A Baseband Mixed-Signal Receiver Front-End for 1Gbps Wireless Communications at 60GHz

David Amory Sobel

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2008-45
http://www.eecs.berkeley.edu/Pubs/TechRpts/2008/EECS-2008-45.html

May 2, 2008

A Baseband Mixed-Signal Receiver Front-End for 1Gbps Wireless Communications at 60GHz

by

David Amory Sobel

A.B. (Harvard University) 1997
M.S. (University of California, Berkeley) 2000

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

in

Engineering – Electrical Engineering and Computer Sciences

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in Charge
Professor Robert W. Brodersen, Chair
Professor Borivoje Nikolic
Professor Paul K. Wright

Spring 2008

The dissertation of David Amory Sobel is approved:

_____

Chair Date

_____

Date

_____

Date

University of California, Berkeley

Spring 2008

A Baseband Mixed-Signal Receiver Front-End for 1Gbps

Wireless Communications at 60GHz

by

David Amory Sobel

Abstract

A Baseband Mixed-Signal Receiver Front-End for 1Gbps

Wireless Communications at 60GHz

by

David Amory Sobel

Doctor of Philosophy in Engineering–Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Robert W. Brodersen, Chair

In 1995, the FCC allocated the spectrum from 59GHz to 64GHz as an unlicensed band, and shortly thereafter, extended this unlicensed band to 57-64GHz, thus providing 7GHz of unlicensed spectrum for general purpose use. Similarly, regulatory bodies across the globe have also set aside multi-GHz blocks of spectrum at 60GHz for unlicensed use. The presence of a true multi-GHz worldwide band has sparked immense interest in developing high-throughput 60GHz communications systems. Furthermore, the demonstration of CMOS mm-wave circuitry in recent years has made feasible the possibility of a highly integrated, all-CMOS 60GHz transceiver.

However, the design of such a system is made quite difficult by several factors: first, the 60GHz indoor channel has increased transmission loss due to both the increased free-space loss and the poor transmissivity of common building materials. As a result, systems employing omnidirectional antennae cannot achieve Gbps rates, and more complex systems employing adaptive beamforming must be architected. Secondly, the

limited performance of CMOS mm-wave components requires new approaches to system and architecture design, as the high level of system performance must be maintained with relatively low quality mm-wave circuits.

The work presented in this dissertation focuses on the analysis, specification, and design of a baseband system and architecture for 1Gbps, 60GHz wireless receiver. System exploration is conducted in order to find a suitable architecture for a high-performance system composed of limited-performance CMOS mm-wave circuitry. A mixed-signal baseband receiver design is proposed in order to minimize the overall power dissipation and implementation complexity. The proposed system performs much of the signal processing related to the task of synchronization in the analog domain, so that the signal is properly conditioned prior to sampling and quantization. This approach reduces the resolution requirements of the high-speed analog-to-digital converters (ADCs), dramatically reducing system power consumption.

The proposed architecture was designed and fabricated in a 90nm standard digital CMOS process. The design consists of a mixed-signal carrier phase rotator block with 500Ms/s DDFS DAC, a 16-tap mixed-signal complex decision-feedback equalizer running at 1Gs/s, and two 2Gs/s, 4-bit flash ADCs. This analog front-end achieves a peak SNDR of 24.8dB at the full Nyquist rate, has linearity exceeding 38dB, and dynamic range greater than 31dB. Overall chip power consumption is 55mW.

_____

Professor Robert W. Brodersen
Dissertation Committee Chair

To Eli, who continually fills my heart with joy.

Table of Contents

Table of Contents ii
Acknowledgements v
1 Introduction 1
1.1 Trends in High Data-Rate Wireless Systems 1
1.2 Recent Developments in 60GHz Regulatory and Industrial Landscape ... 3
1.3 Research Goals and Contributions 4
1.4 Organization of the Dissertation 5
2 The 60GHz Communications Channel 7
2.1 Review of Directional Antenna Technology 7
2.1.1 Definition of Terminology 8
2.1.2 Types of Directional Antennae 11
2.1.3 Performance of a Directional Antenna 17
2.2 Key Characteristics of the 60GHz Channel 21
2.2.1 Reduced Performance of CMOS mm-wave Components 22
2.2.2 Oxygen Absorption and Material Penetration 23
2.2.3 Path Loss and Antenna Directionality 26
2.2.4 Feasibility of a High Directionality Antenna 31
2.2.5 Regulatory Issues and Spectral Efficiency 32
2.3 Literature Review of 60GHz Channel Studies 34
2.3.1 Material Properties 35
2.3.2 Channel Properties 36
2.3.3 Summary of Literature Review 42
2.4 Conclusions 44
3 System Design Considerations 46
3.1 RF Circuit Limitations 47
3.1.1 Power Amplifier Linearity 47
3.1.2 Local Oscillator Phase Noise 52
3.2 Modulation Schemes 54

3.2.1 QPSK 54
3.2.2 Quadrature Amplitude Modulation 56
3.2.3 OFDM 58
3.2.4 Minimum-Shift Keying 61
3.2.5 Comparison 64
3.3 Baseband Architecture Considerations 65
3.3.1 Performance Requirements of MSK Baseband Receiver 66
3.3.2 Baseband Architecture Selection 71
3.4 Synchronization Algorithms for Mixed-signal Receivers 83
3.4.1 Feedback Synchronization Algorithms for Mixed-signal Systems 84
3.4.2 Carrier Phase Estimation 91
3.4.3 Symbol Timing Estimation 97
3.4.4 Channel Estimation 100
3.5 System Simulations 105
3.6 Conclusions 112
4 Circuit Implementation 113
4.1 Block-level Architecture 114
4.2 Carrier Phase Rotator 115
4.2.1 Gilbert Quad VGA 116
4.2.2 Feedback Tuning Circuit 122
4.2.3 DDFS Current-mode DAC 125
4.3 Decision-feedback Equalizer 128
4.3.1 DFE Tap Structure with Programmable Current Bias 129
4.3.2 Programmable Current Source Load 131
4.3.3 Programmable Delay Line 132
4.4 Summing Current Buffer 133
4.5 Transimpedance Track-and-hold Amplifier 136
4.5.1 THTIA amplifier 138
4.6 Flash ADC 140
4.6.1 ADC Architecture and Mismatch Analysis 141
4.6.2 Preamplifier 150
4.6.3 Averager 151

4.6.4 NMOS Comparator 152
4.6.5 CMOS Comparator 153
4.6.6 Digital Logic 154
5 Measurement Results 155
5.1 ADC Performance 156
5.2 Carrier Phase Rotator Performance 161
5.3 Decision-Feedback Equalizer Performance 164
5.4 Conclusions 172
6 Conclusion and Future Work 173
Bibliography 175

Acknowledgements

As with any endeavor, this research project could not have been completed without the support and assistance of many others. Foremost, I would like to thank my research advisor, Professor Robert Brodersen for his guidance, his support, and his unfailing ability to steer his students into fascinating and impactful research areas. Also, I would like to thank Professor Jan Rabaey, who along with Professor Brodersen, created the Berkeley Wireless Research Center. I feel truly privileged to have had the opportunity to study in such a rich environment. None of this would have been possible without all of their effort. I would also like to thank Professor Bora Nikolic for providing advice and insight into the area of high-speed serial communications; the course he taught on the topic was directly applicable to many of the techniques I investigated in my research. Thanks is also due to Professor Nikolic and to Professor Paul Wright for reviewing this dissertation and providing helpful feedback.

No department could operate without a support staff, and we were lucky to have such superb people. I must first thank Gary Kelson for his tireless efforts in securing silicon access on my behalf. There was a point in time during my research that I was a "man without a foundry", and were it not for Gary's efforts, I still would be waiting to tape out my chip. Tom Boot deserves special mention for always being willing to go above and beyond the call of duty in order to help the students deal with the vast bureaucracy that is the UC system. I also want to thank Kevin Zimmerman, Brian

Richards, and Brad Krebs for supporting the vast array of computing resources needed at the center.

I would like to thank TSMC for generously donating their foundry access, tools, and technology libraries in support of my work. I would like to particularly thank Dr. Fu-Lung Hsueh for championing our work within TSMC.

An advantage of being a graduate student for as many years as I have been one is that I have had the opportunity to work with and learn from so many brilliant and talented individuals. A disadvantage of being a graduate student for as many years as I have been one is that singling out each person who has had a positive impact on my graduate career would double the length of my dissertation (and perhaps take yet another year to complete!) Therefore, I must offer my apologies to those who go unnamed. However, quite a few of my colleagues deserve mention for their friendship, assistance, and company over these past years. During my first years at Berkeley, I had the good fortune of collaborating with Dennis Yee, Chinh Doan, Sohrab Emami, Johan Vanderhaegen, Brian Limketkai, and Sayf Alalusi. Our discussions—technical and otherwise—were a highlight of my "first tour of duty" at Berkeley. My Ph.D. research was more of a solitary effort, but my experience at Berkeley during those years was by no means solitary, due in large part to the company of many brilliant and friendly colleagues. My years as a Ph.D. student were enriched by many fellow students, most notably Ian O'Donnell, Josie Ammer, Mike Sheets, Tufan Karalar, Nate Pletcher, Simone Gambini, Luca De Nardis, and Danijela Cabric. Whether we were discussing our research, global politics, or the proper pronunciation of Italian vulgarities, your companionship will be missed.

Working part-time at Broadcom throughout much of my Ph.D. research may have added to the duration at my tenure at Berkeley, but it was a decision that greatly added to the quality of my work as a student. When I first joined Broadcom, I told my managers that I intended to eventually return to school to pursue a Ph.D., and I would like to thank them for fully supporting me in that decision. Special thanks are due to Haideh Khorramabadi, Neil Kim, Pieter Vorenkamp, and Tom Kwan for creating a situation where I could simultaneously pursue my academic endeavors while continuing to be a part of the engineering team within the company. I would like to single out my manager, Tom Kwan, to thank him for his support, advice, mentorship, and encouragement during my entire tenure at Broadcom. I would also like to thank all my colleagues in the analog team at Broadcom; I have learned so much from you over these years, and the skills and practices I have gleaned from our interactions have enhanced my work as both a Broadcom employee and a student.

I consider myself lucky to have such a loving and supportive family. I want to thank my parents for encouraging me to pursue my dreams and instilling in me the self-esteem and persistence to keep at it when those dreams feel somewhat out of reach. Your unfailing love and support has been a source of strength to me over the past 32 years. I also want to thank my sister, Heather, for her unconditional kindness. Your smile and your warmth always brightens my day. My brother Eric deserves thanks for simultaneously being my idol, my rival, my tormentor, and my best friend. While you are no longer physically with us, you will always be present in my heart, and I will cherish the 21 years that we shared. I have also been lucky enough to pick up some

wonderful “new” family members along the way. My in-laws, Ed and Glori Chaika, welcomed me into their family with open arms (even if those arms only came up to my waist) and have treated me like a son ever since.

Finally, I want to thank my wife, Melissa. First, I want to thank you for your patience. I know that this whole Ph.D. thing took a little longer than expected, as I know that there is a spreadsheet floating around somewhere that said I’d be graduating sometime in 2006. Thank you for supporting me throughout my entire Ph.D. career and making the countless small (and not-so-small) sacrifices that enabled me to complete this work. Most importantly, thank you for your love. Your presence by my side has enriched my life immeasurably, and I look forward to all of our adventures and misadventures together over the years. I dedicate this work to you with all my love.

1 Introduction

1.1 Trends in High Data-Rate Wireless Systems

Wireless technology has had a dramatic impact on the way we work, live, and play. Fifteen years ago, it was unthinkable to imagine that an individual could have data connectivity anywhere other than at a terminal or desktop that was tethered to an Ethernet port or a dial-up modem. Today, it is commonplace to assume that, through the use of wireless communications technology, a person can have voice, video, and data access anywhere in his office, his home, or just about any corner of the globe. With the rapid commercialization of wireless local area networking (WLAN) technologies such as 802.11 over the past several years, achievable indoor wireless data rates have scaled from roughly 1Mb/s to over several hundreds of Mb/s. This growth in bandwidth has enabled a host of new technologies and applications, including real-time audio and video streaming. Furthermore, the use of wireless, instead of wired, communications technology has helped spread the overall adoption of the technology in the consumer market, as ease of installation and customer satisfaction is greatly increased by the elimination of unwieldy and unsightly wires and cables [1].

As wireless access has become an increasingly important part of our everyday lives, the demand for wireless bandwidth similarly increases. Applications that were typically run over traditional wired networks—such as web surfing, file and printer sharing, and voice over IP (VOIP)—can now be easily run on 802.11 wireless networks. Similarly, recent developments in ultrawideband (UWB) technologies show promise for

delivering “wireless USB” connectivity between computers and storage-intensive peripherals like digital cameras, camcorders, and external hard disk drives at rates of up to 480Mb/s [2].

However, there are new wireless applications that demand even greater bandwidth than either 802.11 or UWB can provide. With the widespread adoption of HDTV’s and sources of HD content (such as HD set-top boxes, Blu-Ray DVD players, HD camcorders, et al), there is growing demand for real-time wireless streaming between these various HD devices. Similarly, wireless connectivity between a PC (such as a laptop) and a display (either a monitor or an LCD projector) is desirable in many environments. Additionally, the prevalence of mobile personal video players (such as the iPhone, et al) motivates for the capability to “instantaneously” synchronize and transfer large media files between the mobile device and a personal computer over a high-speed wireless link. As shown in Figure 1.1, these applications require throughput exceeding 1 Gb/s, well beyond the capacity of existing wireless systems. Therefore, other technologies must be developed in order to accommodate these new applications. These applications require short-range (on the order of 2-20m) wireless communications capable of handling throughputs of 1Gb/s and above in typical indoor residential and corporate environments. This dissertation presents a study of the key issues involved with enabling this new class of wireless technology and presents the design of a receiver for such a system.

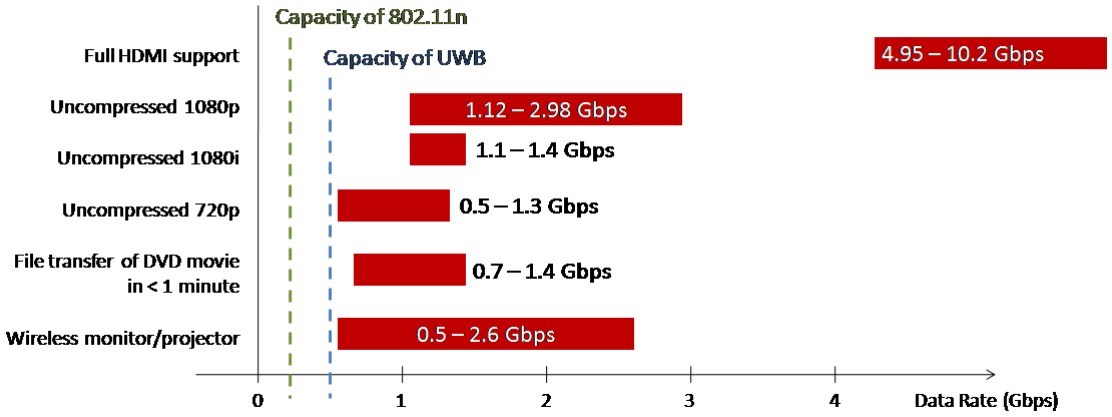

Figure 1.1: New wireless applications and their associated datarate requirements

1.2 Recent Developments in the 60GHz Regulatory and Industrial Landscape

In 1995, the FCC allocated the spectrum from 59GHz to 64GHz as an unlicensed, band [3]. Shortly thereafter, the FCC amended their rules to extend this unlicensed band to 57-64GHz, thus providing 7GHz of unlicensed spectrum for general purpose use [4]. The availability of a true multi-GHz band has sparked immense commercial interest in developing 60GHz technology in order to meet the demands of these new high-bandwidth wireless applications. In 2005, IEEE has organized the 802.15.3c task group to develop a standard for a 60GHz wireless personal area network (WPAN) with bandwidths in excess of 1Gb/s [5]. Around that same timeframe, several companies began research and development of 60GHz wireless technology for commercial applications. One of the key technical driving forces behind the commercial interest in 60GHz technology was the demonstration of the feasibility of implementing 60GHz

circuitry in standard, low-cost CMOS [6], [7]. However, due to the limitations of CMOS technology at 60GHz in comparison to other, more expensive materials, new system architectures and techniques must be investigated in order to make a high data-rate, all-CMOS solution feasible. A major aspect of the research described in this dissertation is the analysis and design of a receiver architecture that can accomplish this goal.

1.3 Research Goals and Contributions

The goal of the research presented here is to investigate system design issues related to the design of an all-CMOS, 1Gbps, 60GHz wireless receiver and to propose both a system architecture and a baseband receiver implementation that is capable of meeting this high throughput requirement with low power dissipation and implementation complexity. This research is centered around two key theses: first that the overall system design and architecture in inherently linked to the underlying circuit implantation, and that the system must be designed with the strengths and weaknesses of CMOS technology in mind at all times. The second thesis is that the judicious partitioning of the signal processing functions of the receiver baseband between the analog and digital domains is necessary for a truly low-power implementation. Therefore, the scope of this work covers both high-level system design and low-level circuit implementation research, as well as an investigation of the interplay between system and circuit design. In summary, the research contributions include:

- An analysis of modulation schemes appropriate for use with a 1Gbps, 60GHz all-CMOS receiver.

- The design of a mixed-signal baseband receiver architecture to reduce overall power dissipation and complexity.
- A full analysis and simulation of the proposed receiver architecture in order to determine necessary system specifications.
- The design and implementation of the proposed receiver in a 90nm digital CMOS process, including the design of blocks such as:
    - A novel mixed-signal carrier phase correction circuit.
    - A high-speed, low-power Gs/s decision-feedback equalizer using a new adaptive tap sizing and positioning scheme.
    - A low-power, 4-bit 2Gs/s ADC using a new staggered-averaging technique to minimize the impact of device mismatch.

1.4 Organization of the Dissertation

The rest of this dissertation is organized as follows: Chapter 2 examines various issues that are unique to 60GHz communications. The 60GHz indoor channel is analyzed in depth, and key characteristics of this channel are discussed. Also, a review of prior work measuring and characterizing the indoor 60GHz channel is presented. In Chapter 3, system issues related to the design of a baseband receiver architecture are discussed. The impact of CMOS circuit performance at 60GHz is discussed, and modulation schemes and architectures that are amenable to an all-CMOS implementation are investigated. Also, a mixed-signal receiver architecture that promises low power consumption and implementation complexity is proposed and simulated. Chapter 4 details the circuit

implementation of the various blocks within the receiver. Chapter 5 presents measurement results of the fabricated baseband receiver chip, and the dissertation closes with conclusions and suggestions for future research in Chapter 6.

2 The 60GHz Communications Channel

The 60GHz communications channel presents several new challenges and opportunities to implement very high bandwidth wireless systems. In many ways, the 60GHz channel has characteristics that are so dramatically different from lower frequency channels, that a wireless transceiver designed for this channel might look quite different from traditional transceiver designs. In this chapter, the 60GHz communications channel is introduced and a qualitative discussion of the key characteristics of the 60GHz channel is presented. This chapter opens with a brief review of directional antenna technology, because—as it will be shown—an approach utilizing directional antenna at 60GHz becomes both feasible and required in order to design a wireless system capable of Gb/s throughput.

As interest in 60GHz communications has grown, there have been numerous studies and measurement campaigns quantifying various parameters of the 60GHz indoor channel and its dependence on other physical parameters such as room size, building materials, and antenna configuration. A survey of these studies is presented in this chapter, with an emphasis on the key channel parameters that impact the system design of a high throughput 60GHz transceiver.

2.1 Review of Directional Antenna Technology

In order to understand the challenges and opportunities present in 60GHz wireless communications, a basic understanding on antenna technology is required. In particular, the utilization of directional antenna at 60GHz is a necessity, so a brief discussion of the

basics of directional antenna technology follows. In depth discussions and analyses of antennae, directional antennae, and their applicability to 60GHz wireless communications can be found in [8] and [9].

