, which were then up-converted with a mini-circuits ZX-05762H-S+ passive mixer. The receiver input power was calibrated for cable losses using the real-time oscilloscope. During these measurements, we discovered that the receiver baseband stages were being saturated by DC offset voltages propagated by the diode/mixer and the isolation and track/hold amplifiers. This issue was under-estimated at design time: offsets for the individual stages were calculated but , it was assumed that the high-pass characteristic of




Figure 6.6. RF DCO Tuning Curve

the FIR filter would block their propagation. In reality, the offsets reaching the FIR filter are large enough that the filter is saturated. Since different offsets appear on different channels, a fraction of the comparators operate with limited input swing and unacceptable BER. This effectively decreases the receiver sampling rate and demands the input pulse duration to be increased to 2nS. Furthermore, since all channels operate with a significant residual offset, the test was performed using the dual-comparator option, increasing the baseband power consumption by $190\mu W$ with respect to the nominal value. Finally, optimum sensitivity performance was obtained when the FIR filter was configured to have [1; -1] response. The baseband gain and comparator thresholds were set to achieve a false-alarm rate lower than 10^{-4} , and the bit-error rate was then measured by feeding the receiver with a repetitive pulse train and measuring the missed-detection probability. The waterfall curves of the receiver under these conditions are reported in figure 6.7. The measured peak power sensitivity in diode mode is -16.5dBm, and it improves to -24.5dBm in mixer mode. These values respectively correspond to 47mV and 18.7mV input voltage at the antenna, and to duty-cycled power dissipations of -43.5dBm and -51.5dBm. Compared to simulations, a sensitivity degradation of approximately 8dB is observed. 5dB can be attributed to the increased losses in the matching network due to the FIB. The remaining 3dB loss is likely due to the degraded baseband noise performance resulting from the offset accumulation and saturation.




Figure 6.7. Measured received waterfall curve in wired test

6.3.2 Duty-cycled Measurements

Duty cycling of the receiver was also measured. In the measured parts however, the digital back-end is capable of maintaining lock conditions only for a few hundreds of pulse transmissions. Testing was therefore done by manually over-riding the loop through the SPI. Figure 6.8 shows the measured receiver power dissipation and missed-detection rate of the receiver as a function of the receiver duty-cycling factor. A minimum on-time of $72nS$ per pulse reception interval is required by the receiver to obtain reliable reception, with a window width limited by slower-than expected transients on the bias lines . At this width, the total receiver duty-cycled consumption at 1Mbps is $275\mu W$ in diode mode, reduced by a factor of 7 only from the peak $1.7mW$ consumption. This result indicates that the bias network should not be duty-cycled, but rather scaled, in order to reduce the receiver turn-on time and hence the total power dissipation. Alternatively, we could stagger the power-on of the bias network with respect to that of the receiver by changing the digital duty-cycle controller .

6.3.3 Interference Rejection Measurements

The interference measurements could only be performed on a non-FIBbed part, and using a configuration that is significantly different from that used for sensitivity measurement. In particular, to prevent the accumulated offset from saturating the baseband stages, the gain of the isolation amplifiers and track and holds was reduced, degrading sensitivity. In addition,




Figure 6.8. Measured power dissipation and error rate as a function of duty-cycling window width

using two comparators per channel in the presence of input offset results in unbalanced thresholds. In the presence of LO feedthrough, the comparator with the smaller noise margin trips, resulting in false positive errors and large interference sensitivity. As a result, a single-comparator per channel was active during this measurement. The combination of decreased gain and effective sample rate also forced the input pulse duration to be increased to 8nS. The desired signal is summed to a continuous wave interferer by a passive combiner and fed to the input of the chip. The measured false-negative rate in diode mode is shown in figure 6.9 for two different input frequencies. As seen in the figure, the same interferer rejection behavior is obtained for equal values of interference offset from the desired signal carrier, showing that the behavior is not set by the RF pre-filter selectivity. From the graph, we can obtain that the effective RF receiver bandwidth is (for 10^{-3} error rate) is approximately 100MHz, corresponding to a rotation of 1.5π during the pulse-period. The expected bandwidth from behavioral simulations is 70MHz, which is reasonably close given the many approximations made to arrive at this conclusion. In figure 6.10, we show the behavior of missed-detection rate versus Signal-to-Interferer-Ratio (SIR) for different offset frequencies. As expected, this curve is monotonic for interferers that are close-in ($\Delta F = 25MHz$), while the curve corresponding to $\Delta F = 125MHz$ shows a maximum BER degradation around SIR=6dB as predicted by our analysis. The same tests were also run for the receiver operating in mixer mode, although the input levels had to be adjusted to prevent saturation. The combined interference tolerance mask of the mixer/diode combo is shown in figure 6.11. Note that since anti-podal signaling is not employed in this test, the mixer-mode performance is degraded by 6dB. We see however that the mixer interference sensitivity is complementary to that of the diode, and that except for the measurement at $\Delta F = -50MHz$, an improvement of




Figure 6.9. Measured false alarm rate versus interferer carrier frequency in diode mode for two different values of desired signal carrier




Figure 6.10. Measured missed-detection rate as a function of Signal-to-Interferer Ratio for different offset frequencies




Figure 6.11. Required SIR versus offset frequency to achieve $5e-3$ B.E.R. for diode and mixer mode

over 4dB is obtained in in-band interference sensitivity, at the price of out-band interference sensitivity.

6.3.4 Wireless Tests

We verified the system through wireless test using commercial antennas . In a first test the transmitter was emulated using the same setup used in wired transmissions, comprising a sinusoidal signal source, an Agilent 81134A pulse generator and a Mini-Circuits up-conversion mixer. The transmitter was adjusted for an output power of 5dBm and then its position relative to the receiver was swept and the false-negative rate recorded. Antennas in the TX and RX were kept oriented for maximum gain (6dBi at 6GHz). The measured results are shown in figure 6.14. While the receiver has a worse sensitivity than the desired value, we can still obtain communication at 4cm range exploiting the antenna directive gain (6dBi). This test was conducted in the BWRC lab, in the presence of EMI from instrumentation and laptop computers and shows that the low receiver sensitivity in diode mode guarantees a good robustness. The same measurement could not be reproduced in mixer mode, because as a result of the improved sensitivity, many out-of-band interferers cause spurious detections.

Figure 6.15 shows the same wireless tests, measured with a chip configured as transmitter and one as receivers (Setup shown in figure 6.13). Both chips were fibbed, and as a consequence, their matching-network peak gain frequencies are offset. This results in a further range reduction when compared to the test performed with the instrumentation transmitter, where the TX carrier could be finely adjusted for maximum gain. The measured range in diode mode is 1.25cm, and it improves to 3.75cm in mixer mode.




Figure 6.12. Board picture and Skycross and custom monopole antennas




Figure 6.13. Test Setup




Figure 6.14. Required SIR versus offset frequency to achieve $5e-3$ B.E.R. for diode and mixer mode




Figure 6.15. Wireless Link Measurement with TX-RX pair

6.3.5 Receiver Performance Summary

The measured receiver performance is summarized and compared to simulated values in Table 6.2. Table 6.1 reports a comparison between this work and the one reported in [11], operating in the 0-1GHz band. Optimizing the link for high-gain results in an active power dissipation reduction of a factor ≈ 2 and in a greater than 1.5x range extension². In addition, the low gain required, combined with the dual-mode operation improves significantly the required interferer tolerance. The proposed timer also obtains over a factor of 2 lower power than the PLL-based solution in [11]. At 250Kbps, the total link energy consumption

$$E_{Link} = \frac{P_{tx} + P_{timer}}{f_{clk}} + \frac{P_{rx} + P_{timer}}{f_{clk}}$$

is also reduced by a factor of 2.

As mentioned above, the degraded performance compared to the simulated results is due to the higher-than expected startup time and input-referred offsets. A redesign should focus on reducing these issues. The most effective way to reduce the offset of the mixer/diode is to employ a calibration DAC feeding the output of this stage. This option was not implemented for the increased design and testing complexity, but is preferable over employing AC coupling capacitors as it guarantees faster startup. The bias startup could be sped up employing a staggered bias-core power-up sequence, or changing the biasing scheme of the diode-mixer itself. One possibility that was explored in the design, but not pursued, was that of using the matching network inductor as a bias choke (See figure 6.16). If this choice is made, care must be taken so that the impedance of the bias network at RF is small enough that no RF-bias leakage is present, and that the quality factor of the inductor is not degraded. Simulations show that configuration shown in figure, consisting of a 15pF decoupling capacitor coupled with an R-C low pass at 200MHz serves this purpose. With these adjustments in place, the receiver could achieve a reduction of over a factor of 2 in power dissipation, while further enhancing the interferer robustness due to the reduced listening window.

6.3.6 Comparison with prior art

While the work in [11] constitutes the most fair comparison point for this system (as it also targets a wireless range of 5 cm) , in order to put this work in context we also compare this system with other low power ultra-wideband transceivers. This comparison is reported in Tables 6.3 and 6.4. Since these comparisons are performed across different systems, they must be analyzed carefully. Concerning the receivers ,it is clear that two main trends exist in these publications: about half of the surveyed ([25],[37],[9]) do not utilize aggressive duty-cycling of the RF front end. This results in good power efficiency at high pulse repetition rate ($\approx 16Mbps$). The energy efficiency however is not scalable to lower pulse repetition rates. In addition, while the use of spreading ([25],[37]) enables superior interference robustness, it comes at a proportional cost in energy/bit.

²Referring to the case of wireless link using two fibbed chips. Range improvement assuming aligned resonant frequencies is actually 5x

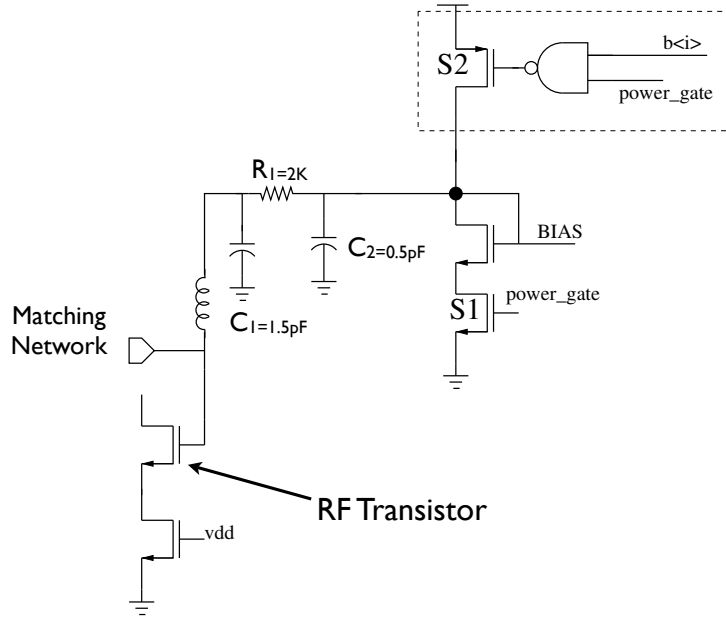


Figure 6.16. Alternate Bias Scheme

Parameter	Measured Value	Simulated Value
Sensitivity	$-16.5dBm / -22.5dBm$	$-24dBm / -32dBm$
Transmitter output power	$3dBm$	$5dBm$
Max. interference power	$-29.5dBm$	$-30dBm$
RF Power Consumption (On State)	$600\mu W / 3.6mW$	$600\mu W / 3.6mW$
Baseband Power Consumption (On State)	$1.1mW$	$1mW$
Duty Cycling window width	$72nS$	$28nS$
RX Average Power ³	$254\mu W$	$90\mu W$
RX Average Power ⁴	$674\mu W$	$257\mu W$
TX Average Power	$26\mu W$	$28\mu W$
Timer Power	$50\mu W$	$40\mu W$
Range	$\leq 1cm$	$1.5 \rightarrow 3.75cm$