2.1.1 Definition of Terminology

When discussing directional antenna, there are a few basic terms that are frequently used. The first is the notion of an *isotropic antenna.* An isotropic antenna is defined as "a hypothetical lossless antenna having equal radiation in all directions" [8]. A *directional antenna*, by contrast, is defined as an antenna "having the property of radiating or receiving electromagnetic waves more effectively in some directions than in others" [8]. Therefore, all physically realizable antennas are by necessity directional in nature, as even the simple half-wave dipole is a more effective radiator in directions perpendicular to the dipole axis than in directions parallel to it. However, the term *directional antenna* is "usually applied to an antenna whose maximum directivity is significantly greater than that of a half-wave dipole" [8]. An antenna, such as the half-wave dipole, which has directionality in one plane but radiates equally in all directions in the other plane is referred to as an *omnidirectional antenna.*

Any directional antenna can be characterized by its *radiation pattern*. The radiation pattern indicates the gain of the antenna in a given plane as a function of the angular offset of the direction of propagation. As an example, the radiation pattern of a half-wave dipole is shown in Figure 2.1 (a) and (b). Assuming the dipole is oriented vertically, the radiation pattern in the azimuth plane is omnidirectional (see Figure 2.1 (b)). The

radiation pattern in the elevation plane (see Figure 2.1 (a)) does show some directionality as discussed above. The maximum directivity of this antenna is 1.64 (linear), or 2.15dBi.

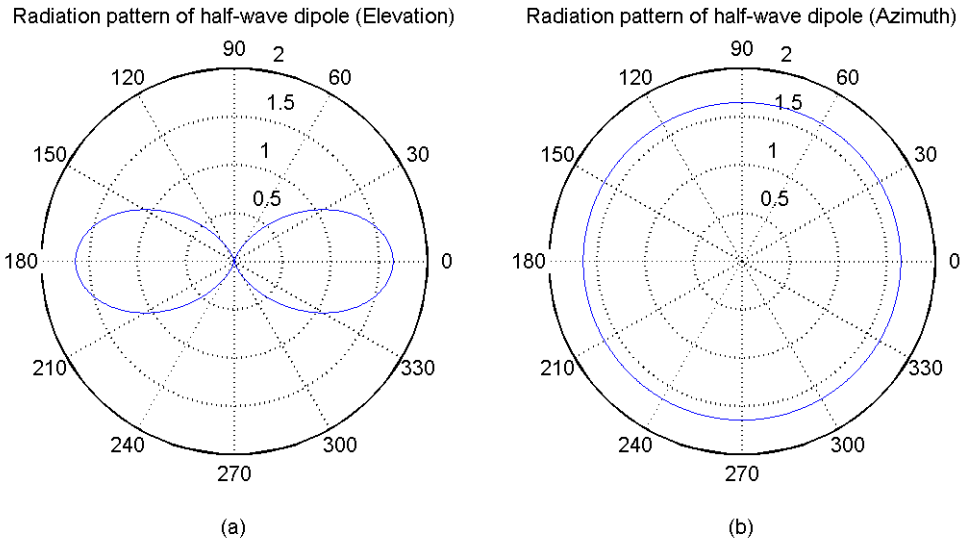

Figure 2.1: Radiation patterns of a half-wave dipole in (a) elevation plane and (b) azimuth plane

The *directivity* of a directional antenna can be derived from its radiation pattern. The directivity is defined as "the ratio of the radiation intensity in a given direction from the antenna to the radiation intensity averaged over all directions" [8]. Therefore, the directivity of an antenna is a function of the angle of observation; however, it is implicitly understood that the direction of interest is the direction of maximum radiation intensity. In this context, directivity can be defined as:

$$D_{max} = \frac{U_{max}}{P_{rad,avg}} = \frac{U_{max}}{P_{rad}/_{4\pi}} \tag{1.1}$$

where U_{max} is the maximum radiation intensity (in W/unit solid angle) and P_{rad} is the total power radiated by the antenna. Therefore, a higher directivity indicates that a higher percentage of the total radiated power is being transmitted in one particular direction.

Closely related to the notion of directivity (and often used interchangeably) is the concept of *antenna gain*. The definition is almost identical, but antenna gain is defined as the ratio of the radiation intensity in a given direction to the power *accepted* by the antenna (normalized by 4π) . As with directivity, antenna gain is a function of the angle of observation, but it is also typically assumed to refer to the direction of maximum radiation intensity. In this context, antenna gain can be expressed as:

$$G_{max} = 4\pi \frac{U_{max}}{P_{in}} = 4\pi \frac{U_{max}}{P_{rad}/e_{cd}} = e_{cd} \frac{U_{max}}{P_{rad}/4\pi} = e_{cd} D_{max} \tag{1.2}$$

where e_{cd} is the radiation efficiency of the antenna (≤ 1). As can be seen from (1.2), the antenna gain is just the product of the antenna directivity and the antenna efficiency. If the antenna is a perfect, lossless radiator then the two quantities will be equal; in all real cases, however, the gain of an antenna is (somewhat) less than its directivity because of losses internal to the antenna. Typically, the gain of a particular antenna is specified in dBi, which corresponds to the ratio of that antenna's gain to the gain of a lossless isotropic antenna (expressed in dB) [8].

The *half-power beamwidth* (HPBW) is another measure of the directivity of an antenna. The HPBW is the measure of the angle (in radians) over which the radiation intensity of the antenna is at least one-half the value of its maximum value. It follows that

an antenna with a higher directivity by necessity must have a smaller HPBW. In fact, the following lower and upper bounds on HPBW can be intuitively derived:

$$\frac{2\pi}{D_{max}} \leq HPBW \leq \frac{4\pi}{D_{max}} \tag{1.3}$$

The inverse relationship between antenna directivity and HPBW has significant ramifications when it comes to wireless link design; in particular, a highly directive antenna will require increasingly precise antenna alignment in order to harness the extra gain that such an antenna can provide. Antenna propagation studies have shown that an antenna typically needs alignment to better than 30% of its HPBW in order to be effectively used [10].

2.1.2 Types of Directional Antennae

Directional antennae can come in many shapes, sizes, and configurations, but they can be subdivided into two distinct classes: *fixed* and *adaptive*. A *fixed* directive antenna is an antenna that whose radiation pattern is determined solely by its mechanical structure, and hence the basic shape of its radiation pattern can be discerned by merely looking at it. Well-known examples include horn antennae, waveguide aperture antennae, and parabolic reflector antennae, (more commonly known as "dish antennae"). As mentioned above, the pattern of a fixed directive cannot be altered, although the pattern can be redirected by physically moving the antenna. An example of this are the rotating radar dishes typically seen at airports and on military equipment.

An *adaptive* directive antenna is, by contrast, one whose radiation pattern can be altered, typically by electrical means. The most common type—and the one of greatest

applicability to the 60GHz space—is the adaptive antenna array. An antenna array is a set of identical (and typically omnidirectional) antennae that are distributed in space in a particular arrangement and each driven by a unique waveform. The total field of the array is determined by vector addition of the fields radiated by the individual elements [8], and the placement and excitation of the array can be configured so that the net array radiation pattern has a high directivity aimed in the intended direction. Unlike fixed antennae, however, the pattern of an adaptive antenna can be altered by changing the relative excitation of the individual elements.

A common example of an adaptive antenna array is the *N-element linear array*. This array is composed of N identical omnidirectional antennae that are placed in a linear fashion with an element-to-element spacing of d. Figure 2.2 shows an example of this array. For the sake of this example, let us further assume that each element in the array is driven by an equal amplitude excitation, but each element is driven by a progressive phase of ϕ. (i.e. element 1 has a phase of 0, element 2 has a phase of ϕ radians, and so forth, up to element N, which has a phase of $(N-1)\phi$. The frequency of excitation is assumed to be fixed, with a wavelength λ, and the *wavenumber*, is defined as the quantity $k=2\pi/\lambda$ [8]. Lastly, in the far-field, the distance from the observation point to each antenna is almost equal, so that the amplitude of the radiation received from each antenna is equal. For mathematical simplicity, the radiation received from each element is assumed to be 1.

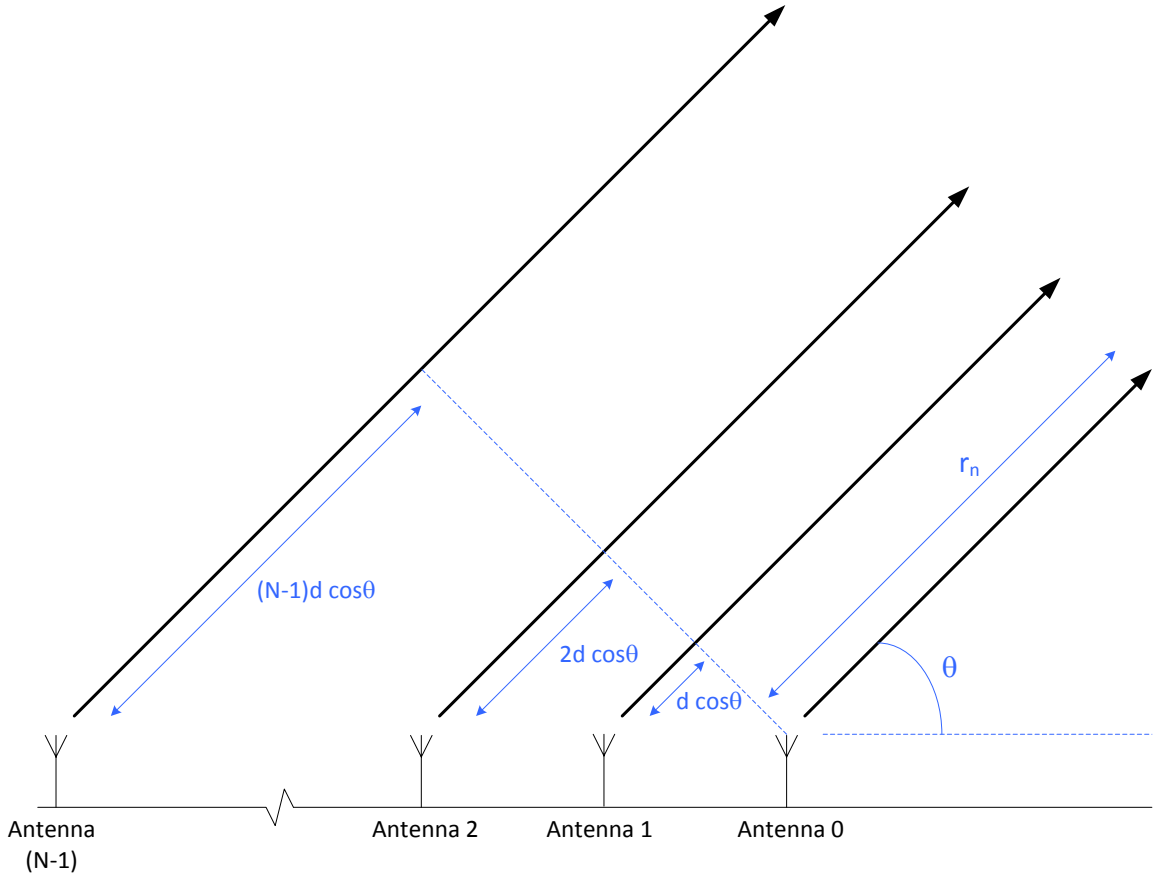

Figure 2.2: N-element linear antenna array (Annotations in blue)

Given these conditions, one can derive the overall antenna pattern as a function of observation angle, θ. The total received E-field is the sum of the E-fields received from each individual element.

$$E_{TOT} = \sum_{n=1}^{N} E_n \tag{1.4}$$

The distance from each antenna to the observation point is:

$$r_n = r + (n-1)d \cdot \cos\theta \tag{1.5}$$

The received E-field from each component is:

$$E_n = e^{-j\,(k\cdot r_n+(n-1)\cdot\phi)} = e^{-jkr}\cdot e^{-j\,[(n-1)(kd\,\cos\,\theta+\phi)]} \tag{1.6}$$

Plugging (1.6) into (1.4) and simplifying, we get:

$$E_{TOT} = e^{-j\left(kr+\frac{N-1}{2}\psi\right)}\cdot\frac{\sin\frac{N\psi}{2}}{\sin\frac{\psi}{2}} \tag{1.7}$$

$$\psi = kd\cos\theta + \phi$$

If one ignores the complex phase shift due to the first term in (1.7), the result is termed the *array factor* (AF) of the antenna and describes how the magnitude of the radiation pattern varies based on the observation angle, θ. The array factor is:

$$AF = \left[\frac{\sin\left(\frac{N}{2}\psi\right)}{\sin\left(\frac{1}{2}\psi\right)}\right] \tag{1.8}$$

where ψ is as defined in (1.7). An example of an array pattern is plotted in Figure 2.3. In this example, N=16 antennae are used with a spacing, d, equal to $\lambda/4$; ϕ is set to 0 radians, so $\psi = \frac{\pi}{2}cos\theta$.

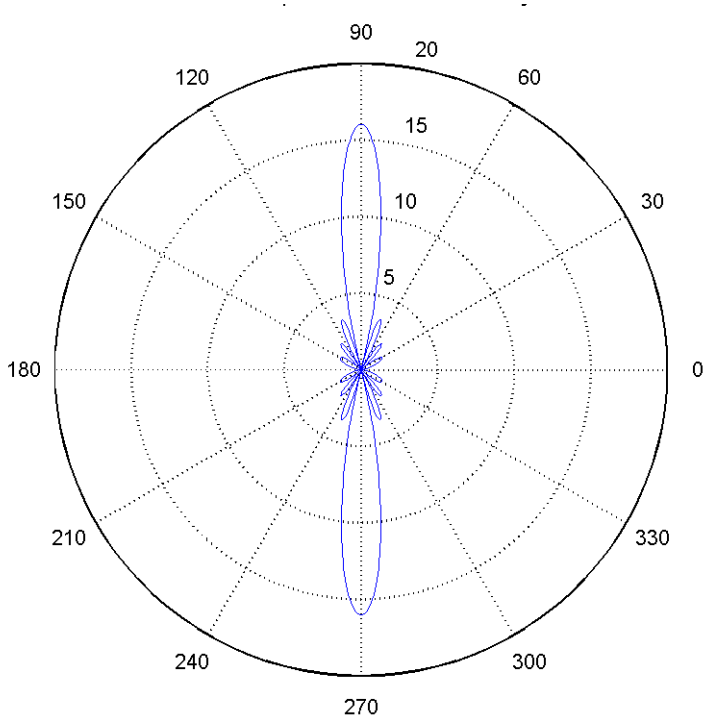

Figure 2.3: Radiation pattern of linear 16-element array, ϕ=0

The array factor is one of two parameters of an antenna array that determines its overall radiation pattern, and hence its directivity and antenna gain. The other factor is called the *element factor* (or EF), and is the radiation pattern of the individual antenna element used to compose the array. In the linear array mentioned above, the individual antenna are assumed to be 0dBi isotropic radiators. In fact, any antenna type could be used in an array; if a non-isotropic antenna is used, then the overall array radiation pattern is merely a product of the array factor and the element factor [9].

There are two important conclusions to be taken from equations (1.7) and (1.8). The first is that the array pattern is electrically programmable through the choice of the progressive phase, ϕ. In fact, the direction of maximum radiation can be set by enforcing

the condition that ψ=0. This leads to the condition that: $\phi_{opt} \triangleq -kd \cdot \cos\theta_{max}$, where θ_{max} is the desired angle of maximum radiation and ϕ_{opt} is the optimal choice for the progressive phasing between antennae. For example, Figure 2.4 shows an antenna pattern where ϕ is set to -π/4 and d is λ/4 and N=16.

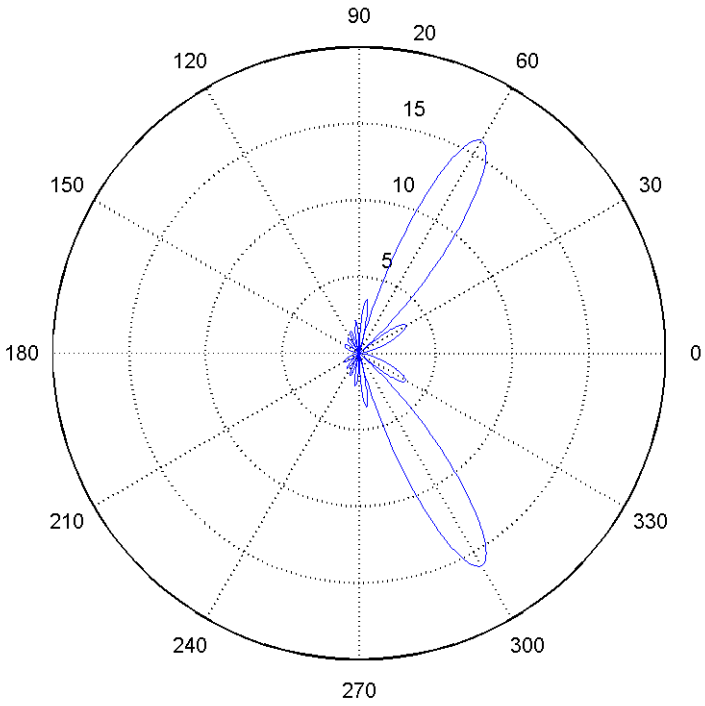

Figure 2.4: Radiation pattern of linear 16-element array, ϕ=-π/4

As expected, the peak directivity is achieved at θ = +π/6, or 60°. Also, a second directivity peak is achieved at θ = -π/6, as this is also a solution to the condition stated above. The second important conclusion is that the array factor is limited by the number of antenna used. In fact, solving for AF_{max} in equation (1.8) yields that the maximum AF achievable is N. Since the array factor is one of two factors that determine antenna gain, the directivity of an array can be increased in one of two ways: either the number of

antennas comprising the array can be increased or the directivity of the individual antenna elements can be increased. The second approach is of limited practicality, as a highly directive antenna element is, by definition, a fixed directive antenna, so its pointing direction cannot be easily modified. Therefore, highly directive antenna arrays that require a reasonable level of control in its pointing angle are typically composed of a large number of low-directionality components. This places a practical limit on the achievable antenna gain that can be implemented, as the complexity of implementing a large number antenna can be difficult. For discussions of techniques to achieve a 60GHz directional antenna array in conventional CMOS technology, see [9].

2.1.3 Performance of a Directional Antenna

The use of a directional antenna has two significant impacts on the overall characteristics of a wireless link. First, as described in the previous section, the directivity of an antenna (when properly aligned) gives additional gain to the signal over the wireless link. Figure 2.5 and Figure 2.6 give a graphical explanation of the cause of the increased antenna gain.

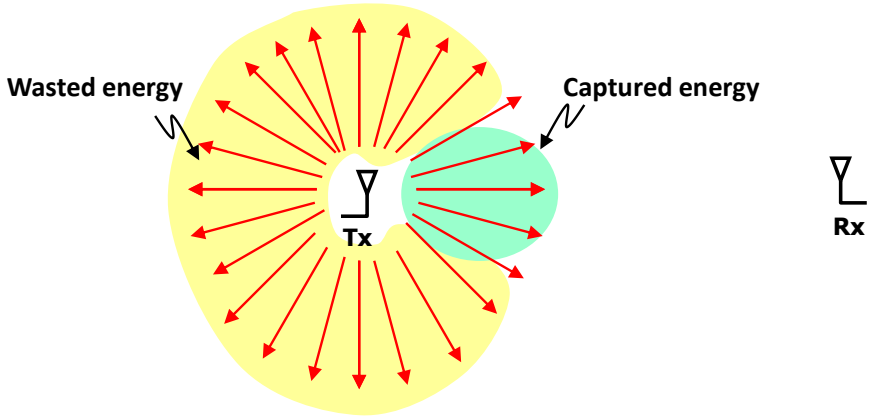

Figure 2.5: Graphical depiction of energy transmitted using omnidirectional antenna

In Figure 2.5, a wireless link is shown where both the receiver and the transmitter use omnidirectional antennae. Since the transmitter radiates energy equally in all directions, much of the transmitted energy is wasted and a relatively small fraction of the transmitted power reaches the receiver. This can be contrasted to the situation in Figure 2.6, where a directional transmit antenna is used. When properly aligned, the transmit antenna radiates a larger fraction of its total energy in the direction of the receive antenna, thereby increasing the total energy transmitted across the wireless link. The gain in received signal energy is exactly the directivity gain discussed in the previous section.

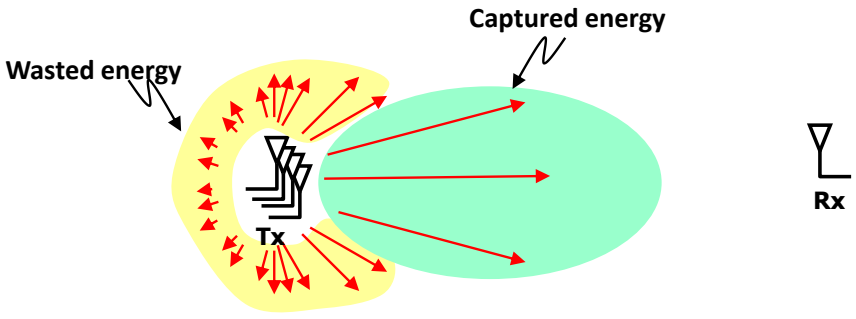

Figure 2.6: Graphical depiction of energy transmitted using directional antenna

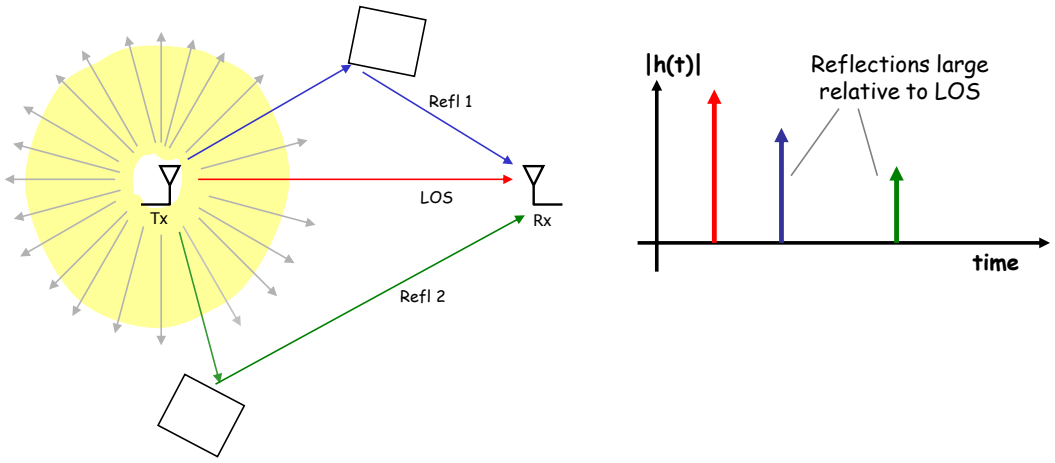

Figure 2.7: Graphical depiction of channel multipath using omnidirectional antenna

The use of a directional antenna also has significant impact on the nature and severity of the channel multipath. Multipath is caused the by reflection of the transmitted energy of off various objects in the vicinity of the wireless link; in indoor channel, reflections may be caused by walls, surfaces such as tables, floors, or ceilings, and various clutter in the local environment. In the graphical depiction of Figure 2.7, omnidirectional antenna are used; the channel has significant multipath due to the strong reflections off of objects near the transmitter or receiver. In Figure 2.8, a directional antenna is used for the transmitter. The directional antenna limits the spatial extent of the signal to the desired path. The same reflections occur as in the previous case, but now more energy is transmitted in the desired path and less energy in these unwanted reflective paths. As a result, the extent of the multipath can be significantly reduced.

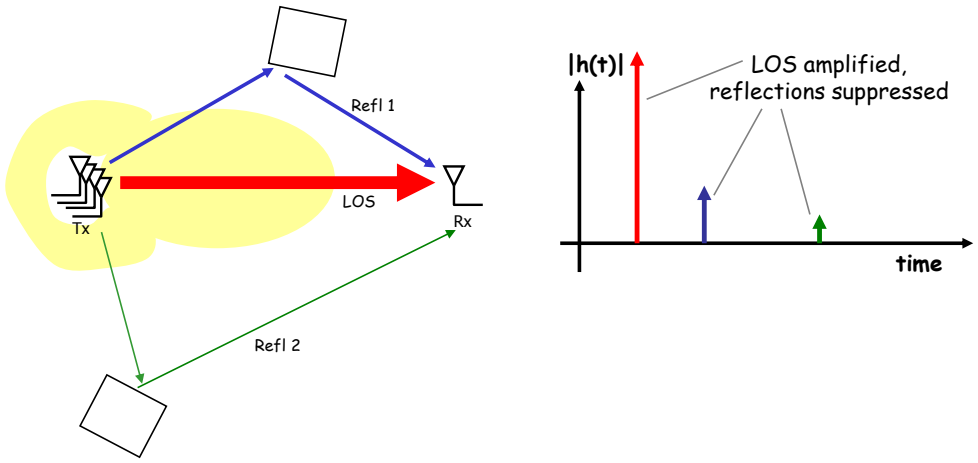

Figure 2.8: Graphical depiction of channel multipath using directional antenna

Measurements at HP Labs [11] confirm this behavior. Figure 2.9 shows a sample indoor channel impulse response with a omnidirectional RX antenna, and Figure 2.10 shows the sample impulse response when an RX antenna with moderate directionality (90° HPBW) is used. Comparison of the two plots shows that the LOS path (at about 9ns) is amplified by roughly 5-10dB and the multipath reflections (at 18 and 22ns) are suppressed by over 15dB.

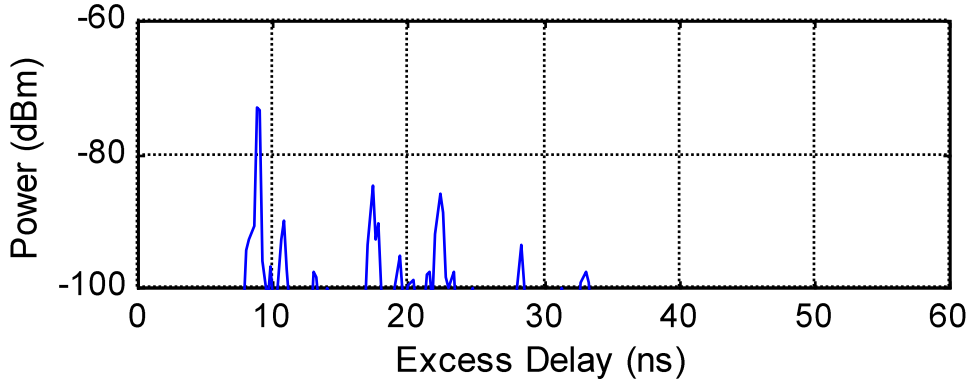

Figure 2.9: Measured channel impulse response with omni RX antenna (from [11])

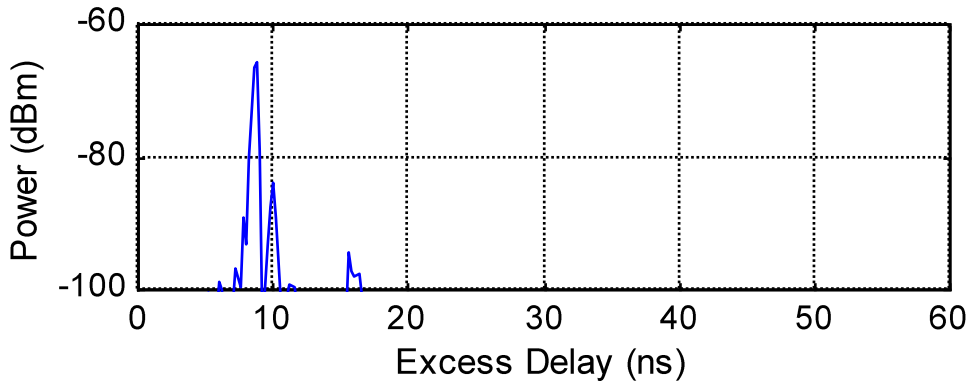

Figure 2.10: Measured channel response with 90° HPBW RX antenna (from [11])

Several studies have been conducted on the effect of antenna directivity on 60GHz indoor channel multipath. The results of these studies are summarized in section 2.3. Of note is the fact that moderately directive antennae can have a very significant impact on reducing the channel multipath [10], [12], [11], [13]. Therefore, when directive antennae are used, complex techniques for combating multipath—such as OFDM—might no longer be necessary. This opens up the system design to alternative approaches that can enable high data-rate systems to be built without requiring the high-performance RF components that OFDM systems typically require.

2.2 Key Characteristics of the 60GHz Channel

The 60GHz channel has some unique characteristics that differ significantly from the characteristics of communications channels in the low-GHz regime. Some of these differences stem from basic electromagnetic or material properties (e.g. oxygen absorption) whereas some of these differences are caused by extrinsic factors such as regulatory issues or practical issues regarding device form factor.

2.2.1 Reduced Performance of CMOS mm-wave Components

In order to enable widespread adoption of 60GHz wireless communications, the transceivers must be fabricated in a low-cost, highly integrable technology such as CMOS. While 60GHz radios have been demonstrated in other technologies such as SiGe [14] and GaAs HBT [15], the goal of the research at the BWRC has been to demonstrate a 60GHz radio entirely in CMOS. Due to the lower electron mobility in silicon, however, CMOS suffers a performance penalty at microwave frequencies when compared to these other semiconductor materials operating at 60GHz or when compared to CMOS operating in the low-GHz regime. Also, the lower breakdown voltage of deep submicron CMOS limits the power handling capabilities the devices; therefore, microwave CMOS power amplifiers (PA's) have reduced output power. Table 2.1 compares reported performance metrics of some key transceiver components for three different cases: CMOS operating at 60GHz, SiGe operating at 60GHz, and CMOS operating at 2.4GHz. The results indicate that the RF components available at 60GHz in CMOS do not match those in SiGe; nor does CMOS performance easily scale from the low-GHz regime up to 60GHz. While continued scaling of CMOS technology will likely enable higher performance 60GHz components, the system design needs to take into account the limited performance of the currently available components. The increased noise figure of the LNA will limit the SNR at the receiver; the limited power handling capability of the PA will limit the received power and may also increase the nonlinear distortion created at the transmitter; the extra phase noise from the VCO can create reciprocal mixing or cause symbol jitter that will impact receiver sensitivity.

Component	Specification	CMOS @ 60GHz	SiGe @ 60-77GHz	CMOS @ 2-5GHZ
LNA	Noise Figure	8.8dB [16]	5.7dB [17]	4dB [18]
	Gain	12dB	23dB	N/A
VCO	Phase noise	-85dBc/Hz @ 1MHz [19]	-95dBc/Hz @ 1MHz [20]	-116dBc/Hz @ 1MHz [21]
PA	1dB compression	+9dBm [22]	+14.5dBm [23]	+24.5dBm [24]
	Power-added efficiency	8.8%	12.8%	31%

Table 2.1: Comparison of reported performance of key transceiver components

The focus of the work presented in this dissertation is the baseband system architecture and circuit design, so the poor performance of CMOS at 60GHz does not directly impact the circuits presented herein. However, the system and baseband architecture of a 60GHz receiver must be robust under these limited circuit performance constraints. Therefore, a key challenge that this research addresses in chapter 3 is determining the techniques and approaches to implementing a very high-performance wireless link with limited performance RF components.

2.2.2 Oxygen Absorption and Material Penetration

It is known that oxygen molecules interact with microwave frequencies between approximately 50GHz and 70GHz [25]. As a result, atmospheric oxygen absorbs microwave radiation in the 60GHz band, resulting in excess attenuation of the transmitted signal by as much as 15dB/km at sea level [26]. Since this attenuation is in addition to the Friis path loss discussed in section 2.2.3, oxygen absorption severely limits the applicability of 60GHz communications to many long-range applications. However, for indoor, short-range applications such as WLAN or point-to-point media streaming, the O_2

absorption issue has little significance. For instance, a 100m transmit distance would only suffer from only 1.5dB of O_2 absorption; however, the Friis path loss equation [see equation (1.9)] predicts a path loss of 108dB over a similar distance, assuming 0dB antenna gain.

It is interesting to note that the O_2 absorption issue has actually been embraced by some applications. In high-data rate applications where a low probability of intercept is required, O_2 absorption can help prevent the signal energy from propagating significantly beyond the desired receiver, preventing other unwanted receivers from receiving the signal. An example of this is satellite-to-satellite communications: as the concentration of oxygen rapidly decreases with altitude above sea-level, the signal attenuation due to oxygen absorption in the satellite-to-satellite link is negligibly small. However, the signal energy would be drastically attenuated as it propagated through earth's atmosphere towards sea level; as a result, earth-based eavesdropping on these satellite links would be extremely difficult.

Of more significance to indoor communications are the transmissivity and reflectivity of common indoor building materials with respect to 60GHz radiation. As will be discussed in more depth in section 2.3.1, certain building materials are significantly more opaque to 60GHz radiation than to signals at lower frequencies Measurement campaigns comparing the material transmissivity at 60GHz to that of 2.5GHz [27], 1.7GHz [28], and 5.85GHz [29]all show that many common building materials have higher penetration loss and lower reflectivity at 60GHz than at these lower frequencies.

As a result of the poor material transmission at 60GHz, partitions in an indoor environment act as isolative boundaries between neighboring areas in a single building. As an example, [30] has determined that the material attenuation through a standard wood door would be more than 16dB. Similarly, measurements conducted at the BWRC, showed that typical interior wall and flooring constructions can block 60GHz attenuation by as much as 24dB. Therefore, partitions in indoor structures, such as walls and floors, prevent appreciable signal energy from propagating from one room to adjacent rooms.

The implications of these findings are threefold: First, unlike low-frequency WLAN systems such as 802.11a/b/g/n, a 60GHz wireless system is not well-suited to the typical residential deployment model, where a single basestation is used to serve many mobile clients scattered throughout the house. The attenuation suffered at 60GHz by propagating through one (or more) interior walls or floors is likely to add too much path loss to make high-bandwidth communications feasible. A corporate setting, however, would be more conducive to a 60GHZ WLAN deployment, as a ceiling-mounted basestation could serve several mobile clients without suffering from material penetration issues, provided that the internal structure of the office was more akin to a “cubicle farm” than several distinct office rooms.

A second consequence of the poor material penetration at 60GHz is a corollary of the first: since walls and floors prevent signal energy from spilling out of one room into the next, each room naturally becomes its own microcell. Furthermore, since the structural partitions isolate the adjacent cells from one another, frequency reuse can be aggressively exploited. In other words, each cell can utilize a large fraction (or perhaps

the entirety) of the 7GHz of available spectrum without worrying about intercellular interference. As a result, typical concerns of spectrum efficiency are mitigated, since the opportunities for frequency reuse are significant.

Lastly, the poor material penetration of 60GHz radiation has an effect on the nature of the multipath present in the indoor channel. Most notably, even with omnidirectional antennae, the delay spread of 60GHz indoor channels tends to be smaller than that of lower frequency channels. The path length of multipath reflections are necessarily shorter at 60GHz, since the reflections of sufficient energy can only come from in-room objects [27], [31], [32]. In [31], the delay spread of several different interior structures were measured at both 1.7GHz and 60GHz. In all cases, the delay spread measured at 60GHz was between two to four times smaller than that measured at 1.7GHz.

2.2.3 Path Loss and Antenna Directionality

The Friis path loss equation describes the free-space path loss incurred by an electro-magnetic wave as it propagates from a transmitter to a receiver. The Friis path loss equation is:

$$L = \frac{P_{RX}}{P_{TX}} = \frac{G_{TX} G_{RX} c^2}{f_c^2 (4\pi d)^2} = \frac{G_{TX} G_{RX} \lambda^2}{(4\pi d)^2} \tag{1.9}$$

where P_{TX} and P_{RX} are the power at the transmitter and receiver, G is the antenna gain, f_c is the carrier frequency, λ is the carrier wavelength, d is the propagation distance. Note that the convention used here defines L as the ratio of the power received to the power

transmitted. Hence, *L* is by definition less than or equal to 1 and a smaller value of *L* indicates a larger amount of path loss..

The Friis path loss equation can be viewed as a power flux equation; the total power exiting a sphere (or any other three-dimensional shape) is independent of the size or position of the sphere but instead must be equal to the total power being radiated by transmitters within that sphere. Therefore, it comes as no surprise that the free-space path loss increases with the square of the propagation distance. Furthermore, it can be shown that the antenna gain, *G*, is proportional to the antenna area, *A*, as follows:

$$G \propto \frac{A}{\lambda^2} \tag{1.10}$$

The power flux argument becomes clear if it is assumed that the transmitter is an omnidirectional source (i.e. G_{TX}=1) and the expression from (1.10) is used for G_{RX}. The path loss equation becomes:

$$L \propto \frac{A_{RX}}{(4\pi d)^2} \tag{1.11}$$

where A_{RX} is the area of the receive antenna. Equation (1.11) shows that the power captured by the receiver is proportional to the surface area of the sphere subtended by receiver antenna; hence the power flux argument becomes obvious.

The Friis equation as shown in (1.9) has led many to conclude that the path loss irrevocably gets worse as the carrier frequency increases [33], [34]. In fact, a direct application of the Friis equation can lead to the conclusion that a 60GHz communications

channel has 22dB more loss than an otherwise identical 5GHz communications channel. This conclusion is somewhat misleading, however, as the additional path loss incurred at higher frequencies is not due to some fundamental effect, but is primarily a consequence of the assumption that antenna gain is a fixed constant independent of frequency. (Typically it is assumed to be 0dB for the case of an omnidirectional antenna.) As indicated by equation (1.10), the area of an antenna *of a fixed gain* scales inversely with frequency. Thus, by the power flux argument of the preceding paragraphs, an antenna of fixed gain will be physically smaller at higher frequencies and will thus capture a smaller fraction of the total radiated power.

The antenna gain, however, is not constrained to be a fixed constant or somehow limited by some a priori fundamental constraint. A more likely constraint is the overall size of the antenna aperture, as physical parameters, such as the size of the wireless device, places an upper-bound on the overall area that can be devoted to the antenna [35]. If it is assumed that the antenna area, *A*, for a given device is fixed, the Friis equation takes the following form:

$$L \propto \frac{A_{TX} A_{RX}}{\lambda^2 (4\pi d)^2} \tag{1.12}$$

where A_{TX} is the area of the transmit antenna. Now it seems as if the path loss gets more benign as frequency increases, in direct contradiction of equation (1.9). From this, one might conclude that the use of higher carrier frequency and higher antenna gain provides a potentially unlimited benefit. In the example comparison between a 60GHz system and a 5GHz system, it now appears that the 60GHz system could achieve an improvement in

channel gain (and hence, SNR) of 22dB! Unfortunately, such an optimistic scenario is not the case either; since an antenna is a passive device, increasing the gain of an antenna can only be achieved by increasing its directionality. As was discussed in section 2.1, directionality can provide many benefits to wireless communications—including the increased antenna gain described here—but it also presents several significant implementation challenges. In particular, the system complexity required to properly produce and align a directional array pattern grows significantly as the antenna gain is increased.

Therefore, it remains to be determined what kind of antenna is suitable for a 60GHz system. The benefit of an omnidirectional antenna is the simplicity of this approach: no additional hardware is necessary in order to create a directional, high-gain antenna pattern, and no scheme to align the pattern is required either. However, the increased path loss suffered at high frequency with an omnidirectional antenna would limit the capacity of the wireless link. An upper-bound of the channel capacity can be approximated using the Shannon capacity calculation for an AWGN channel. The Shannon capacity limit gives the maximum theoretical error-free capacity that can achieved through a channel using arbitrarily complex coding schemes [36]. The Shannon capacity is:

$$C = BW \log_2\left(1 + \frac{P_{RX}}{BW\ N_0}\right) \tag{1.13}$$

where C is the channel capacity in bits/sec, BW is the channel bandwidth is Hz, and N_0 is the input-referred noise level at the receiver in W/Hz. In the case of 60GHz communications, the bandwidth is sufficiently large such that it can be approximated as infinite. Given that approximation, equation (1.13) simplifies to:

$$\lim_{BW\to\infty} C = 1.44\frac{P_{RX}}{N_0} \tag{1.14}$$

Equation (1.14) can be used to as an upper-bound to determine the theoretically maximum capacity that can be achieved. A more conservative approximation of channel capacity can be based on the observation that common modulation schemes (e.g. BPSK, QPSK, and MSK) require a minimum receiver SNR in order to get a reasonable uncoded error rate. For these modulations schemes, in order to get an uncoded BER on the order of 10^{-3}, an SNR at receiver of roughly 7dB is required.

Using the Shannon capacity limit or the 7dB SNR requirement, the capacity of an omnidirectional 60GHz link as a function of transmit power can be determined. Figure 2.11 shows such a plot for a transmit distance of 10m and 20m. In this plot, it is assumed that the receiver has an overall noise figure of 10dB, and there are 10dB of losses from miscellaneous sources such as shadowing loss, board losses, and implementation loss. According to the Shannon capacity limit, the transmitter would require 23dBm (200mW) transmit power to achieve a capacity of 1Gbps at 10m; according to the more realistic 7dB SNR criterion, the transmitter would have to output 30dBm (1W) of power to achieve this capacity! As shown in section 2.2.1, 60GHz power amplifiers with these power levels are simply unrealizable in commercially available CMOS technologies. Therefore, in order to realize a 10m, Gbps link at 60GHz, the path loss dictates that directional antennae are a requirement.