Table 6.1. Summary of measured transceiver performance

Parameter	[11]	This work
Modulation Scheme	BPSK	PPM
Operating Band	0-1GHz	5.6GHz
Sensitivity	$-70dBm$	$-16.5dBm / -22.5dBm$
TX Radiated Power	$3dBm$	$5dBm$
Maximum interference power	$-83dBm$	$-35dBm$
Power Consumption (On State, Diode Mode)	$3.6mW$	$1.7mW$
Duty Cycling window width	50nS	72nS
RX Average Power (Diode Mode, 250Kbps)	$180\mu W$	$62.5\mu W$
TX Average Power	$6.25\mu W$	$6.25\mu W$
Timer Power	$72\mu W$	$40\mu W$
Link energy	1.28nJ/bit	0.61 nJ/Bit

Table 6.2. Transceiver Performance Comparison with 0-1GHz system [11]

Systems [28] and [9] as well this work do not use spreading gain, and as a result suffer from a degraded interferer tolerance, but achieve lower energy consumption. Amongst such receivers, the one presented in this work achieves the best robustness to in-band interference tolerance, measured both as absolute interference power ($-30dBm$) and as peak-signal-to-interferer power ($-13dB$). In addition, notice that receivers [28] and [9] have high peak power consumption. Operation at lower data-rate therefore requires a tighter timing control which could be difficult to achieve.

A similar comparison is also performed for the transmitter in Tab.6.4. Again, care should be taken in comparing results from hardware designed for different operating conditions. In general, transmitters designed for 16Mbps operation require a tighter spectral mask compared to transmitters operating at low data rate control to avoid violations to the FCC mask. In addition, transmitters operating in the higher frequency bands (5-10GHz) require higher power than ones operating in the 3-5GHz band. The designed transmitter exploits the low transmission duty cycle to improve the energy efficiency by over a factor of 2 compared to state-of-the art.

6.4 Timer Measurements

The issues with the digital part locking were not observed in the stand-alone timer testblock, which was therefore used to characterize the performance of the timer section of the transceiver and the operation of the proposed BER estimation algorithm. Figure 6.18 shows the measured jitter histogram of the timer while operating in TX mode at 520MHz internal clock rate and powered by a switching supply. When running at 520MHz internal clock rate, the measured cycle-to-cycle jitter standard deviation is 160pS, and the peak-to-peak value 1.8nS (corresponding to 320ppm). The N-cycle jitter increases linearly with

Parameter	[25]	[37]	[27]	[9]	This work
CMOS Node	$0.18\mu m$	$0.13\mu m$	90nm	90nm	LP-65nm
Operating Band	3-5GHz	0-1GHz	3-5GHz	3-5GHz	5.6GHz
Sensitivity (1Mbps)	-84dBm	-80dBm	$-90dBm$	-90dBm	-44dBm/-52.5dBm
Max.interference power	-32dBm	N/A	-83dBm	N/A	-35dBm
Pulse Rate	16Mpps	32Mpps	100Kpps	16Mpps	1Mpps
Spreading Gain	32	32	1	1	1
Bit Rate	500Kbps	1Mbps	100Kbps	16Mbps	1 Mbps
Peak Power Diss.	N/A	3.3mW	35.8mW	8mW-22mW	1.7mW-4.6mW
Average Power Diss.	36mW	3.3mW	$250\mu W$	8mW-22mW	$300\mu W$
Energy/Pulse	2.25nJ	110pJ	2.5nJ	500pJ	300pJ
Energy/Bit	72nJ	3.3nJ	2.5nJ	500pJ	300pJ
Includes Timer	No	No	No	Yes	Yes

Table 6.3. Performance Comparison with published UWB Receivers

Parameter	[16]	[13]	[47]	This work
Operating Band	3-10GHz	3-5GHz	3-5GHz	5.6GHz
CMOS Node	90nm	90nm	90nm	LP-65nm
Pulse Rate	16MPps	16.7MPps	16.7MPps	1MPps
Radiated Energy/Pulse	0.2pJ	2.7pJ	260fJ	2.88pJ
Total Energy/Pulse	40pJ	47pJ	19pJ	25pJ
Leakage Power	N/A	$98\mu W$	$180\mu W$	$1.7\mu W$
Efficiency	0.5%	5.74%	1.37%	11.3%

Table 6.4. Transmitter Performance Comparison with published UWB Receivers




Figure 6.17. Measured Timer Power Dissipation versus timing resolution at different supply voltages

the number of cycles N , limited by power supply and circuit $1/f$ noise. The timer was also fed a PPM-modulated waveform, to characterize its robustness in the presence of a noisy input phase signal and verify the performance of the proposed conditionally recirculating architecture. The histogram in figure 6.19 displays the measured distribution of the time interval between which the receiver duty-cycling window is opened, and the pulse arrival time. The nominal opening of the receive window is set to 5nS before the estimated pulse arrival time. The measurement reveals a mean delay of 6.2nS, with a 1.3nS variance (.26%)

Figure 6.17 shows the timer power dissipation as a function of the timing resolution for an output frequency of 2MHz(i.e. the inverse of the achievable clock rate). At the nominal power supply of 0.75V and clock rate of 400MHz, the measured power dissipation is $50\mu W$. Reducing the supply voltage degrades the timing resolution but reduces power dissipation. At 0.65V, an internal clock rate of 300MHz (corresponding to an effective baseband sampling rate of 1.5GHz) is obtained at a power dissipation of $26\mu W$. At 0.5V, the power dissipation is reduced to $\approx 3\mu W$, but the timing resolution is limited to 50nS ($\pm 10\%$ of the 500nS wake-up interval). Since the bulk of the timer power is consumed in the divider, an optimized re-design would push the design towards increasing the number of oscillator stages to decrease the divider clock rate.