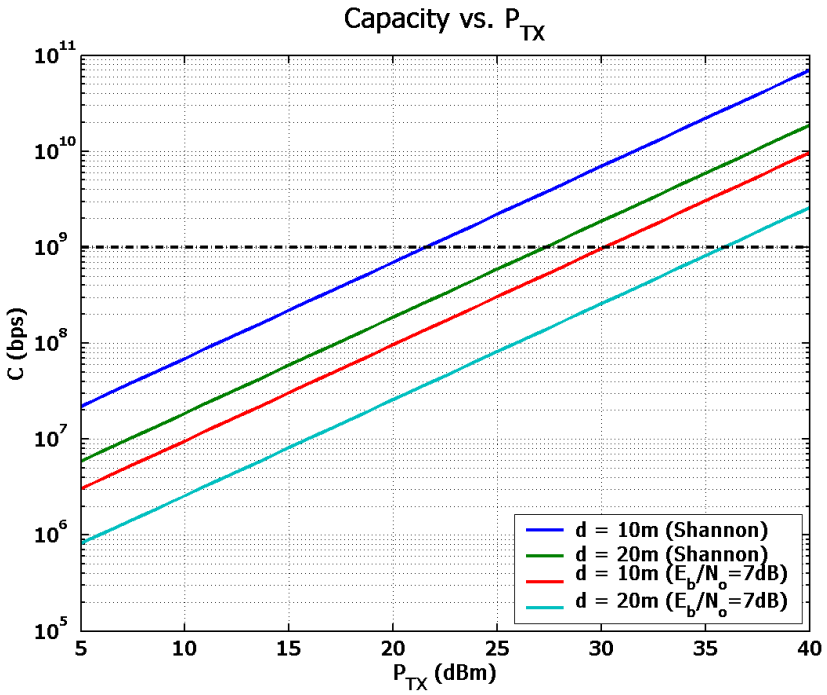

Figure 2.11: Link capacity vs. transmit power for omnidirectional antenna

2.2.4 Feasibility of a High Directionality Antenna

The previous sections demonstrate that the increased path loss at 60GHz and the limited performance of 60GHz CMOS circuits necessitate the use of directional antennae. Luckily, at 60GHz the use of directional antennae becomes much more practical than at lower frequencies. As discussed in section 2.2.3, the area required for an antenna of a particular gain is inversely proportional to the square of the carrier frequency. For instance, a rectangular aperture antenna with 12dBi directivity could be synthesized in a 5mm x 5mm area at 60GHz, whereas a similar antenna at 2GHz would be 16.8cm on a side [8]. Therefore, for small form-factor mobile devices where PC board real estate is at

a premium, operating at 60GHz allows for the use of directional antennae that would be prohibitively large at lower frequencies.

2.2.5 Regulatory Issues and Spectral Efficiency

In 1995, the FCC allocated the spectrum from 59GHz to 64GHz as an unlicensed, band [3]. Shortly thereafter, the FCC amended their rules to extend this unlicensed band to 57-64GHz, thus providing 7GHz of unlicensed spectrum for general purpose use [4]. Furthermore, regulatory bodies across the globe have also set aside multi-GHz blocks of spectrum at 60GHz for unlicensed use; Japan has allocated 59-66GHz as an unlicensed band and Europe has allocated 57-66GHz. (See Figure 2.12 for the allocation of unlicensed spectrum in various world-wide regions.)

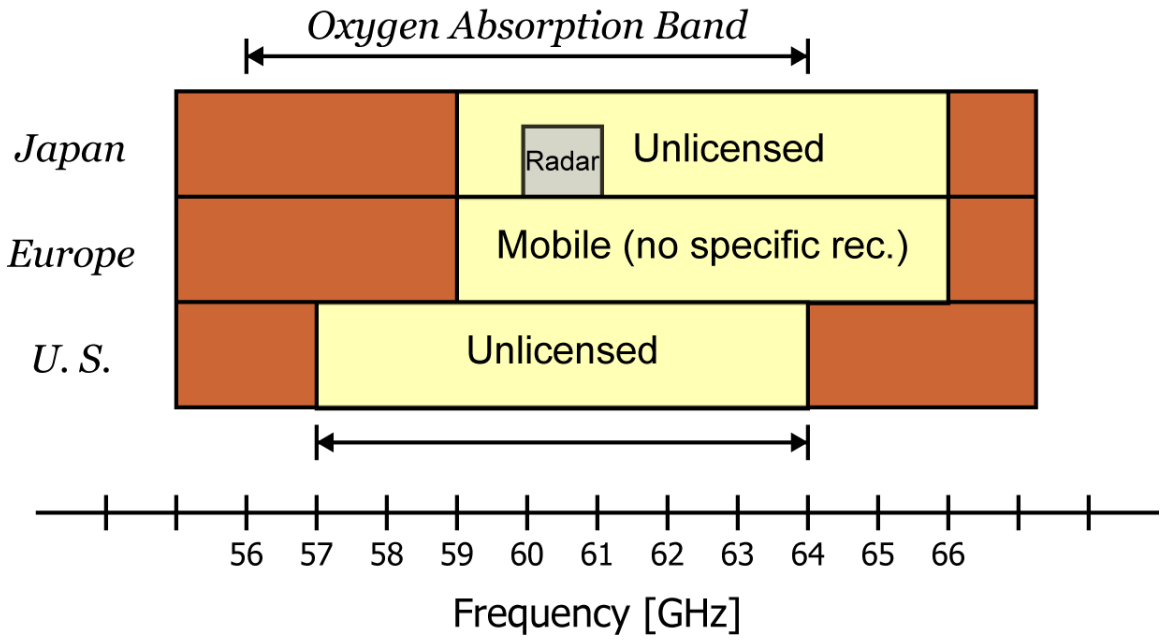

Figure 2.12: Worldwide allocation of unlicensed spectrum around 60GHz

The presence of a true multi-GHz worldwide band has sparked immense commercial interest in developing 60GHz communications links. Commercial issues aside, the availability of such a large bandwidth at 60GHz also has some profound technical implications as well. With such a wide bandwidth available for communications, bandwidth scarcity is no longer an issue, and spectral efficiency is no longer a parameter of utmost importance. Therefore, the system is free to explore schemes that trade-off circuit or system complexity for spectral efficiency.

In sharp contrast, the evolution of the 802.11 standard from .11b to .11a/g to .11n was driven primarily by the desired to "cram" more throughput through a fixed 20MHz channel bandwidth. As the standard evolved, the spectral efficiency went from roughly 0.5b/sec/Hz for 802.11b to upwards of 10b/sec/Hz for some draft .11n proposals. With this increase in spectral efficiency, however, comes a dramatic increase in the specifications of the analog and RF components. With higher spectral efficiency, higher SNR's are required, which in turn mandates higher performance components: the noise figure on the LNA must be reduced; the phase noise of the oscillator must be improved; and the linearity of all the components—most notably the power amplifier—must be dramatically improved.

However, as discussed in section 2.2.1, the current state of 60GHz CMOS circuits does not display the kind of performance necessary for a system with high spectral efficiency. Therefore, designing a 60GHz system with low spectral efficiency becomes not just a choice, but a *necessity*.

2.3 Literature Review of 60GHz Channel Studies

Recent interest in 60GHz communications has led to a number of studies of the properties indoor 60GHz channel in an effort to both better understand the nature of the 60GHz channel and to characterize it is an accurate, yet compact, manner. The studies published in literature fall into one of two categories: measurement-based campaigns and simulation campaigns based using ray-tracing tools. Each technique has their own set of advantages and complement the other. Measurement-based campaigns generate the required physical data of 60GHz propagation, transmission, and reflection that form the basis of any understanding of the 60GHz channel. However, measurement campaigns are labor-intensive and as a result are limited in their scope and in the diversity of physical environments measured. Simulation-based campaigns can build upon the measurement campaigns by applying the propagation and material properties discerned from these earlier studies and quickly and easily applying them to a variety of physical environments and room geometries. As will be discussed below, these studies indicate that the indoor channel properties are heavily dependent on room geometry and configuration, so a wide variety of environments need to be simulated. As mentioned above, channel simulations generally have used optics-based ray-tracing tools to simulate the wave propagation and generate a deterministic channel model for a particular environment [10]. One limitation for ray-tracing simulators is that they typically do not model diffusion, diffraction or other scattering mechanisms; however, many studies of the 60GHz channel indicate that the these scattering mechanisms typically do not occur in 60GHz indoor channels and

that a ray-traced simulation with only specular reflections will produce valid data [37], [38], [39].

2.3.1 Material Properties

Many of the measurement studies characterize the reflectivity and transmissivity of common indoor building materials. The transmissivity of a material is typically specified as the transmission loss (in dB/cm) that 60GHz radiation incurs while propagating through a given distance of that material. The reflectivity of a material is typically specified as the loss incurred by reflecting off the surface of that material. Table 2.2 below summarizes the material property measurements of a few measurement campaigns.

	Ref [37]		**Ref** [11]	**Ref** [30]	**Ref** [27]	**Ref** [38]
Material	**Trans (dB/cm)**	**Refl (dB)**	**Refl (dB)**	**Trans (dB/cm)**	**Trans (dB/cm)**	**Trans (dB/cm)**
Concrete	>6	6.5	N/A	6.67	N/A	6.7
Plasterboard	2	5.5	14.8	1.5	2.4	N/A
Wood	6.3	18	14	4.2	N/A	4.2
Chipboard	N/A	N/A	N/A	5.15	N/A	5.1
Acrylic glass	4.2	5.5	7.3	1.03	N/A	N/A
Glass	5.1	12	N/A	5.05	11.3	6.03
Wire-mesh glass	>14	3	N/A	N/A	31	N/A
Office whiteboard	N/A	N/A	N/A	N/A	5	N/A

Table 2.2: Transmission and reflection loss of common building materials at 60GHz

A few conclusions can be reached by looking at the table above. First, material transmission at 60GHz is poor, particularly through exterior structural elements such as concrete and wood. Therefore, 60GHz communications is not suitable for short-range building-to-building links where LOS is not guaranteed. Even in indoor environments with most building materials of typical thickness, room-to-room isolation is usually

greater than 20dB [30]. Therefore, 60GHz links seem most suitable to single-room environment, and a microcellular approach utilizing a high degree of frequency reuse is realizable.

Additionally, non-metallic building materials tend to be poor reflectors of incident 60GHz radiation. With the exception of wire-mesh glass and tiles, all of the 32 materials tested in [37] had a reflection loss greater than 5dB, with typical losses exceeding 10dB. As a result, the amplitude of most multipath reflections in 60GHz channels will be relatively small. For example, the extensive measurements in [12] conclude that in the absence of strong reflectors, "the reflected multipath components are at least 10dB below the LOS component." Similarly, in [31], the delay spread of several different interior structures were measured at both 1.7GHz and 60GHz. In all cases, the delay spread measured at 60GHz was between two to four times smaller than that measured at 1.7GHz. Additional measurements of 60GHz indoor channel properties are discussed in the next section.

2.3.2 Channel Properties

The aim of the measurement and simulation campaigns summarized here was to extract meaningful properties of the 60GHz channel. While the exact metric extracted differed from study to study, the metrics can be roughly grouped into the following categories:

- **Temporal characterization of channel multipath**: A variety of metrics were extracted to determine the temporal characteristics of the channel

multipath. Some metrics included RMS delay spread ([10], [38], [40], [12] [41]) Rician k-factor ([40], [10], [11]) and 90% settling time ([11]).

- **Spatial characterization of channel multipath**: Some of the studies characterized the spatial nature of the 60GHz channel multipath. For instance, [12] conducted extensive spatial measurement using highly directional antenna to extract angle of arrival information. Other studies ([10], [11], [13], [42], [43]) measured the impact of antenna directivity, alignment, and polarization on the temporal properties of the multipath channel.
- **Path loss**: Most studies conducted a set of location-specific path loss measurements. For instance, [12] reported the path loss as a function of distance in an interior corridor. This is the simplest metric extracted, and while it provides the least level of detail, the overall path loss in an indoor environment is of crucial importance.

2.3.2.1 RMS delay spread and Rician K with Omnidirectional Antennae

The two most relevant temporal metrics of the 60GHz multipath channel are RMS delay spread and Rician k-factor. The RMS delay spread, *R*, is a power-weighted measure of the temporal duration of the multipath reflections and is defined as follows:

$$R = \sqrt{\frac{\sum_{k=1}^{N}(\tau_k - t_a)^2 a_k^2}{\sum_{k=1}^{N} a_k^2}}$$
$$t_a = \frac{\sum_{k=1}^{N} \tau_k a_k^2}{\sum_{k=1}^{N} a_k^2} \quad (1.15)$$

where N is the number of multipath components, τ_k is the arrival time of the kth multipath ray, a_k is the amplitude of the kth multipath ray, and t_a is power-weighted average arrival time. High RMS delay spreads indicate channels whose multipath is significant over longer periods of time; hence, channels with low R relative to the symbol period have less intersymbol interference to contend with.

The Rician k-factor, K, is a ratio of the power in the dominant (LOS) signal path versus the total power in all other multipath components. Hence, K, is defined as:

$$K = \frac{a_{\max}^2}{(\sum_{k=1}^{N} a_k^2) - a_{max}^2} \quad (1.16)$$

A high K is indicative that most of the received power is contained in the dominant LOS component; hence the eye diagrams of high-K channels are cleaner and "more open". Higher K's are desirable since they indicate a lower probability of degradation from multipath.

Many studies have measured and analyzed the RMS delay spread (RDS) of typical indoor 60GHz channels. [44] measured the delay spread in several typical indoor environments with omnidirectional antenna and found that over all environments (including both LOS and NLOS situations), RDS fell within the range of 15-100ns. Furthermore, the RDS was a strong function of the physical environment; as expected, rooms with highly reflective walls had much higher RDS than rooms with those with less

reflective walls. For example, a “lecture hall” with concrete and wood walls gave RDS measurements in the range of 15-25ns, whereas a “computer room” with smooth metal walls had an RDS between 40 and 50ns. The campaign in [45] expanded this study to measure the impact of room size on RDS. The three rooms mentioned above were considered “small rooms”, each with dimensions of roughly 35 x 32 x 10 ft. A larger room, with dimensions of 140 x 133 x 23, had a delay spread between 50 and 70ns. The longer delay spread of larger rooms is expected; since typical wall material is poorly reflective at 60GHz, multipath components will only persist for a few “bounces” before dying out. Therefore, the multipath components in a larger room will have a later arrival time with respect to those in a smaller room.

The ray-tracing campaign in [46] corroborates the results above. Omnidirectional antenna were used and the room size was roughly 80 x 32 ft. Delay spread was about 10-12ns with a transmit distance of 10m and line-of-sight (LOS). Furthermore, the delay spread was proportional to transmit distance; as the transmit distance was decreased to 2m, RDS decreased to about 2-4ns. Once the LOS path was blocked, however, simulated RDS increased to 40-50ns. The Rician K-factor (RKF) was also simulated and was found to decrease with increasing transmit distance; at 2m, K was roughly 10-15dB and at 10m K was 6-8dB under LOS conditions.

Williamson, et al conducted a similar ray-tracing campaign of a typical office environment in [10]. This study included the effect of office partitions causing blockage to the LOS path. The omnidirectional transmitter was in a fixed location in the room, and the omnidirectional receiver was located at various points in the room. Including both

LOS and NLOS configurations, the median delay spread was roughly 23ns and 90% of locations has a delay spread of less than 27ns.[1] The median Rician K-factor was roughly 5dB and 90% of locations had a RKF of 1dB or better.

The author later conducted a channel measurement campaign at Hewlett Packard Laboratories [11]. While the office environment under investigation was different from that studied in his previous work, the measured results were largely in line with those from [10]. In this study, an omnidirectional RX antenna was used, and the TX antenna had a 90^{o} HPBW; also LOS was maintained for all measurements. The median delay spread was 10ns and 90% of locations has a delay spread of less than 15ns. The median Rician K-factor was roughly 4dB and 90% of locations had a RKF of 1dB or better.

2.3.2.2 Channel Spatial Characteristics and the Impact of Antenna Directivity on Delay Spread and Rician K

As discussed qualitatively in section 2.1, directive antennae have to potential to reduce channel path loss and to reduce channel multipath. Reducing path loss leads to higher received power, and reducing channel multipath leads to lower delays spread and higher Rician K. As with the studies presented in the previous section, the exact impact of antenna directivity is a strong function of the physical environment under study, but definite trends can be observed across a wide range of measurements.

[1] Note that these statistics indicate that 40% of the locations in the room has a delay spread between 23ns and 27ns.

The Williamson report [10] mentioned above also performed the ray-tracing simulations of the same office environment assuming directional antenna at the TX and RX. With both antennae having a 60^o HPBW (12 dBi), the median delay spread reduces to roughly 6ns and 90% of locations have an RDS less than 13ns. Similarly, the median Rician K is 18dB and 90% of locations have an RKF better than 7dB. Antenna alignment is shown to be important, but provided that the misalignment is less than 30% of the HPBW, no performance degradation is noticeable. If misalignment is equal to the HPBW, then the delay spread will roughly double and Rician K will plummet by as much as 10-15dB. The measurements later conducted at HP Labs [11] corroborate these simulation results; again, the office environments are different, and different antenna configurations are used, but the results are largely in line. A setup where the TX antenna is omnidirectional and the RX antenna has a 90° HPBW yields a median delay spread of 10ns, and 90% of the RDS values are below 15ns.

Smulders et al [13] also performed ray-tracing simulations with antennas of varying directivity. TX and RX antennas with 10^o, 30^o, and 60^o (corresponding to 25dBi, 16dBi, and 9dBi) were used. In LOS environments, the 7dBi configuration has an RDS of less than 7ns, and the 25dBi configuration has an RDS of less than 1ns. However, 5° misalignment of the 25dBi antenna can cause RDS to increase to roughly 10ns and can cause a reduction in received power by as much as 25 to 30dB.

Xu, Rappaport, et al [12] conducted an extensive study of the spatial characteristics of the 60GHz indoor channel in a typical office/university environment. They used a highly directional RX antenna and, at each location measured, rotated the RX antenna in

a complete circle and measured the channel properties at each rotation angle. They performed the measurement in both intra-room and inter-room scenarios in order to extract LOS and NLOS characteristics, respectively. When the antenna were properly aligned, the RDS in all of their LOS and NLOS measurements ranged from 5ns to 20ns; however, when alignment was not guaranteed, the RDS roughly doubled. Also, they defined a spatial peak-to-average ratio, where the peak power was defined as the maximum power received during a particular spin measurement (i.e. at proper antenna alignment), and the average power was defined as the average power received over all rotation angles in that spin measurement. The PAR reported for most measurements was in the range of 12 to 14dB. This PAR metric can be viewed as the average amount of available RX antenna gain that is lost if the RX antenna is improperly aligned. Therefore, it becomes clear that antenna alignment is critical if the antenna gain is needed to maximize the received power.

2.3.3 Summary of Literature Review

A few key points about the 60GHz channel can be culled from the wealth of data available in the literature. They are summarized below:

- **Common building materials significantly attenuate 60GHz transmission**. Many indoor building materials are relatively opaque to 60GHz signal radiation, especially when compared to lower frequencies. In the absence of a strong reflective path, the extra path loss incurred in NLOS environments would significantly degrade the overall performance of the wireless link.

- **Common building materials are poor reflectors at 60GHz**. With the exception of metallic objects, most other building materials do not reflect 60GHz radiation very well. Also, the reflections tend to be specular in nature, rather than diffuse. As a result, the multipath at 60GHz will be smaller than the multipath at lower frequencies.
- **Configurations with omnidirectional antennae will require additional techniques to mitigate multipath, even in LOS conditions**. Delay spreads with omnidirectional antenna were measured in the range of 15-50ns in LOS conditions. Rician K-factors in the 1 to 5dB range would be common. When compared to the desired data rate of 1GB/s, the delay spread is big enough that significant efforts would be required to compensate the multipath. For instance, an OFDM approach would require over 100 subcarriers, and an equalizer might require over 100 taps.
- **Directional antenna can significantly decrease the channel multipath**. Moderately directive antennae (6 to 12dBi) can reduce the delay spread below 10ns and maintain a Rician K above 10dB, even in some NLOS scenarios. Highly directive antennae (16dBi and above) can further reduce the delay spread to 1 to 5ns range or below.
- **Directive antenna rely upon proper alignment of beam pattern to be effective**. Directive antennae can increase the received power due to their antenna gain. However, this benefit is lost if the antennae are not properly aligned. In fact, misalignment can cause upwards of 20-30dB of additional path loss, which is well

in excess of the benefit provided by the antenna gain. Also, the delay spread and Rician K degrade with misalignment, and alignment sensitivity increases with antenna gain. A rule of thumb is that the pointing error must be less than 30% of the HPBW in order to suffer minimal performance degradation.

2.4 Conclusions

In this chapter, the 60GHz indoor channel was introduced, and key characteristics were discussed. It was shown that the relatively poor performance of CMOS at 60GHz and the large amounts of path loss at this frequency makes high SNR wireless communications very difficult. The availability of several GHz of unlicensed bandwidth, however, affords the system designer the opportunity to architect a system that trades off RF component performance for spectral efficiency. It was also demonstrated that the large path loss at 60GHz prevents the use of an omnidirectional antenna for Gb/s communications; fortunately, directional antenna are small enough at 60GHz to be realizable in small form-factor portable devices.

A literature review of the 60GHz channel was presented. Measurements and simulations performed in these studies indicated that the use of a directional antenna greatly suppressed the channel multipath in indoor settings. Delay spreads of less than 10ns and Rician K's of greater than 10dB were typical with moderately directive antennae, and highly directive antennae could achieve delay spreads below 5ns. However, the efficacy of directive antennae were limited by two key factors in our desired application: first, in order for the antenna to be electronically steerable, an

adaptive array composed of a multitude of low-directionality antennae must be used; in order to keep implementation complexity relatively low, the array cannot be composed of too many elements, hence limiting its maximum directivity. Secondly, the sensitivity to antenna misalignment grows with increasing directivity; therefore, in the presence of imperfect alignment, the benefit gained from using a directional antenna is subjected to diminishing returns.

A 16-element antenna array (in a 4x4 grid) presents a suitable candidate for the desired application. It achieves a maximum directivity of 12dBi, which will increase the received power by 24dB when used on both transmitter and receiver. It has a HPBW of 75°, which represents a good compromise between multipath suppression and alignment requirements. A discussion of the techniques to implement such an array in CMOS can be found in [9].

3 System Design Considerations

In this chapter, the system issues related to the design of a baseband receiver architecture are discussed. The chapter opens with a detailed discussion of the performance limitations of mm-wave CMOS circuits and the impact of these limitations on system performance and modulation choice. Several candidate modulation schemes are then introduced in section 3.2, and the tradeoffs between these schemes are discussed. Simulations of these candidate schemes with nonideal mm-wave circuit models is used to directly compare the feasibility of each of these schemes, and it is found that Minimum-Shift Keying (MSK) is the most suitable modulation scheme for this application. In section 3.3, different baseband architectures are considered, and a new mixed-signal receiver architecture is proposed. This architecture is capable of operating under indoor multipath scenarios and uses mixed-signal equalization and carrier recovery in order to minimize the dynamic range requirements key circuit blocks, thus lowering the overall system complexity and power consumption of the entire receiver. The architecture proposed here is contingent upon robust digital synchronization and equalization algorithms in order to properly function; these algorithms are introduce in section 3.4, and issues specific to a mixed-signal architecture are discussed. Based on the proposed architecture, the system specifications are analyzed by means of system-level simulations. The simulation model, results, and system specification is discussed in section 3.5.

3.1 RF Circuit Limitations

As introduced in section 2.2.1, CMOS circuits have reduced performance capabilities at 60GHz, especially when compared to lower frequency operation. (See Table 2.1.) The system design of a receiver must take this limited performance into account when targeting 1Gb/s system performance, as traditional architectures and techniques might no longer be feasible. There are two circuit blocks whose performance has been identified as critically limiting factors in 60GHz performance: the power amplifier (PA) and the voltage-controlled oscillator (VCO). Performance limitations of these two circuit blocks and their impact on system performance are discussed below.

3.1.1 Power Amplifier Linearity

The power amplifier (PA) is typically the final active component in a transmitter and is typically directly connected to the transmit antenna or to a related passive transmit devices (i.e. bandpass filter, transmit/receive switch, et al). As such, the PA sets the upper limit for the amount of power that can be transmitted into the channel. As discussed in section 2.2.3, due to the large path loss at 60GHz and the high throughput specifications of this system, a significant amount of transmit power is required from the transmitter and, hence, the PA. This requirement is further complicated by the fact that the PA must be able to output these high power levels without significantly distorting the constellation shape. Therefore, both the maximum power handling capability and its linearity are crucial specifications for a 60GHz PA.

The linearity and power handling capabilities of a PA is typically specified by its 1dB compression point (P_{1dB}) and saturation power (P_{sat}). The saturation power is the maximum output power that the PA is capable of delivering . The 1dB compression point is measured as the PA output power level at the point where the large-signal gain is 1dB less than its small-signal gain. Both these characteristics are examples of *AM/AM distortion*; in the case of gain compression and output saturation, changes in the input amplitude have a nonlinear effect on the output amplitude. In addition to AM/AM modulation, PAs also exhibit *AM/PM modulation***,** where the phase of the output waveform is nonlinearly affected by the input amplitude. Both of these nonlinear modulation mechanisms can be modeled in the manner described below [47].

Let the input signal to the PA be defined as follow:

$$x(t) = r(t)\cos[\omega_0 t + \varphi(t)] \tag{2.1}$$

where r(t) is the signal envelope, ω_0 is the carrier frequency, and φ(t) is the signal phase. The output of the PA is then:

$$y(t) = A[r(t)]\cos\{\omega_0 t + \varphi(t) + \Phi[r(t)]\} \tag{2.2}$$

Where the term A[r(t)] models the nonlinear gain (AM/AM) behavior of the PA, and Φ[r(t)] models the amplitude-dependent phase-shift (AM/PM) behavior of the PA.

Prior research has developed models that capture the nonlinear and saturating characteristics typical to solid-state PAs. The Rapp model [48] encapsulates AM/AM distortion and saturating behavior in the follow model:

$$A[r] = \frac{a \cdot r}{\left[1 + \left(\frac{a \cdot r}{R_{sat}}\right)^{2p}\right]^{\frac{1}{2p}}} \tag{2.3}$$

where a is the PA small-signal gain, R_{sat} is the output saturation amplitude, and p is an integer. An example of the input-output relationship of the Rapp model is shown in Figure 3.1(a). Also, there are similar models for AM/PM distortion. The model proposed by Saleh [47] is:

$$\Phi(r) = \frac{\alpha \cdot r^2}{1 + \beta \cdot r^2} \tag{2.4}$$

where α and β are parameters of the PA. Figure 3.1(b) shows typical AM/PM distortion from the Saleh model.

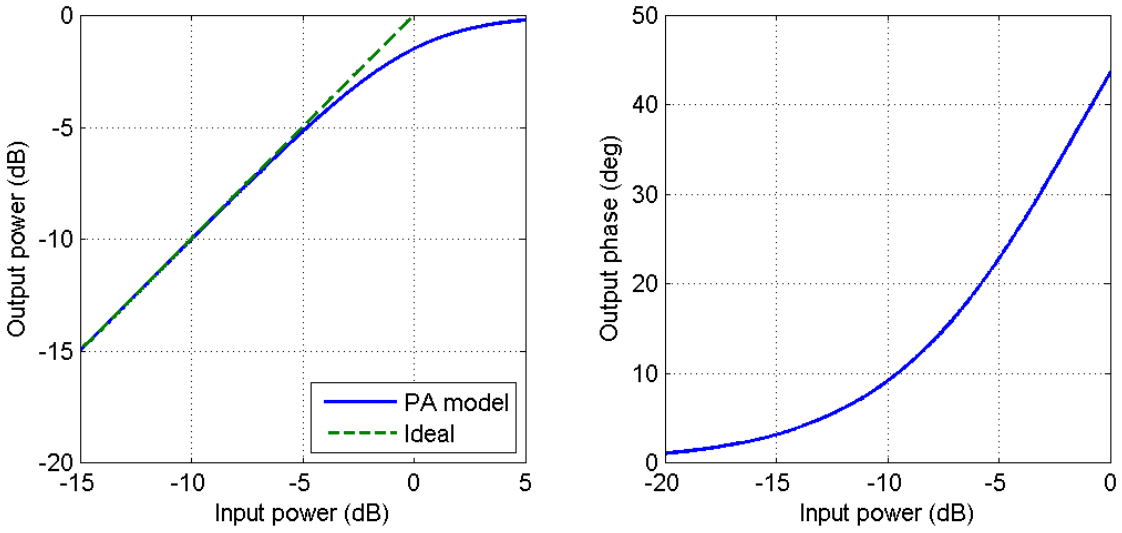

Figure 3.1: (a) AM/AM distortion model and (b) AM/PM distortion model

As is generally expected with nonlinear amplifiers, the nonlinearity can be avoided by operating at lower power levels, (i.e. moving to the left on the curves in Figure 3.1). However, as was discussed above, the 60GHz amplifier must be able to deliver high output power levels, on the order of +10dBm. Therefore, the system design must devise techniques that are robust to the nonlinear behavior exhibited by the PA. A careful

inspection of the equations, however, yields another solution: it is important to note that the nonlinearities in the PA are only excited by amplitude fluctuations in the envelope of the transmitted signal and not by fluctuations in the transmitted signal itself. Therefore, if the amplitude fluctuations in the signal envelope can be kept sufficiently small, the nonlinearities of the PA will not have significant impact on the overall PA behavior. The AM/AM nonlinearity will only manifest itself as a reduced (but linear) PA gain and the AM/PM distortion will cause an increased (but fixed) phase shift through the PA.[1]

PA nonlinearity will have several deleterious effects on the transmitted signal. First, it will cause spectral regrowth, where the carefully shaped transmit spectrum will spill out into neighboring spectrum, violating any predefined spectrum mask and potentially causing interference in adjacent channels. This is not a particularly important issue at 60GHz, as the large amount of spectrum available can allow for significant guard bands to be allocated between adjacent channels. However, this effect does point to a potential problem with intercarrier interference in a multicarrier system (i.e. OFDM). In other words, each individual OFDM subcarrier will experience spectral regrowth, and the distortion created by this phenomenon will spill over onto adjacent subcarriers.

Constellation warping, another detrimental effect of PA nonlinearity, will manifest itself in both single-carrier and multicarrier systems. The AM/AM distortion will compress the signal constellation in a nonlinear fashion, and the AM/PM distortion will

[1] Strictly speaking, this is not entirely true. The nonlinearity of the PA will be excited by the oscillation in the signal itself at the carrier frequency. However, for constant envelope modulation schemes the distortion created by PA nonlinearity will only fall at harmonics of the carrier frequency and hence not impact system performance.

create an amplitude-dependent rotation. Both these effects will be more severe with transmitted signals have higher peak-to-average ratios (PARs) in the transmitted waveform, as these signal will exercise a wider range of the PA nonlinear transfer function illustrated in Figure 3.1. As an example, a 64-QAM signal with 4dB backoff from the output saturation level is passed through the nonlinear PA described in Figure 3.1. The ideal and corrupted constellations for this example are shown in Figure 3.2; as can be observed, even with 4dB backoff, severe nonlinear constellation compression occurs due to the saturating nature of the PA. Also, a significant amplitude-dependent rotation can also be seen in the constellation.

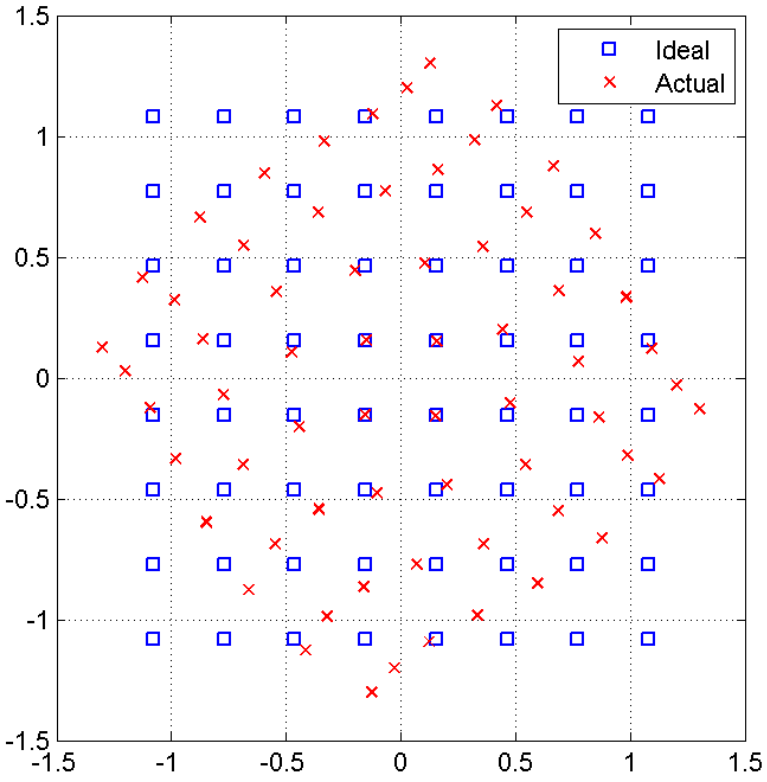

Figure 3.2: Impact of nonlinear PA on QAM64 constellation

3.1.2 Local Oscillator Phase Noise

A local oscillator (LO) is used in both the transmitter and receiver to perform the required frequency translation to and from the 60GHz carrier frequency. An ideal LO will generate a single-tone at the desired frequency, but in actuality, a real LO will typically also have phase noise (among other non-idealities) in its output.

The phase noise of a free-running oscillator (FRO) perturbed by white noise sources can be expressed as [49]:

$$\mathcal{L}_{FRO}(\Delta f) = \frac{K}{(\Delta f)^2} \tag{2.5}$$

where Δf is the frequency offset from the carrier and K is a parameter of the oscillator. Equation (2.5) predicts infinite phase noise power at DC, which implies that the DC phase error of the FRO is unbounded. While this is the case for a FRO, most LOs use some sort of phase-lock mechanism (e.g. a phase-lock loop) to keep the low-frequency phase noise bounded. Also, communications links utilizing coherent detection require some sort of carrier phase recovery (CPR) scheme in order to drive the low-frequency phase noise to a suitably low level. The combination of these two mechanism filters the FRO phase noise by a second-order filter whose poles are at the loop bandwidths of the PLL and CPR block. Therefore, the phase noise seen at the detector is:

$$\mathcal{L}_{DET}(\Delta f) = \frac{K}{(\Delta f)^2} \cdot \left(\frac{\frac{(\Delta f)^2}{f_{o,PLL}^2}}{1 + \frac{(\Delta f)^2}{f_{o,PLL}^2}} \right) \cdot \left(\frac{\frac{(\Delta f)^2}{f_{o,CPR}^2}}{1 + \frac{(\Delta f)^2}{f_{o,CPR}^2}} \right) \tag{2.6}$$

where $f_{o,PLL}$ and $f_{o,CPR}$ are the loop bandwidths of the PLL and CPR, respectively.

Coincidentally, LO phase noise degrades the signal in ways quite similar to PA nonlinearity. Since the LO output is convolved with the incoming signal, the phase noise skirts can cause spectral of the signal leakage into adjacent bands. Similarly, multicarrier systems suffer increased susceptibility to intercarrier interference caused by phase noise.

Also, LO phase noise can cause rotational constellation noise in both single-carrier and multicarrier systems. Any phase noise that is not cancelled by the PLL or CPR will look like carrier phase error at the slicer. Since the uncancelled phase noise changes on a symbol-to-symbol basis, the resultant output constellation will appear to have phase "jitter". An example of the impact of phase noise on a signal constellation is shown in Figure 3.3. In this example, 64-QAM is subjected to 3° RMS phase noise, and the phase jitter is easily observed.

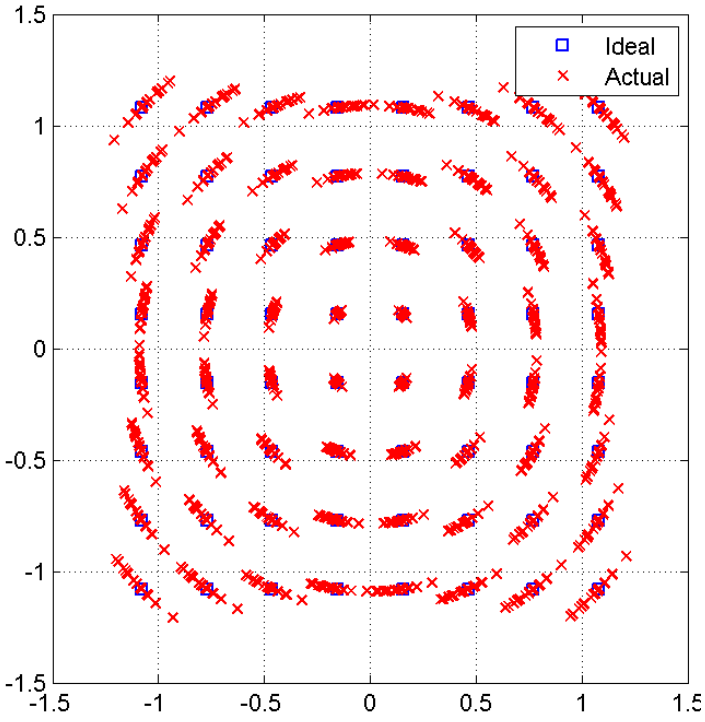

Figure 3.3: Impact of VCO phase noise on QAM64 constellation

3.2 Modulation Schemes

Given the dominant mm-wave circuit limitations discussed above, different candidate modulation schemes can be qualitatively and quantitatively compared. For each of these modulation schemes, a simulation is performed where the signal is passed through a non-ideal PA and LO model as discussed above. No other sources of noise or distortion are considered, as the purpose of these simulations is to determine which modulation scheme is most robust to the limited performance of CMOS mm-wave circuits. (Other factors, such as circuit thermal noise, channel multipath, or ADC quantization error are modeled in section 3.5 to get a more realistic estimate of actual system performance.) Table 3.1 summarizes the parameters used for the PA and LO models.

Power Amplifier		**Local Oscillator**	
p (from Rapp model)	4	**Phase noise at 1MHz offset**	-85dBc/Hz
R_{sat} (from Rapp model)	+1.1dB	**$f_{o,CPR}$**	500kHz
α (from Saleh model)	1.3 radians	**$f_{o,PLL}$**	1MHz
β (from Saleh model)	1		
PA output power	At P_{-1dB}		

Table 3.1: Parameters for the non-ideal circuit models used in simulation

3.2.1 QPSK

QPSK represents the "baseline" case for the choice of modulation scheme. Before rolloff filters are considered, it encodes 2bits/symbol, so it has reasonable bandwidth efficiency, and the SNR required at the slicer for 10^{-3} BER is a moderate 7dB. The peak-

to-average ratio (PAR) of a QPSK symbol is 3dB, so the PA will exhibit some nonlinearity and signal compression. QPSK has been shown to have moderate sensitivity to phase noise, as it suffers a roughly 3dB sensitivity loss in the presence of 10° RMS phase noise [50].

A simulation of a QPSK constellation passing through a nonideal PA and LO is shown in Figure 3.4. Several key points can be observed: first, the gain compression from the PA results in a smaller average amplitude than desired; none of the output symbols have a magnitude equaling that of the ideal constellation points. This will be observed in all the simulations, as one effect of using a nonlinear PA is reduced signal gain. Rotational symbol jitter from LO phase noise can also be observed. Additionally there is a slight rotational tilt in the counter-clockwise (positive phase) direction; this cannot be caused by the phase noise, as the high-pass filtering of the phase noise removes its DC phase content. Instead, this is an artifact of the non-linear AM/PM distortion in the PA, which imparts an increasingly positive phase shift to the signal as its magnitude increases. Lastly, some small amount of AM/AM distortion can be observed in the radial spread of the “constellation cloud”.