To verify the B.E.R. estimation functionality, the design was fed with a forbidden sequence and the B.E.R. flag output was observed. The result is shown in figure 6.20, which shows how the occurrence of the forbidden sequence $\{1/2, 1, 1/2\}$ triggers the error detection flag. Data demodulation also operates correctly.




Figure 6.18. Transmit mode cycle-to-cycle jitter




Figure 6.19. Measured distribution of window opening to pulse arrival delays




Figure 6.20. Demonstration of error-detection functionality

Chapter 7

Conclusions

To summarize, this thesis presented a design methodology for cm-range wireless data network. The contributions of this thesis can be summarized as follows:

- The development and validation of a model for cm-range wireless channels, enabling the system designer to choose a channel-loss optimal transmission frequency , and the experimental verification of this framework.

By applying this methodology to the design of a wireless link with 5cm range and 1cm^2 antennas, we showed that the use of a carrier frequency such that antennas operate in the *radiating near field* results in minimum path loss and hence in a lower power and more integrated solution compared to operating in the 0-1GHz frequency band.

- The design of a communication scheme based on PPM modulation that merges the synchronizer and the data slicer portions of the back-end, and enables online bit-error rate estimation.

- The development of an interference-robust , low-overhead receiver architecture for low-power Ultra-Wideband Receivers, and its demonstration in a baseline 65nm CMOS process. The prototype achieves an energy consumption of $300\mu\text{W}$ at 1Mbps in receive mode and $75\mu\text{W}$ in transmit mode, including all digital and bias circuits, and does not require a crystal-based timer to operate .

This transceiver and its design methodology also demonstrate that *bit-level duty-cycling* is an effective technique to reduce the power consumption of RF VCOs beyond what is possible by using transistor sizing only, extending the portion of the design space in which power dissipation can be effectively traded off with sensitivity .

7.1 Research Directions

The techniques developed in this thesis can be extended in several ways. We present here two possible extensions-one regarding the digital communication aspects discussed in this thesis, the other concerning the implementation .

7.1.1 Power-efficient TCM for wireless sensor networks

In particular, while simple differential PPM was used in this work, more complex Trellis Coded Modulation schemes (TCM) could be applied to ensure sufficient transition density in the code while maintaining the ability to estimate B.E.R.. Such codes should be coupled with a higher precision timer to maintain reception fidelity leading to a trade-off between synchronizer power and circuit power. A literature search on this topic revealed that while TCM per-se is well understood and researched, its co-design with synchronization and reception hardware to reduce power consumption has not been extensively explored .

7.1.2 Circuit Implementation

An obvious (an already discussed) development in this context is to re-design the receiver to mitigate the issues with offsets and power-on time that were found during measurements. Such re-design would easily more than half the power dissipation , and improve the interference robustness and sensitivity to be close to the design values. On a longer term, we notice that at the sensitivity limit , the gain of from the RF input to the regenerative comparator input approaches unity. Power dissipation could therefore be reduced by moving the comparator bank closer to the antenna, for example using sub-sampling and absolute value detection. However, communication techniques have to be developed to guarantee interference robustness when using this architecture.

Appendix A

Matching Network Impedance Calculations

A.1 Receiver Side Impedance

First consider the conventional L-type matching network shown in figure A.1. Let the inductor impedance be

$$Z(L) = R_s + jX$$

at the frequency of interest. Using conventional series-to-parallel transformation, we find that in order for the network to achieve impedance matching to the source R_A , $R_s = \frac{R_A}{1+Q^2}$, $X = \frac{R_A Q}{1+Q^2}$, with $Q = \frac{X}{R_s}$. Since both quality factor and minimum value of passives are usually fixed, this network does not let independently choose R_A and Q , i.e. this problem does not in general have a solution. For example, for $R_A = 50$, $\omega = 2\pi \cdot 6GHz$, and an inductor quality factor of 10, the solution requires an inductance value of 130pH, which is not well controlled. Since for $Q \geq 1$ the inductance value is a decreasing function of Q , increasing inductance requires decreasing the quality factor. This can be done adding resistance (increasing the loss) or capacitance (decreasing the net magnetic energy stored). Increasing resistance is clearly the worse option, as it fixes the maximum achievable gain to $\frac{1}{2}$. When capacitance is added on the other hand, the voltage gain from the input to the inductor positive terminal is given by

$$A_v = \frac{1}{2} \left(\frac{j\omega L + R_s}{j(\omega L - \frac{1}{\omega C_1}) + R_s} \right) = \frac{1}{2} \frac{j(\frac{C_2}{C_1} + 1) + R_s \omega C_2}{j + R_s \omega C_2}$$

Since $\omega L - \frac{1}{\omega C_1} = \frac{1}{\omega C_2}$ to ensure resonance.