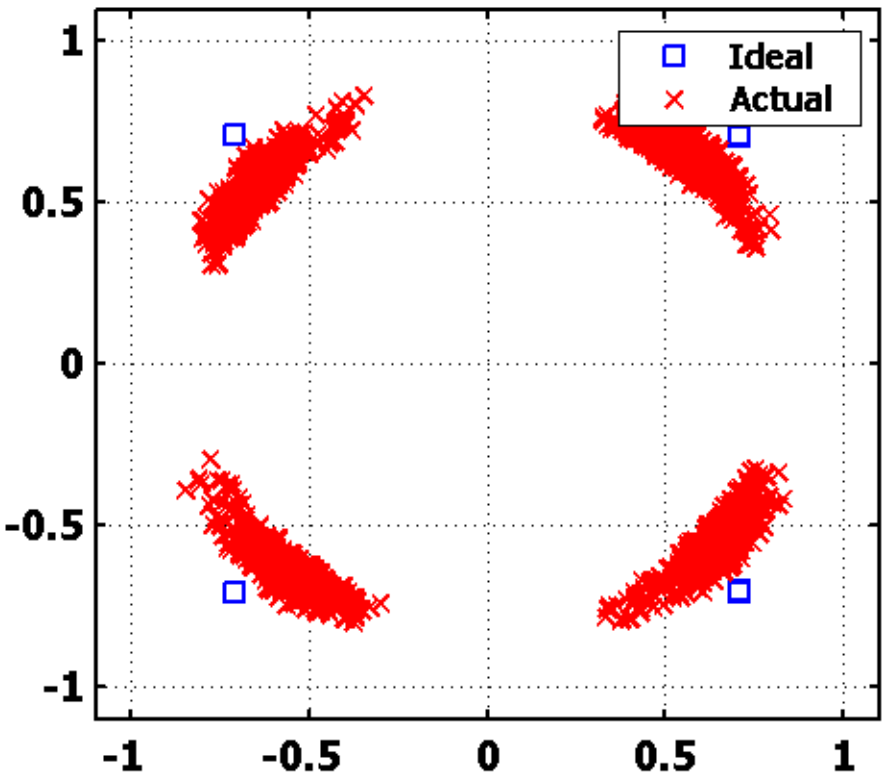

Figure 3.4: QPSK constellation with nonideal PA and VCO

The SNR of the output constellation was measured to be roughly 15dB. As noted earlier, this is before circuit thermal noise, quantization error, or other implementation errors are considered. Therefore, the mm-wave circuit nonidealities should not significantly impact QPSK's ability to function at a 10^{-3} BER ($SNR_{req'd}$ = 7dB), but may limit its capability to function at more stringent bit error rates

3.2.2 Quadrature Amplitude Modulation

QAM is considered here largely as a "straw man". At its core, QAM achieves higher spectral efficiency at the price of high circuit performance requirements, which is

in direct conflict with the previously discussed goals of 60GHz CMOS operation. Therefore, as expected, QAM performance will suffer significantly in the presence of mm-wave circuit nonidealities.

16-QAM is used here as the example QAM scheme. It encodes 4 bits/symbol, but it requires 12dB of SNR at the slicer to achieve 10^{-3} BER, 5dB more than the other schemes considered here. The PAR is roughly 5.5dB, so the linearity of the PA will be significantly tested. Also, since the constellation symbols are much more closely spaced, QAM is more sensitive to the rotational symbol jitter caused by LO phase noise.

A simulation of a 16-QAM constellation passing through a nonideal PA and LO is shown in Figure 3.5. As expected, the output constellation is heavily distorted by the circuit nonidealities. LO phase noise and PA distortion are easy to observe. Individual constellation points can no longer be distinguished from one another, an indication of unacceptably high BER.

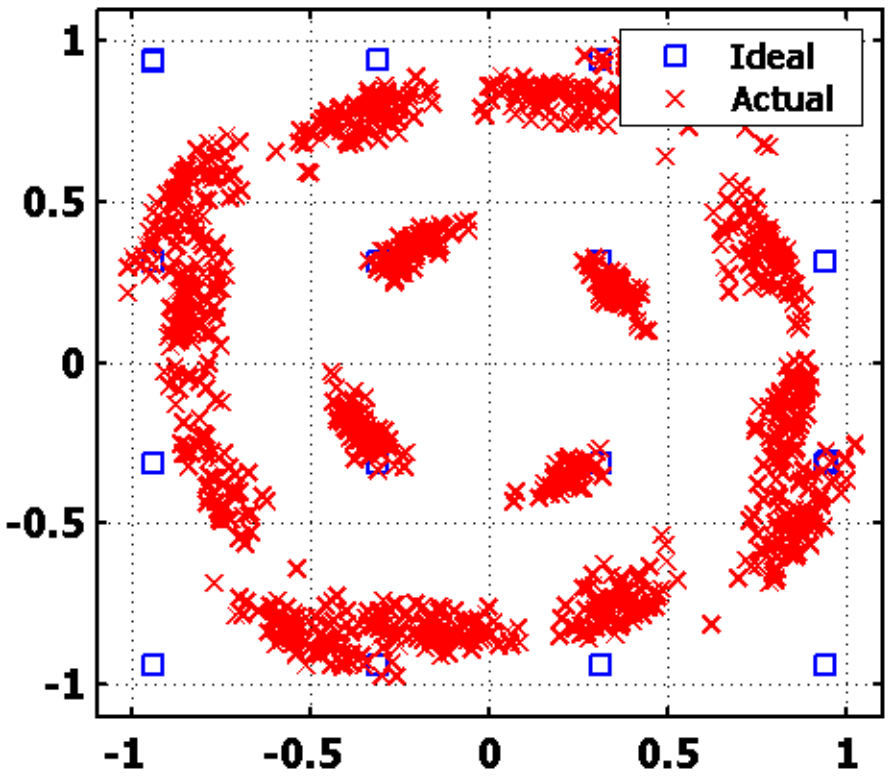

Figure 3.5: QAM16 constellation with nonideal PA and VCO

3.2.3 OFDM

Multicarrier systems such as OFDM have been successfully used in several commercial wireless communications systems, such as 802.11a/g/n, MBOA-UWB, WiMax, and DVB-T. The primary advantage of OFDM is its robustness to multipath; by spreading the communication over many low-bandwidth subcarriers, each subcarrier can be made to look frequency-flat, even in the presence of significant multipath. The low symbol rate enables the use of a guard interval, hence removing any intersymbol interference and greatly easing channel equalization. Also, coding across subcarriers enables the system to perform robustly in the presence of deep channel fades.

The use of multiple subcarriers imposes significant performance constraints on the mm-wave circuitry. In the case of a 60GHz indoor beamformed channel targeting 1Gbps, more than 200 subcarriers would be required with a subcarrier spacing of no greater than 2.2MHz.[2] The large number of subcarriers results in an OFDM signal that has a very high PAR, because the PAR grows with the square-root of the number of subcarriers if no special techniques are employed. Even if PAR reduction techniques are used, the resultant PAR for this system would likely be greater 10dB. Therefore, the PA nonlinearity would likely have a significant impact on the output waveform. The PA nonlinearity could be avoided by utilizing a sufficient amount of PA backoff, but the backoff would cause less signal power to reach the receiver, impairing system throughput.

Also, the performance of an OFDM with small subcarrier spacing would degrade in the presence of LO phase noise. As shown in equation (2.5), LO phase noise increases at smaller offsets from the carrier. Therefore, smaller subcarriers would suffer from increased intercarrier interference cause by the LO phase noise.

A simulation of an OFDM signal constellation passing through our nonideal PA and LO is shown in Figure 3.6. This example waveform uses QPSK subcarriers and is constructed to have a PAR of roughly 10dB. Again, we can see that the nonideal mm-wave circuits degrade the fidelity of the output waveform. While the individual

[2] As discussed in chapter 2, RMS delay spread might be as large as 15ns. Therefore, the channel impulse response can be expected to die down within 3σ, or 45ns. If the guard interval is set to 45ns and is 10% of the OFDM symbol time, then the symbol time is 450ns, corresponding to a 2.2MHz subcarrier spacing. If QPSK is used and 1Gb/s is the throughput target, then at least 500MHz total bandwidth is required. Therefore, at least 225 (=500MHz/2.2MHz) subcarriers are required.

constellation "clouds" are discernable, the calculated SNR of this waveform is only 9.5dB. As mentioned previously, this waveform does not include any other noise or distortion effect other than those from the PA and LO. Since QPSK-OFDM requires an SNR of 7dB to hit 10^{-3} BER in an ideal AWGN channel, the performance degradation caused by the nonideal PA and LO are likely to limit overall system performance in a real implementation environment. Therefore, given the assumptions about the performance of available mm-wave CMOS components [16], an OFDM system would not be able to meet the system-level performance targets.

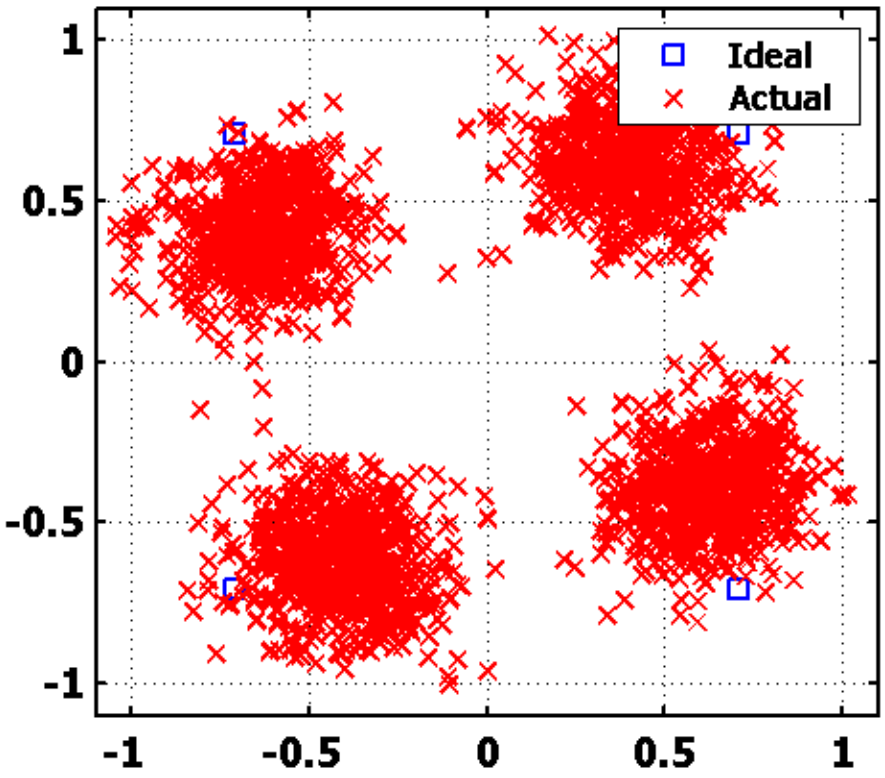

Figure 3.6: OFDM constellation with nonideal PA and VCO

3.2.4 Minimum-Shift Keying

Minimum-Shift Keying (MSK) is a specific type of Frequency-Shift Keying (FSK) where the modulation index, h, is equal to 0.5. MSK modulation ensures both a constant envelope and a continuous phase at symbol boundaries. The bandpass representation of an MSK symbol is:

$$s_{RF}(t) = \cos\left[2\pi\left(f_c + \frac{1}{4T}d_n\right)t - \frac{1}{2}n\pi d_n + \theta_n\right], \quad nT \le t \le (n+1)T$$
$$\theta_n = \frac{1}{2}\pi \sum_{k=-\infty}^{n-1} d_k \tag{2.7}$$

where d_n is the data symbols (±1), T is the symbol period, and θ_n is the accumulated phase trajectory of all the past data symbols. Translating the signal at carrier down to its baseband equivalent, an MSK symbol takes the form of (2.8). Figure 3.7 shows the baseband equivalent of a sample MSK waveform. It can be seen that MSK looks very much like staggered QPSK, but with sinusoidal pulse shape, g(t). Not surprisingly, MSK exhibits the same power efficiency as QPSK, so it requires SNR of 7dB at the slicer to maintain a 10^{-3} BER.

$$s_{BB}(t) = \sum_{n=-\infty}^{\infty} [d_{2n}\, g(t - 2nT) - j \cdot d_{2n+1}\, g(t - 2nT - T)]$$
$$g(t) = \begin{cases} \sin\dfrac{\pi t}{2T}, & 0 \le t \le 2T \\ 0, & \text{otherwise} \end{cases} \tag{2.8}$$

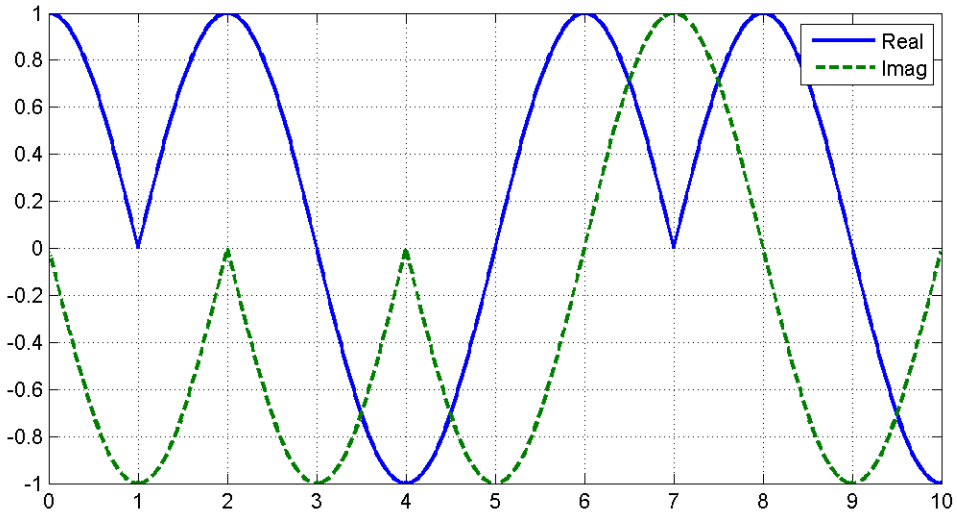

Figure 3.7: Baseband equivalent of sample MSK waveform

The primary advantage of utilizing MSK is its property of constant envelope. Since the information is encoded in the phase and frequency of the signal, rather than in its amplitude, the PAR of an MSK signal is 0dB. Therefore, as mentioned in section 3.1.1, the PA nonlinearity only results in a (slightly) reduced PA gain, but no signal distortion occurs.[3] Furthermore, it has been shown that MSK is significantly more robust to phase noise than QPSK. Whereas 10° RMS phase noise results in a 3dB sensitivity loss for QPSK, the same phase noise only leads to a 1dB sensitivity loss for MSK [50]. Compared to QPSK, MSK does have slightly lower spectral efficiency, but this is not a significant constraint in the 60GHz band.

[3] In reality, other circuits in the transmit path—such as filters with a non-constant group delay—may increase the PAR of the signal at the PA a little beyond 0dB. In that case, the nonlinearity of the PA will cause constellation distortion. However, careful design of the other transmit path circuits can keep the PAR low enough so that this second-order effect is negligible.

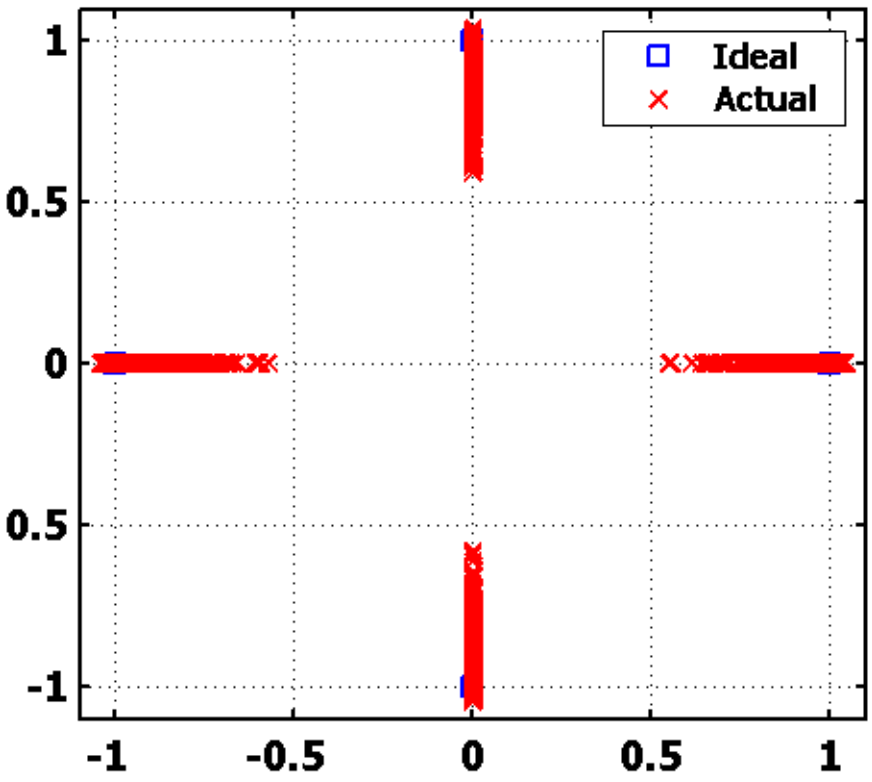

Figure 3.8: MSK constellation with nonideal PA and VCO

A simulation of an MSK signal constellation passing through our nonideal PA and LO is shown in Figure 3.8. As expected, the nonlinear PA only causes a slight reduction in gain, and the LO phase noise creates some spreading of the constellation points. The SNR of the output constellation was measured to be roughly 24dB. Therefore, the mm-wave circuit nonidealities should not significantly impact MSK's ability to function at BERs well below 10^{-6} BER ($SNR_{req'd}$ = 14dB).

3.2.5 Comparison

Table 3.2 summarizes the relevant characteristics of the different candidate modulation schemes. As expected, high-order modulation schemes (i.e. 16-QAM) are not suitable for 60GHz implementation due to the stringent requirements placed on the various mm-wave components; the 16-QAM constellation is irrecoverably distorted by the nonidealities in the PA and LO. The high peak-to-average ratio and tight subcarrier spacing of OFDM also imposes difficult requirements on the mm-wave components; after the inclusion of PA and LO nonidealities, the output signal is degraded to an SNR of 9.5dB, before other limitations like circuit noise are even considered. In contrast, an MSK scheme is quite robust to the mm-wave circuit issues discussed previously. The constant envelope property of MSK signaling makes in immune to PA nonlinearity, and its signal structure makes it much less sensitive to phase noise than other modulation schemes. Furthermore, MSK has the same power efficiency as its QPSK and OFDM counterparts; therefore, no additional burden is placed on the noise requirements of the receiver circuits. The only significant downside of utilizing an MSK signaling scheme is the hardware complexity required to equalize the channel. Unlike OFDM, which uses an efficient FFT-based algorithm to manage channel multipath, an MSK scheme (and most other single-carrier schemes) requires a relatively inefficient time-domain equalizer. For channels that have very long channel impulse responses, the extra hardware required for an equalizer can be prohibitively expensive; however, as discussed in Chapter 2, the use of directive antennae at 60GHz can reduce the channel multipath to levels that can be reasonably handled with an equalizer in a single-carrier system.

Modulation	Single-carrier QPSK	High-order modulation (16-QAM)	OFDM-QPSK	Constant Envelope (MSK)
SNR_{req} (BER=10^{-3})	7dB	12dB	7dB	7dB
PAR_{TX}	~3dB	~5.5dB	~10dB	0dB
PA linearity req't	Moderate	High	High	Low
Sensitivity to Phase Noise	Moderate	High (Symbol Jitter)	High (ICI)	Low
SNR_{out} after RF nonidealities	15dB	Bad	9.5dB	24dB
Channel Equalization Complexity	High (Equalizer)	High (Equalizer)	Moderate (FFT)	High (Equalizer)

Table 3.2: Comparison of modulation schemes

Therefore, MSK is deemed the most suitable modulation scheme for a 60GHz 1Gbps indoor wireless link. MSK provides a good compromise between mm-wave performance requirements and baseband architecture complexity. Issues and challenges with implementing the baseband architecture with minimum complexity and power dissipation are discussed in the next section.

3.3 Baseband Architecture Considerations

The previous section compared several different candidate modulation schemes for the 60GHz, 1Gbps indoor wireless link and settled upon MSK as the most suitable scheme for an all-CMOS implementation. This section discusses the functional requirements of the MSK baseband receiver and discusses various implementation tradeoffs in the baseband design.

3.3.1 Performance Requirements of MSK Baseband Receiver

The primary goal of the MSK receiver is to detect and demodulate the received MSK signal with a suitably low bit-error rate (BER). The chosen design target is an uncoded BER of 10^{-3}; if this BER target is met, then a coding scheme can be added that will reduce the BER to acceptable levels for wireless transmission. A secondary goal of the receiver design is achieve this BER target with a minimum received SNR. As mentioned previously, MSK in an ideal AWGN channel requires 7dB SNR to achieve a BER of 10^{-3}, representing a lower bound for required SNR.

In order to be able to suitably detect the received signal, the receiver needs to be appropriately configured in order to demodulate the incoming signal properly. A set of tasks, referred to in this work by the term *synchronization*, must be performed in order to properly configure the receiver to its environment. These tasks include functions such as symbol timing recovery, automatic gain control (AGC), carrier phase and frequency recovery, and channel response equalization. Symbol timing recovery and AGC are standard functions that must be incorporated into almost every digital communications receiver in existence; carrier recovery is necessary because MSK requires coherent detection in order to perform correctly without significantly sensitivity loss; as will be discussed below, channel equalization is required because the multipath observed in indoor 60GHz channels is significant, even with the use of directive antennae.