Figure A.1. Initial L-matching network

A.2 Transmit Side Impedance

We now use the circuit in figure A.1 to evaluate the impedance in transmit mode at the receiver resonance, i.e. $\omega_p = \sqrt{\frac{C_1+C_2}{LC_1C_2}}$. At this frequency, notice that $-\omega_p^2 LC_1 + 1 = -\frac{C_1+C_2}{C_2} - 1 = \frac{C_1}{C_2}$. Therefore, called Y_p the admittance seen at the negative terminal of capacitor C_2 , we have

$$Y_p = \frac{1}{j(\omega L - \frac{1}{\omega C_1}) + R_s} + \frac{1}{R_A} = \frac{j\omega_p C_1}{-\frac{C_2}{C_1} + j\omega_p R_s C_1} + \frac{1}{R_A}$$

After some algebra:

$$Z_p = (Y_p)^{-1} = \frac{R_A(j\omega_p R_s C_1 - \frac{C_2}{C_1})}{(j\omega_p (R_s + R_A) C_1 - \frac{C_2}{C_1})} = \frac{R_A}{\omega_p^2 (R_A + R_s)^2 C_1^2 + (\frac{C_2}{C_1})^2} ((\frac{C_2}{C_1})^2 + \omega_p^2 R_s (R_A + R_s) C_1^2) + j \frac{C_2}{C_1} (\omega_p R_A C_1) \quad (\text{A.1})$$

Appendix B

Statistical Jitter Transfer Analysis

We mentioned in Chapter 4 that the synchronizer exhibits a signal dependent gain in the jitter transfer path due to the quantization of the time measurement . We define *Gain* the ratio of the standard deviation of the time measurement to the standard deviation of the input measurement, and assume the input has normal distribution $N(\mu, \sigma)$. Since the output is quantized, we first need to compute its pdf. Denoted by $F_N(x)$ the standard Normal CDF, and T_s the measurement quantization step, and by $T_m \in Z$ the normalized measured time interval, we have that

$$Pr(T_m = K, \mu) = F_N\left(\frac{(K+1)T_s - \mu}{\sigma_{in}}\right) - F_N\left(\frac{KT_s - \mu}{\sigma_{in}}\right) \quad (B.1)$$

Which is shown in figure B.1 for different values of $\frac{T_s}{\sigma_{in}}$. Since the proposed clock-recovery scheme never achieves exact frequency lock, we can model the incoming edge phase as having an i.i.d. phase offset Θ with respect to the recovered clock. Alternatively said , $\mu \in U(-\frac{T_s}{2}, \frac{T_s}{2})$. We therefore compute the output noise variance using Bayes rule by averaging the variance of the output (conditional to μ) with respect to the unknown phase μ . This is given by

$$k_{avg} = \sum_k k Pr(T_m = k, \mu) \quad (B.2)$$

$$\sigma^2(\mu) = \sum_k (k Pr(T_m = k, \mu) - k_{avg})^2 \quad (B.3)$$

$$\sigma(T_m) = \sqrt{E_\mu\{\sigma^2(\mu)\}} = \sqrt{\frac{1}{T_s} \int_{-T_s/2}^{T_s/2} \sigma^2(\eta) d\eta} \quad (B.4)$$

Defined $G = \frac{\Delta\sigma(T_m)}{\Delta\sigma}$, values of $\sigma(T_m)$ and G are shown in figure B.2.




Figure B.1. Probability density function of time measurement for different values of input jitter variance σ_{in} to quantization resolution T_s

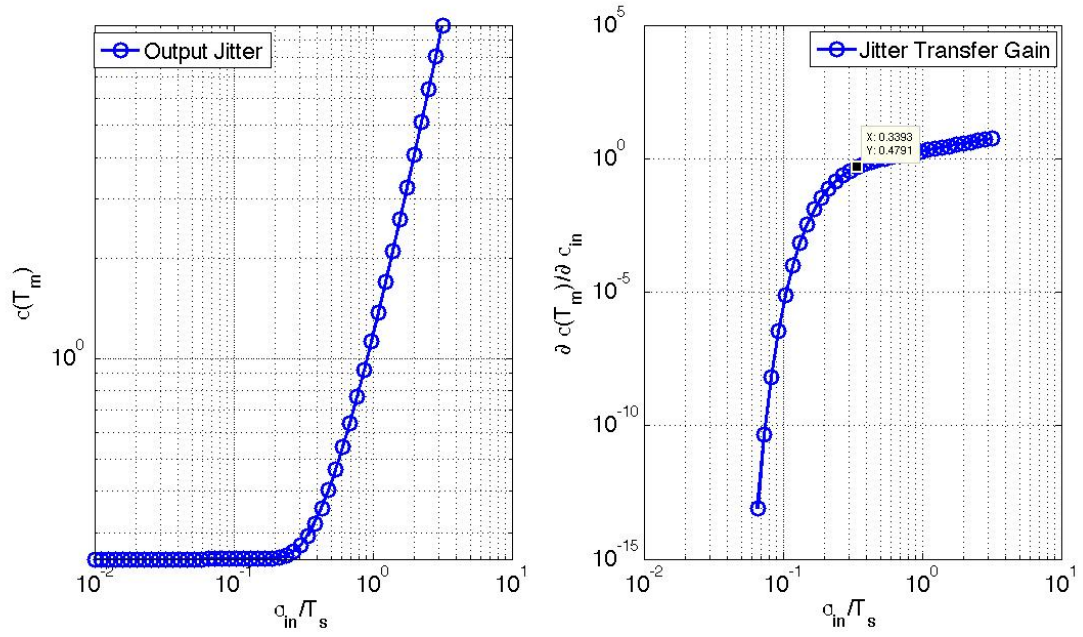


Figure B.2. Input-Output jitter transfer characteristic(left) and small signal gain(right)