As discussed in Chapter 2, indoor 60GHz environments suffer from multipath. Directional antennae can help minimize the amount of multipath observed, but this technique alone is not sufficient to create a receiver that is robust to various indoor

environments. The channel study conducted in [10] indicated that indoor environments with 60° HPBW antennae exhibit RMS delays spreads ranging from about 5ns to 15ns and Rician K ranging from 15dB to 5dB, with the variation caused by the local physical environment. Figure 3.9 shows the performance of an ideal MSK receiver without an equalizer in a variety of typical indoor multipath environments [10]. Each simulation data point is the ensemble average of 20 Monte Carlo runs of 10000 MSK symbols each. Each Monte Carlo run has a randomly generated, time-invariant channel impulse response with the channel parameters as noted. Perfect carrier phase recovery and timing synchronization are assumed. As can be seen in this figure, even at a "best case" channel environment, the MSK receiver suffers a 2dB performance penalty at 10^{-3} BER levels when an equalizer is not used. At more realistic channel response scenarios, the performance degradation is quite severe: BER's of 10^{-3} cannot be achieved, as the BER error floor for {K=5dB, Trms=15ns} channels is roughly 2×10^{-2}, an unacceptably high level.

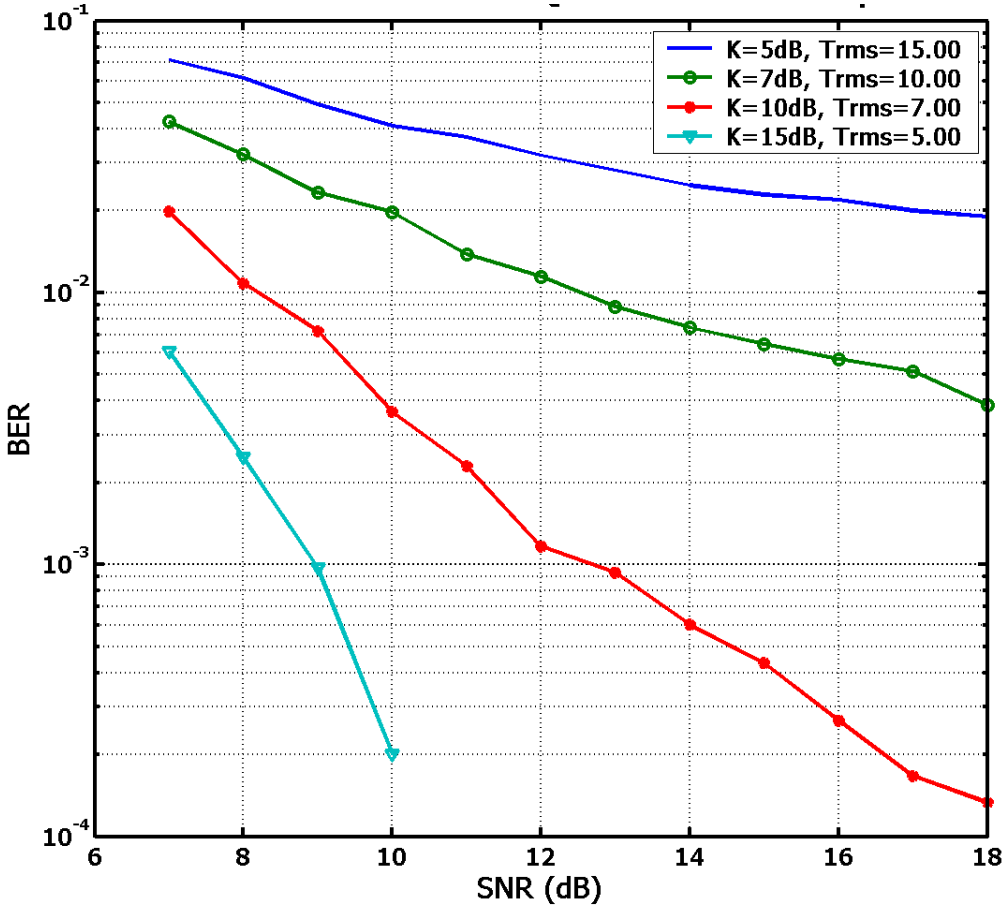

Figure 3.9: BER vs. SNR for MSK without an equalizer

The performance of Figure 3.9 can be contrasted to the performance of an MSK receiver with an ideal equalizer, shown in Figure 3.10. As before, perfect carrier phase recovery and timing synchronization are assumed. As expected, there is a slight performance penalty at channels with large amounts of multipath, but the SNR required for 10^{-3} BER is still roughly 9.25dB even in the most dispersive channels. Therefore, an equalizer is required for the MSK receiver, as the performance penalty for not including one is far too steep.

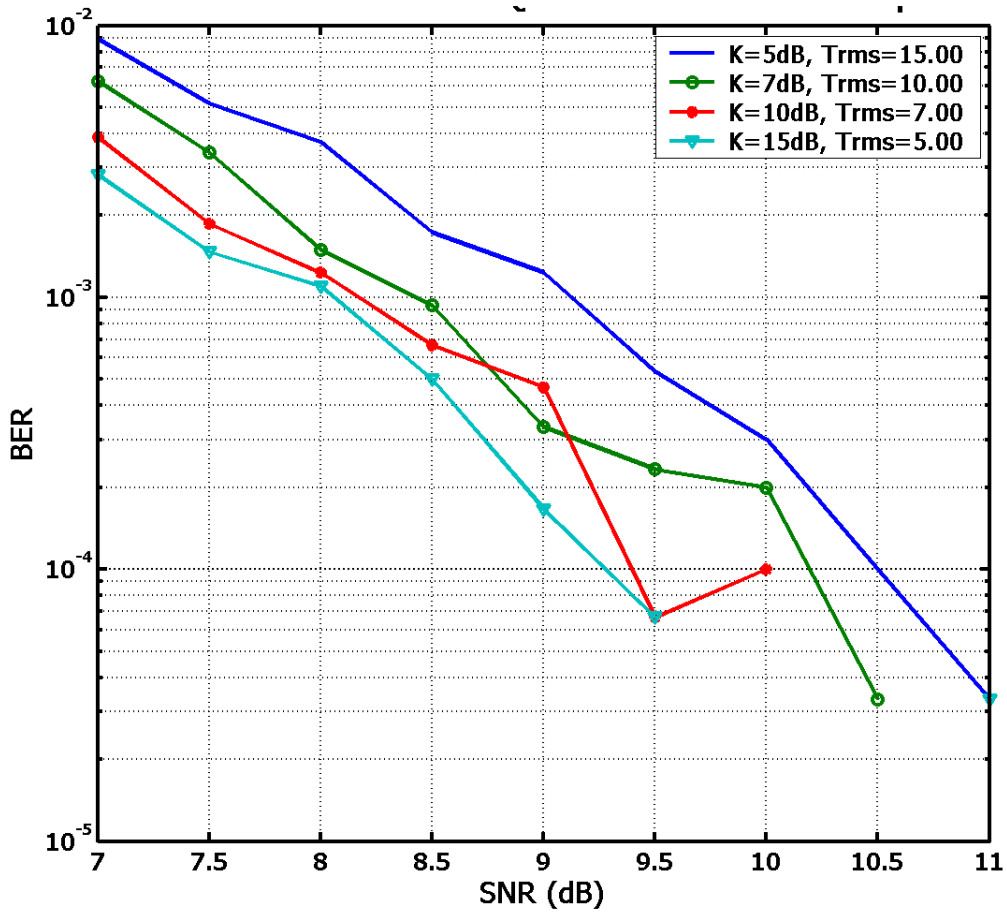

Figure 3.10: BER vs. SNR for MSK with an equalizer

In the previous simulations of MSK performance without and with an equalizer, perfect carrier and timing synchronization were assumed. In reality, phase and timing sync will have some residual errors, which will impact the performance of the receiver. Figure 3.11 shows the performance of the MSK receiver in the presence of carrier phase error; these simulations assume a {K=5dB, Trms=15ns} channel and the presence of an equalizer. It can be observed that there is minimal performance degradation for carrier phase errors up to roughly 5°.

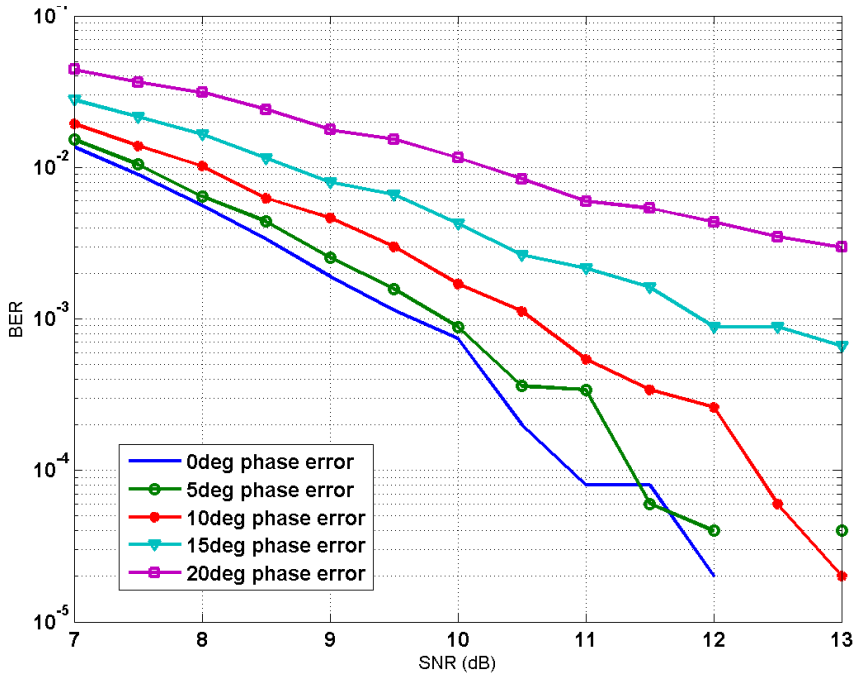

Figure 3.11: MSK performance in K=5dB, Trms=15ns channel with carrier phase error

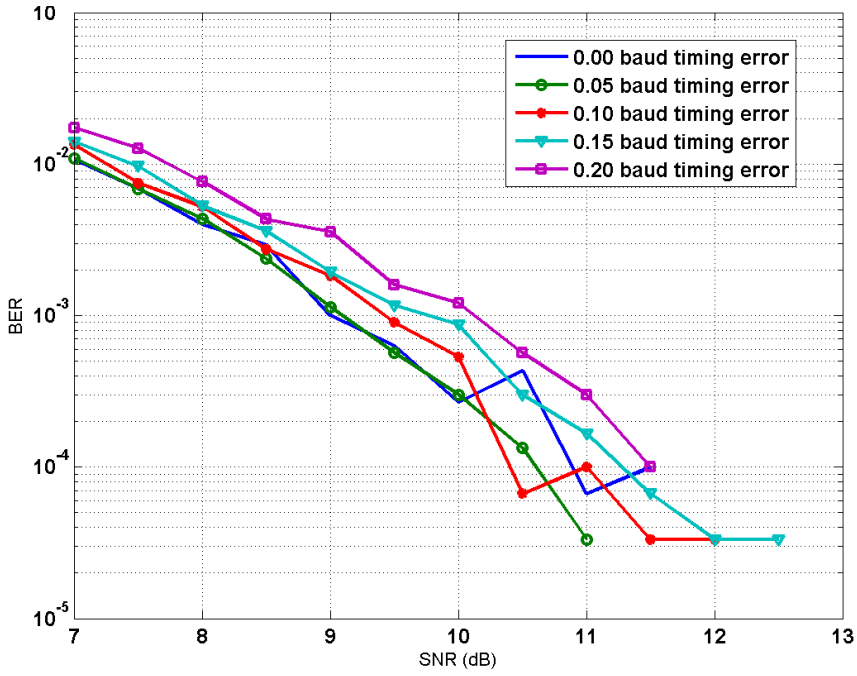

Figure 3.12: MSK performance in K=5dB, Trms=15ns channel with symbol timing error

Figure 3.12 shows the performance of the MSK receiver in the same channel in the presence of symbol timing offset. Performance is surprisingly robust to symbol timing errors; a timing error of 0.1 baud (100ps, given a 1ns symbol period) degrades receiver sensitivity by less than 0.5dB.

In summary, the baseband receiver must be capable of demodulating and detecting the received MSK signal with a BER below 10^{-3}. In order to make that goal achievable, the receiver requires a set of synchronization functions to estimate and correct for channel multipath, carrier phase/frequency offset, and symbol timing offset. These synchronization blocks are used to properly condition the received signal such that detection can be performed with a minimal loss in receiver sensitivity. The next section discusses candidate architectures for implementing the MSK baseband receiver.

3.3.2 Baseband Architecture Selection

As discussed previously, the baseband receiver must be capable of synchronizing to the received signal and then demodulating and detecting it. Since MSK is a low-order modulation scheme and the desired throughput is 1Gbps, the resultant MSK symbol rate is 1Gsym/s. This high baud rate requires high-speed circuitry in order to process the incoming signal. In particular, the converter circuitry (i.e. the analog-to-digital converters) are required to run at some multiple of the symbol rate. Since high-speed ADCs can dominate the overall power consumption of the baseband receiver, an architecture that minimizes the performance requirements of the ADCs is required. Therefore, techniques that can minimize the ADC resolution requirements or the sampling rate requirements can significantly decrease the total power consumption of the

baseband receiver. However, these techniques cannot result in a significant receiver performance loss. Also, any architecture employed must be adaptable and robust to the varying channel conditions that it might encounter.

3.3.2.1 The Role of Synchronization in a Receiver

Before different candidate architectures can be considered, the concept of synchronization must be expanded upon. The concept of synchronization was first introduced in section 3.3.1. In this dissertation, the term "synchronization" is used to refer to all the auxiliary tasks the receiver must perform in order to be able to correctly demodulate and detect the incoming signal with optimal sensitivity. Therefore, symbol timing synchronization, carrier phase/frequency recovery, automatic gain control, and channel equalization are all considered "synchronization" functions in this work. For the sake of clarity, this work subdivides the synchronization task into two subtasks, dubbed "estimation" and "correction". The estimation subtask involves estimating the synchronization parameter that needs to be accounted for; the estimation subtask operates on the incoming signal to calculate its estimate, but it does not alter the received signal in any way. The correction subtask uses the estimate provided by the estimation subtask to actually correct the received signal to properly synchronize the receiver.

An example or two can help illustrate this distinction. In a DFE, the estimation subtask consists of the (typically LMS) adaptation routine that is used to estimate the channel impulse response. The correction subtask would then correspond to the actual signal processing that equalizes the received signal to remove the channel multipath.

Similarly, in an AGC block, the estimation subtask would consist of something akin to a peak detector that would determine the appropriate gain required by the VGA, whereas the correction subtask would correspond to the gain provided by the VGA itself.

With this definition of synchronization in mind, different baseband architectures can be considered.

3.3.2.2 Signal Partitioning: Digital vs. Analog

The different types of architectures under consideration vary predominantly in where the synchronization subtasks are performed. In particular, each architecture contains a different partitioning of the required signal processing between the analog and digital domains. In short, there are 3 types of architectures under consideration, dubbed "digital", "mixed-signal", and "analog" architectures, based on the overall partitioning of the synchronization functions.

3.3.2.3 "Mostly Digital" Architecture

What is termed here as a "mostly digital" architecture is the type of architecture that has become typical in most digital wireless receivers. As depicted in Figure 3.13, very little synchronization functionality is performed in the analog domain; instead, the received signal is down-converted by the RF front-end and converted into the digital domain with minimal further analog processing aside from filtering and gain. Both synchronization subtasks—estimation and correction—are performed in the digital domain, prior to digital demodulation and detection.

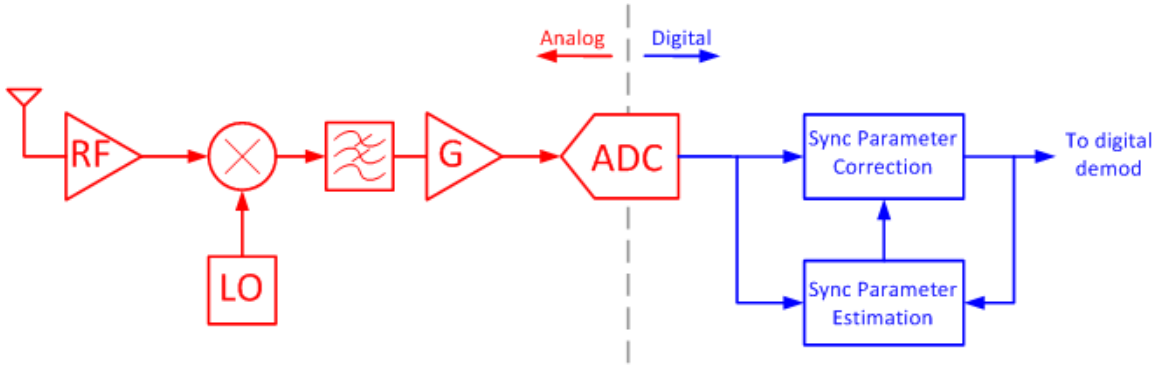

Figure 3.13: "Mostly digital" baseband architecture

There are several well-known advantages to using this mostly-digital approach. First, since much of the signal processing is performed in the digital domain, highly flexible and robust digital algorithms can be employed for synchronization and detection. The estimation subtask of the synchronization procedure can utilize maximum likelihood algorithms to generate highly accurate estimates of the synchronization parameters. Also, digital signal processing does not suffer from circuit noise as analog circuits do, resulting in high-fidelity signal processing.

However, a key downside of the mostly digital approach is the stringent performance requirements that it typically places on the interface circuitry, most notably the analog-to-digital converter. Since this architecture performs most of its signal processing in the digital domain, the signal is not properly synchronized at the ADC input. As a result, the ADC typically requires increased resolution and oversampling to accommodate the increased dynamic range seen at the ADC input. For example, it is entirely possible for a receiver to have an all-digital AGC loop that includes a digital

VGA; however, this system will need a very high resolution ADC in order to accommodate for the fact that the AGC loop does not gain up small signals seen at the ADC input. Similarly, a fully digital DFE will require increased resolution in the ADC as well; since the channel multipath has not yet been corrected at the ADC input, the ADC will see have to convert a signal with an increased dynamic range, necessitating additional ADC bits (see Figure 3.16 on page 79). As mentioned previously, the increased ADC resolution requirements can have significant impact on the overall power dissipation of the receiver.

3.3.2.4 "Mostly Analog" Architecture

At the other end of the spectrum from the "mostly digital" architecture is the "mostly analog" architecture, shown in Figure 3.14. In this architecture, both synchronization subtasks—estimation and correction—are performed in the analog domain. Examples of this sort of architectures are common in some of the early all-analog systems, such as FM and AM radio and some earlier analog cellular standards.

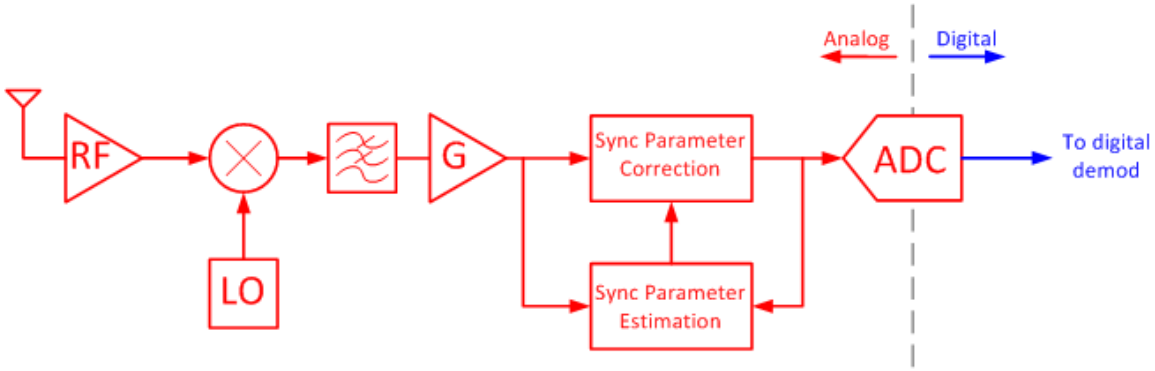

Figure 3.14: "Mostly analog" baseband architecture

A benefit to this approach is that the received signal is conditioned prior to conversion by the ADC; therefore, the performance requirements of that block are reduced, resulting in a significant power savings. However, there are several disadvantages to this approach. The estimation subtask of synchronization is now performed in the analog domain; as a result, only simple analog structures can be used to estimate the synchronization parameters and maximum likelihood algorithms could no longer be employed. For example, an AGC loop would require an analog peak detector, and carrier recovery would be performed with some of analog frequency locking mechanism like a Costas loop. However, it is difficult to even envision an analog block that could perform the channel response estimation that is necessary for channel equalization. Furthermore, since all the calculations performed by the estimation blocks are done in the analog domain, the parameter estimates are corrupted by circuit noise.