Bibliography

- [1] A.Oncu and M.Fujishima. 19.2mW 2Gbps CMOS Pulse Receiver for 60GHz Band Wireless Communication. In *VLSI Circuits Symposium, 2008*. IEEE, 2008.
- [2] A.Poon, S.O'Driscoll, and T.H.Meng. Optimal Operating Frequency in Wireless Power Transmission for Implantable Devices. In *Engineering in Medicine and Biology Society (EMBS)*. IEEE, 2007.
- [3] C.A. Balanis. *Antenna Theory: Analysis and Design*. Wiley, 1997.
- [4] B.Cook, A.Molnar, and K.Pister. Low Power RF Design for Sensor Networks. In *Radio Frequency Integrated Circuits (RFIC) Symposium*. IEEE, 2005.
- [5] B.Razavi. *RF Microelectronics*. Prentice-Hall, 1997.
- [6] C.K.Yang, R. Farjad-Rad, and M.Horowitz. A 0.5 μ m CMOS 4.0Gbit/s Serial-Link Transceiver with Data Recovery Using Oversampling. *IEEE Journal of Solid State Circuits (JSSC)*, (5), 1998.
- [7] NFC Consortium. NFC Forum Type Tags White Paper V1.0. NXPBVTypeTagsWhitePaper-Apr09.pdf.
- [8] B.W. Cook, S. Lanzisera, A. Berny, A. Molnar, and K.S. J. Pister. Low-Power 2.4-GHz Transceiver With Passive RX Front-End and 400-mV Supply. *IEEE Journal of Solid State Circuits (JSSC)*, (12), 2006.
- [9] D.C. Daly, P.Mercier, Manish Bhardawaj, Alice Ston, Voel Voldman, Richard Levine, and A.P. Chandrakasan. A Pulsed UWB Receiver SoC for Insect Motion Control. In *Solid-State Circuits Conference, 2009. ISSCC 2009. Digest of Technical Papers. IEEE International*. IEEE, 2009.
- [10] D.C.Yates, A.S.Holmes, and A.J.Burdett. Optimal Transmission Frequency for Ultralow-power Short-Range Radio Links. *IEEE Transactions on Circuits and Systems I:Regular papers*, (7), 2004.
- [11] J.Rabaey D.Guermandi, S.Gambini. A 1V 250Kpps 90nm Transceiver For cm-Range Wireless Communications. In *European Solid-State Circuits Conference*. IEEE, 2007.

- [12] D.Ham and A.Hajimiri. Concepts and Methods in the Optimization of Integrated LC VCOs. *IEEE Journal of Solid-State Circuits*, (6), 2001.
- [13] D.Wentzloff and A.P.Chandrakasan. A 47pJ/pulse 3.1-to-5GHz All-Digital UWB Transmitter in 90nm CMOS. In *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International*. IEEE, 2007.
- [14] E.A.Klumperink, S.Louwsma, G.M.Wienk, and B.Nauta. A CMOS Switched-Transconductor Mixer. *IEEE Journal of Solid-State Circuits*, (8), 2004.
- [15] F.Lu, H.Samueli, J.Yuan, and C. Svensson. A 700MHz 24-bit Pipelined Accumulator in 1.2 μ m CMOS For Application as a Numerically Controlled Oscillator. *IEEE Journal of Solid-state Circuits*, (8), 1993.
- [16] A. Fort, J. Ryckaert, C. Desset, P. De Doncker, P. Wambacq, and L. Van Biesen. Ultra-Wideband Channel Model for Communication Around the Human Body. *Selected Areas in Communications, IEEE Transactions on*, (4), 2006.
- [17] B. Gyselinckx, C. Van Hoof, J. Ryckaert, R.F. Yazicioglu, P. Fiorini, and V. Leonov. Human++: Autonomous Wireless Sensors for Body Area Networks. In *Custom Integrated Circuits Conference, 2005. Proceedings of the IEEE 2005*. IEEE, 2005.
- [18] R. R. Harrison, P. T. Watkins, R. J. Kier, R. O. Lovejoy, D. J. Black, B. Greger, and F. Solzbacher. A Low-Power Integrated Circuit for a Wireless 100-Electrode Neural Recording System. *IEEE Journal of Solid State Circuits (JSSC)*, (1), 1995.
- [19] D. Hopkins, A. Chow, R. Bosnyak, B. Coates, J. Ebergen, S. Fairbanks, J. Gainsley, J. Ho, R. Lexau, F. Liu, T. Ono, J. Schauer, I. Sutherland, and T. Drost, R. Circuit Techniques to Enable 430Gb/s/mm² Proximity Communication. In *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International*. IEEE, 2007.
- [20] ARM Incorporated. ARM Cortex V0 Overview. <http://www.arm.com/products/CPUs/ARM-Cortex-M0.html>.
- [21] J.Bohorquez, A.P.Chandrakasan, and J.Dawson. A 350 μ W CMOS MSK Transmitter and 400 μ W OOK Super-Regenerative Receiver for Medical Implant Communications. *IEEE Journal of Solid State Circuits (JSSC)*, (4), 2009.
- [22] J.M.Rabaey, T.Karalar, J.Ammier, T.Tuan, and B.P.Otis. Pico-Radios for Wireless Sensor Networks: the Next Challenge in Ultra-Low Power Design. In *Solid-State Circuits Conference, 2005. ISSCC 2005. Digest of Technical Papers. IEEE International*. IEEE, 2005.
- [23] A.Chandarakasan J.Powell. Differential and Single Ended Elliptical Antennas for 3.1-10.6 GHz Ultra Wideband Communication. In *Antennas and Propagations*. IEEE, 2004.
- [24] J.Rabaey, S.Gambini, and D.Guermandi. A 3-10GHz Ultra-Wideband Pulser. UC Berkeley Technical Report 2006-136.

- [25] J.Ryckaert, M.Verhelst, M.Badaroglu, S.D'Amico, V. De Heyn, Desset C., P.Nuzzo, B. Van Poucke, P. Wambacq, A. Baschiroto, W. Dehaene, and G. Van der Plas. AC MOS Ultra-Wideband Receiver for Low Data-Rate Communication. *IEEE Journal of Solid State Circuits (JSSC)*, (11), 2007.
- [26] K.Iniewski. *Wireless Technologies:Circuits, Systems and Devices*. CRC Press, 2007.
- [27] F.S. Lee and A.Chandrakasan. A 2.5 nJ/bit 0.65 V Pulsed UWB Receiver in 90 nm CMOS. *IEEE Journal of Solid State Circuits (JSSC)*, (12), 2007.
- [28] T.H. Lee. *The Design of CMOS Radio-Frequency Integrated Circuits (2nd ed.)*. Cambridge University Press, 2004.
- [29] J. Liang, C.C. Chiau, X. Chen, and C.G. Parini. Study of a Printed Circular Disc Monopole Antenna for UWB Systems. *IEEE Transactions on Antennas and Propagation*, (11), 2005.
- [30] A. Liscidini, A. Mazzanti, R. Tonietto, L. Vandi, P. Andreani, and R. Castello. Single-Stage Low-Power Quadrature RF Receiver Front-End: The LMV Cell. In *Solid-State Circuits Conference, 2006. ISSCC 2006. Digest of Technical Papers. IEEE International*. IEEE, 2006.
- [31] L.Smaini, C.Tinella, C.Stoecklin, D.Helal, L.Chabert, C.Devaucelle, R.Cattenoz, N.Rinaldi, and D.Belot. Single-Chip CMOS Pulse Generators for UWB Systems. *IEEE Journal of Solid-State Circuits*, (7), 2006.
- [32] C. Marcu, D. Chowdhury, C. Thakkar, Ling-Kai Kong, M. Tabesh, Jung-Dong Park, Yanjie Wang, A. Afshar, B. Gupta, A. Arbabian, S. Gambini, R. Zamani, A.M. Niknejad, and E. Alon. A 90nmCMOS Low-Power 60GHz Transceiver with Integrated Baseband Circuitry. In *Solid-State Circuits Conference, 2009. ISSCC 2009. Digest of Technical Papers. IEEE International*. IEEE, 2009.
- [33] M.Chorost. *Rebuilt: My Journey Back to the Hearing World*. Mariner Books, 2006.
- [34] N. Miura, M. Mizoguchi, D and. Inoue, K. Niitsu, Y. Nakagawa, M. Tago, M. Fukaishi, T. Sakurai, and T. Kuroda. A 1 Tb/s 3 W Inductive-Coupling Transceiver for 3D-Stacked Inter-Chip Clock and Data Link. *IEEE Journal of Solid State Circuits (JSSC)*, (1), 2007.
- [35] M.McCorquodale, J.O'Day, S.Pernia, G.Carichner, S.Kubba, and R.Brown. A Monolithic and Self-Referenced RF LC Clock Generator Compliant with USB 2.0. *IEEE Journal of Solid State Circuits (JSSC)*, (2), 2007.
- [36] M.Verhelst, N.Van Hellenputte, and G.Geelen W.Dehaene. A Reconfigurable, 0.13 μ m CMOS 110pJ/pulse, Fully Integrated IR-UWB Receiver for Communication and sub-cm Ranging. In *Solid-State Circuits Conference, 2009. ISSCC 2009. Digest of Technical Papers. IEEE International*. IEEE, 2009.

- [37] M.Verhelst and W.Dehaene. System Design of an Ultra-Low Power, Low Data Rate Pulse UWB Receiver in the 0-960MHz Band. In *International Conference on Communications(ICC)*. IEEE, 2005.
- [38] K.Sarabandi N.Behdad. Bandwidth Enhancement and Further Size Reduction of a Class of Miniaturized Slot Antennas. *IEEE Transactions on Antennas and Propagation*, (8), 2004.
- [39] N.M.Pletcher, S.Gambini, and J. Rabaey. A $65\text{ }\mu\text{W}$, 1.9 GHz RF to Digital Baseband Wakeup Receiver for Wireless Sensor Nodes. In *Custom Integrated Circuits Conference, 2007. CICC 07 , IEEE*. IEEE, 2007.
- [40] N.M.Pletcher, S.Gambini, and J. Rabaey. A -72dBm sensitivity, $52\text{ }\mu\text{W}$ Wakeup Receiver using Uncertain-IF Architecture. In *Solid-State Circuits Conference, 2008. ISSCC 2008. Digest of Technical Papers. IEEE International*. IEEE, 2008.
- [41] K.O. O, P.Gorday, J.Lin, C.Cao, Y. Su, Z.Li, J. Mechta, J.Brewer, and S.Hwan. Progress Towards a Single Chip Radio. in *Low Power Electronics Design*, 2007.
- [42] Ian O'Donnell, S.Chen, and R.W. Brodersen. A 2.3mW Baseband Impulse-UWB transceiver in CMOS. In *VLSI Circuits Symposium, 2006. VLSI 06 , IEEE*. IEEE, 2006.
- [43] Stephen O'Driscoll, Ada S.Y. Poon, and Theresa H. Meng. A mm-Sized Implantable Power Receiver with Adaptive Link Compensation. In *ISSCC 2009*. IEEE, 2009.
- [44] B.P. Otis, Y.H. Chee, and J.M. Rabaey. A $400\text{ }\mu\text{W} - \text{RX}$, 1.6mW-TX Super-Regenerative Transceiver for Wireless Sensor Networks. In *Solid-State Circuits Conference, 2005. ISSCC 2005. Digest of Technical Papers. IEEE International*. IEEE, 2005.
- [45] Ozgur A. and Leveque O. and Tse D.N.C. Hierarchical Cooperation Achieves Optimal Capacity Scaling in Ad Hoc Networks. *Information Theory, IEEE Transactions on*, (10), 2007.
- [46] S. Pinel, S. Sarkar, P. Sen, B. Perumana, D. Yeh, D. Dawn, and J. Laskar. A 90nm-CMOS60GHz Radio. In *Solid-State Circuits Conference, 2008. ISSCC 2008. Digest of Technical Papers. IEEE International*. IEEE, 2008.
- [47] P.Mercier, D.Daly, and A.P.Chandrakasan. An Energy-Efcient All-Digital UWB Transmitter Employing Dual Capacitively-Coupled Pulse-Shaping Drivers. *IEEE Journal of Solid-State Circuits*, (6), 2009.
- [48] P.Mercier, D.Daly, and A.P. Chandrakasan. An Energy-Efficient All-Digital UWB Transmitter Employing Dual Capacitively-Coupled Pulse-Shaping Drivers. *IEEE Journal of Solid State Circuits (JSSC)*, (6), 2009.
- [49] A.S. Porret, Melly T., Python D., C.C. Enz, and E.A. Vittoz. An Ultralow-Power UHF Transceiver Integrated in a Standard Digital CMOS Process: Architecture and Receiver. *IEEE Journal of Solid State Circuits (JSSC)*, (3), 2001.

- [50] S.Rezeq R.B.Staszewski, J.L.Wallberg et al. All Digital PLL and Transmitter for Mobile Phones. *Solid-state Circuits, IEEE Journal of*, (12), 2005.
- [51] G. Retz, H. Shanan, K. Mulvaney, S. O'Mahony, M. Chanca, P. Crowley, C. Billon, K. Khan, and P. Quinlan. A Highly Integrated Low-Power 2.4GHz Transceiver Using a Direct-Conversion Diversity Receiver in 0.18 μ mcmos for IEEE802.15.4 WPAN. In *Solid-State Circuits Conference, 2009. ISSCC 2008. Digest of Technical Papers. IEEE International*. IEEE, 2008.
- [52] R.G.Meyer. Low Power Monolithic Peak Detector Analysis. *IEEE Journal of Solid State Circuits (JSSC)*, (1), 2007.
- [53] J. Ryckaert, Van Der Plas G., De Heyn V., Desset C., Vanwijnsberghe, G. Van Poucke B., and Craninckx J. A 0.65-to-1.4nJ/burst 3-to-10GHz UWB Digital TX in 90nm CMOS for IEEE 802.15.4a. In *Solid-State Circuits Conference, 2005. ISSCC 2007. Digest of Technical Papers. IEEE International*. IEEE, 2007.
- [54] S.Emami, C.H.Doan, A.M.Niknejad, and R.W.Brodersed. A 60-GHz Down-Converter CMOS Single-Gate Mixer. In *Radio Frequency Integrated Circuits (RFIC) Symposium*. IEEE, 2005.
- [55] Mingoo Seok, S. Hanson, Yu-Shiang Lin, Zhiyoong Foo, Daeyeon Kim, Yoonmyung Lee, Nurrachman Liu, D. Sylvester, and D. Blaauw. The Phoenix Processor: A 30pw Platform for Sensor Applications. In *VLSI Circuits, 2008 IEEE Symposium on*. IEEE, 2008.
- [56] S.Gambini, N.Pletcher, and J.Rabaey. Sensitivity Analysis in AM Detectors. UC Berkeley Technical Report 2008-41.
- [57] W.W. Si, D. Weber, S. Abdollahi-Alibeik, MeeLan Lee, R. Chang, H. Dogan, Haitao Gan, S. Rajavi, Y. Luschas, S. Ozgur, P. Husted, and M. Zargari. A Single-ChipCMOS Bluetooth v2.1 Radio SOC. *IEEE Journal of Solid State Circuits (JSSC)*, (12), 2008.
- [58] S.Pellerano, S.Levantino, C.Samori, and A.Lacaita. A 13.5mW 5-GHz Frequency Synthesizer with Dynamic Logic Frequency Divider. *IEEE Journal of Solid-State Circuits*, (2), 2004.
- [59] S.Rai, J.Holleman, J.N. Pandey, F.Zhang, and B.Otis. A 500 μ W Neural Tag With 2 μ V rms AFE and Frequency-Multiplying MICS/ISM FSK Transmitter. In *Solid-State Circuits Conference, 2009. ISSCC 2009. Digest of Technical Papers. IEEE International*. IEEE, 2009.
- [60] TI Technical Staff. CC2520 Data Sheet. cc2520.pdf.
- [61] S.Ye, L.Jansson, and I.Galton. A Multiple-Crystal Interface PLL with VCO Realignment to Reduce Phase-Noise. *IEEE Journal of Solid State Circuits (JSSC)*, (12), 2002.

- [62] Melly T., A.S. Porret ., Python D., C.C. Enz, and E.A. Vittoz. An Ultralow-Power UHF Transceiver Integrated in a Standard Digital CMOS Process: Transmitter. *IEEE Journal of Solid State Circuits (JSSC)*, (3), 2001.
- [63] L.S. Theogarajan. A Low-Power Fully Implantable 15-Channel Retinal Stimulator Chip. *IEEE Journal of Solid State Circuits (JSSC)*, (10), 2008.
- [64] U.Mengali and A.D'Andrea. *Synchronization Techniques for Digital Receivers*. Kluwer Academic Publisher, 1997.
- [65] B. Wicht, T. Nirschl, and D Schmitt-Landsiedel. Yield and Speed Optimization of Latch-Type Voltage Sense Amplifier. *IEEE Journal of Solid-State Circuits*, (7), 2004.
- [66] W.J.Butera. *Programming a Paintable Computer*. PhD thesis, Massachussets Institute of Technology, 2002.