In contemporary digital wireless systems, this “mostly analog” approach is never used, as the digital modulation techniques used in these systems do not map well onto “mostly analog” architectures. However, it is mentioned here to highlight the benefits of performing the correction subtask in the analog domain.

3.3.2.5 Mixed-signal Architecture

A third candidate architecture is dubbed a mixed-signal architecture because it implements its synchronization functions in a mix of analog and digital circuitry (Figure 3.15). These systems have the estimation subtask implemented in the digital domain in order to harness the processing power of digital estimation algorithms. However, these

systems implement the correction subtask in the analog domain in order to properly condition the received signal prior to quantization, thus minimizing the performance requirements on the ADCs. Therefore, the mixed-signal architecture achieves the benefits of both the "mostly digital" and "mostly analog" approaches without suffering from either systems weaknesses. These types of architectures typically show up in high-speed systems such as disk-drive read channels [51] or high-speed serial link receivers [52], where the large bandwidths required motivate the use of techniques that reduce the ADC requirements.

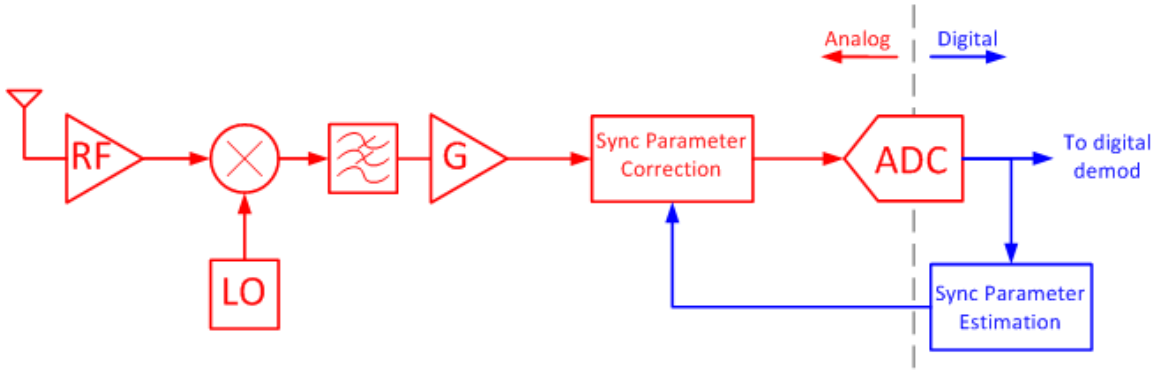

Figure 3.15: Mixed-signal baseband architecture

Implementing a mixed-signal architecture is not without its challenges, however. As can be seen in Figure 3.15, all the synchronization functionality must be implemented using feedback loops. This is due to the fact that the estimate from the digital estimation block must be fed back to control the analog correction block. The use of a feedback loop in the synchronization functionality limits some of the digital synchronization algorithms that can be employed by the estimation block; in particular, algorithms that have a large amount of latency to generate their estimate cannot be used since the excess latency can

destabilize the feedback loop. Also, since there will by necessity be several different synchronization loops present (i.e. timing, carrier phase, and channel equalization), careful design must be used to ensure that the loops do not couple into one another, increasing estimation error and potentially destabilizing one or more loops.

3.3.2.6 Comparison of Digital and Mixed-signal Architectures

A cursory glance of contemporary wireless receivers will show that most commercial receivers utilize a "mostly-digital" approach. In fact, one of the current "holy grails" of radio system design is the notion of a software radio, consisting of little more than an antenna and an ADC. While these sorts of systems can offer a high degree of flexibility, this flexibility comes at the price of increased circuit performance requirements, as detailed in the previous section.

A qualitative comparison of a mostly digital architecture to a mixed-signal architecture can be made by comparing a fully digital equalizer with a mixed-signal equalizer, as shown in Figure 3.16 and Figure 3.17, respectively. (For simplicity, these examples assume the channel gain is 1.) In the fully digital DFE, the multipath channel increases the dynamic range of the received signal, as shown in Figure 3.16. The ADC directly quantizes this multipath-corrupted signal, requiring increased resolution in the ADC to accommodate the increased dynamic range. After quantization, the digital equalizer removes the multipath components, thus clearing up the eye diagram.

The mixed-signal DFE is shown in Figure 3.17. As before, the multipath channel increases the dynamic range of the received signal, but the analog equalizer removes the

signal multipath before quantization. Therefore, the dynamic range of the input signal is reduced in comparison to the digital scheme, and hence a lower resolution ADC can be used. In both schemes, an estimation block in the digital domain provides the equalizer with an estimate of the channel multipath.

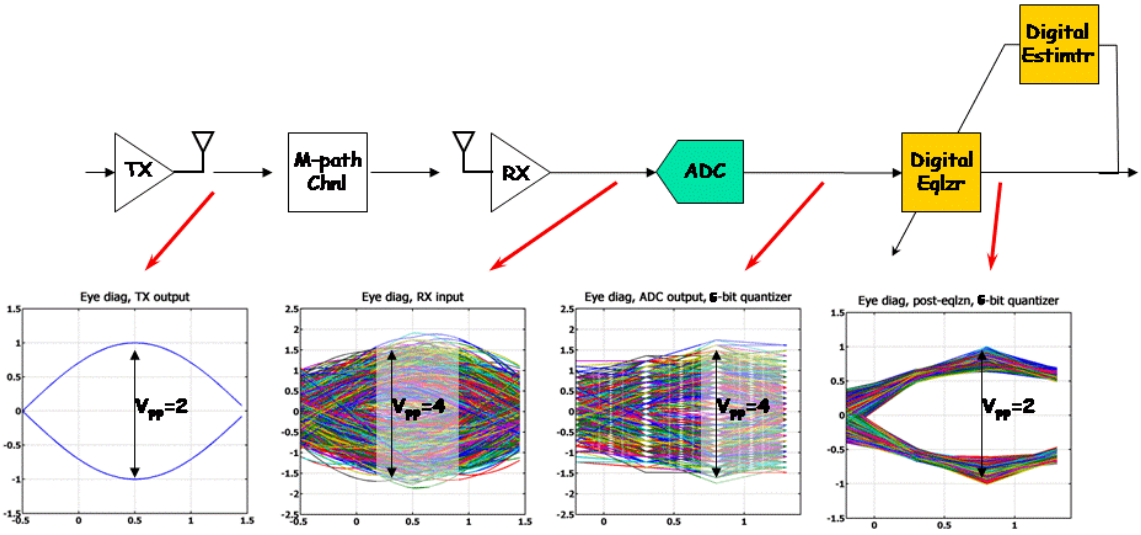

Figure 3.16: Signal levels with a digital equalizer

System level performance of these two candidate architectures can also be compared on a quantitative basis. Receiver sensitivity tests with the standard {K=5dB, Trms=15ns} channel were performed, and the results are shown below. With a "mostly digital" architecture, the system required ADCs with at least 6-bits of resolution in order to have a minimal impact on receiver sensitivity (Figure 3.18). In comparison, a mixed-signal architecture required ADCs with only 4-bits of resolution (Figure 3.19).

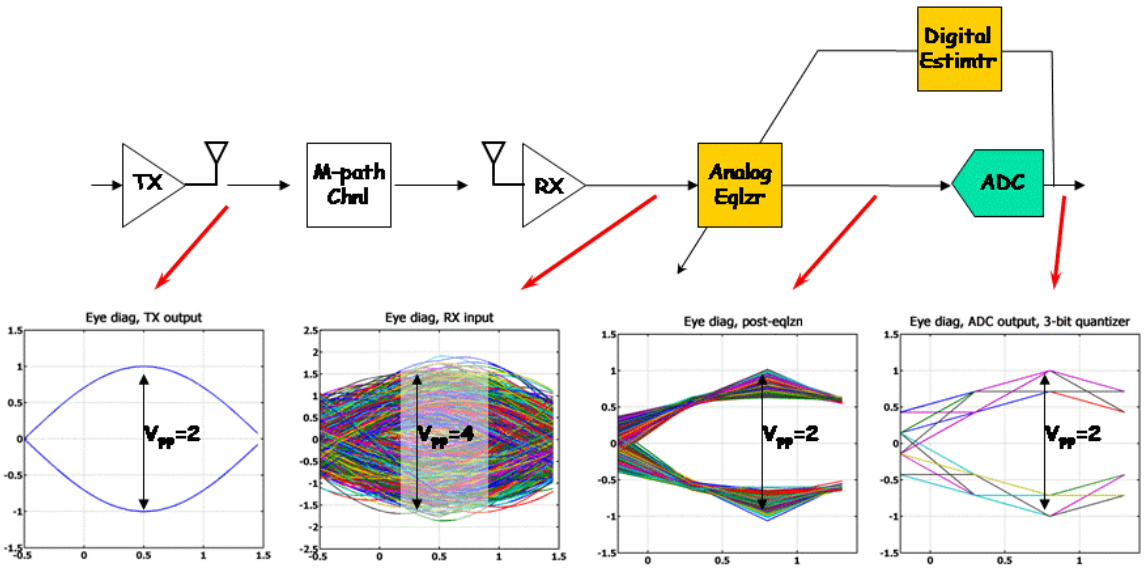

Figure 3.17: Signal levels with a mixed-signal equalizer

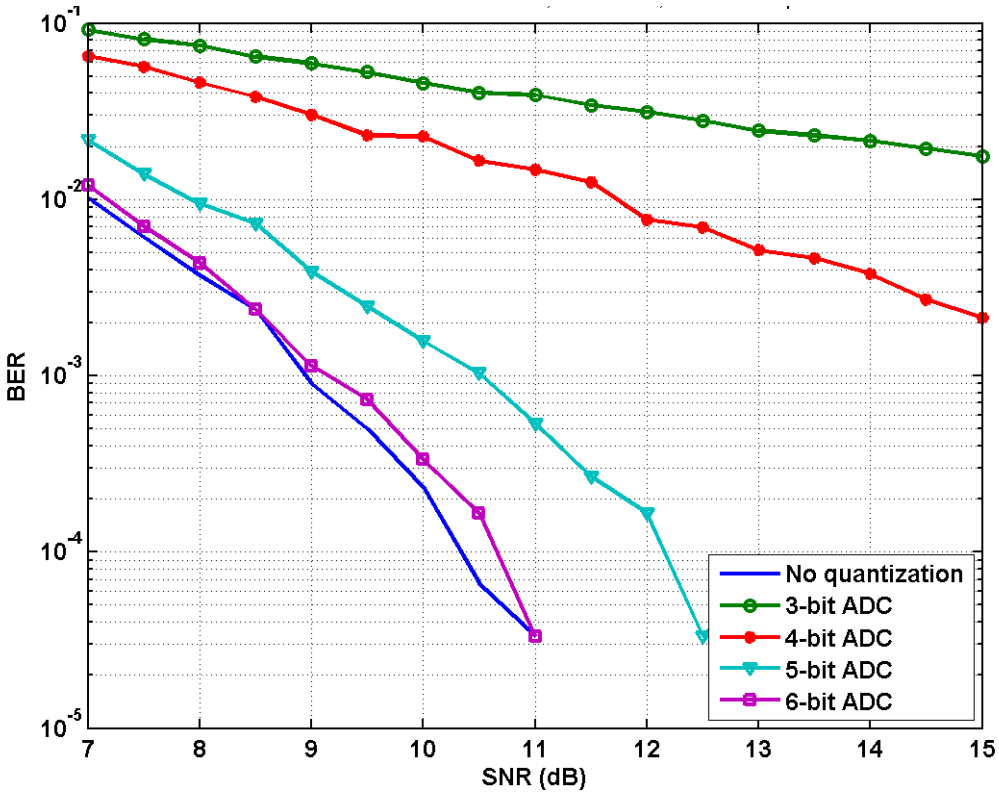

Figure 3.18: BER vs. ADC resolution for a "mostly digital" architecture

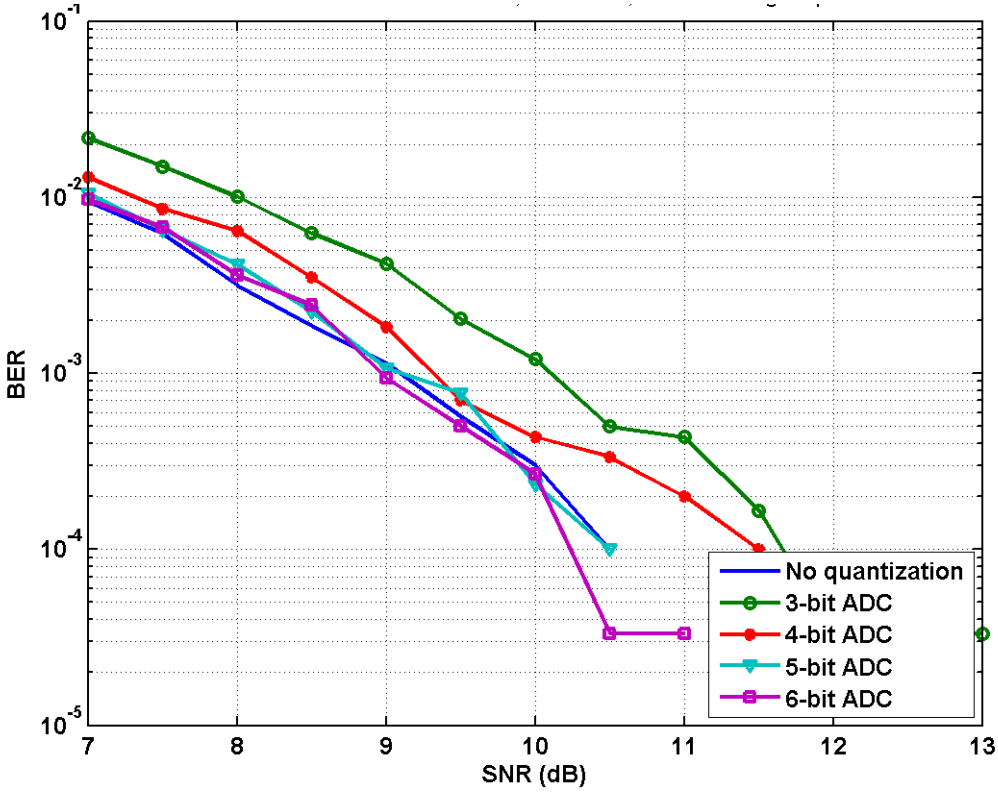

Figure 3.19: BER vs. ADC resolution for a mixed-signal architecture

A brief review of state-of-the-art 6-bit ADCs with sampling speeds above 1Gs/s indicates that a 6-bit 2Gs/s ADC will likely consume more than 50mW. (See Table 3.3) Since coherent MSK detection requires a quadrature baseband, two ADCs would be required, representing over 100mW of power dissipation for the "mostly digital" solution. As will be demonstrated in chapter 4, a 4-bit 2Gs/s ADC can be implemented with drastically lower power dissipation. Also, as will be demonstrated in the next chapter, the additional analog circuitry required to implement a mixed-signal receiver architecture is relatively simple and consumes tens of milliwatts. Therefore, for this MSK receiver

architecture, a mixed-signal architecture looks promising for a low overall system power dissipation.

Ref	F_s (GHz)	# bits	Power (mW)	Projected power of 2 ADC's running at 2Gs/s (mW)
[53]	1.0	6	55	220
[54]	1.25	6	32	104
[55]	3.5	6	98	112

Table 3.3: Recently reported high-speed (>1Gs/s), 6-bit ADCs

A simplified block diagram of the proposed receiver architecture is shown in Figure 3.20, with the baseband section enclosed in the gray rectangle. The signal is down-converted from RF, and then after passing through filters and a VGA, the quadrature signal is passed onto the proposed mixed-signal baseband receiver architecture. A carrier phase rotator (CPR) is the first analog block that the signal passes through; this block corrects for any phase or frequency offset between the transmit and receive local oscillators. After the CPR block, the signal is passed onto the decision feedback equalizer, which removes the multipath reflections from the input signal. Only after these signal processing operations in the analog domain is the signal digitized by the two 4-bit 4Gs/s ADCs. The output of the ADCs is used by the digital estimation blocks to determine the estimate for the symbol timing, channel multipath, and carrier phase/frequency offset. These digital estimates are, in turn, fed back to the appropriate analog correction block in order to appropriately synchronize the received analog signal. In addition, the digital output of the ADCs is passed onto the digital demodulation block (not shown) for demodulation and detection.

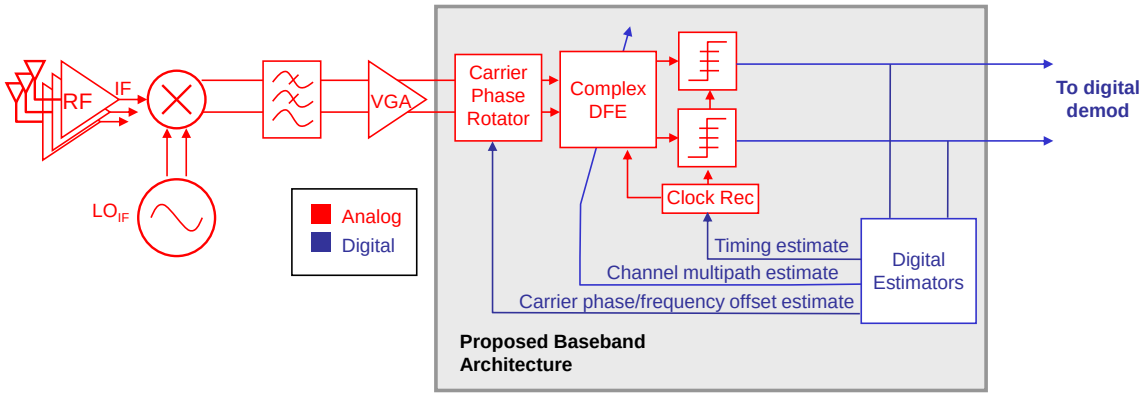

Figure 3.20: Simplified block diagram of proposed mixed-signal receiver architecture

3.4 Synchronization Algorithms for Mixed-signal Receivers

The previous section discussed the functional requirements of the baseband receiver and motivated the selection of a mixed-signal architecture due to the fact that it can achieve high fidelity synchronization and estimation using robust digital algorithms while simultaneously achieving low system power consumption to the reduced performance requirements on the analog interface circuitry (i.e. ADCs). This section discusses the digital synchronization algorithms required for implementing a mixed-signal architecture. In the previous section, perfect estimation of the synchronization parameters was assumed as a precondition to the actual correction performed by the mixed-signal blocks. In this section, the digital estimation algorithms required to generate high accuracy estimates of the synchronization parameters are discussed.

3.4.1 Feedback Synchronization Algorithms for Mixed-signal Systems

A synchronization block can be implemented as either a feedforward or feedback system, depending on whether the estimation block uses the data at the correction block's input or output, respectively [56]. (See Figure 3.21.) As discussed in section 3.3.2, a mixed-signal receiver architecture requires the use of feedback configurations. This is because the estimator block requires digital input data, and the digitization step in a mixed-signal architecture occurs after the correction of the synchronization parameter, which occurs in the analog domain. Therefore, classical feedback theory can be used to analyze and understand the feedback synchronization algorithms used in mixed-signal systems.

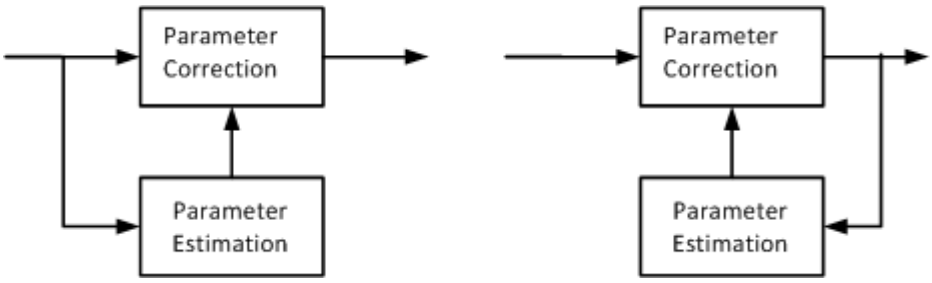

Figure 3.21: Feedforward (left) and feedback (right) estimation

A conceptual block diagram of a feedback synchronizer is shown in Figure 3.22. The uncorrected synchronization parameter is represented by the input, x_i. This term could represent any of the synchronization parameters discussed previously, such as the carrier phase, the symbol timing, or a single multipath channel "tap" value. The role of the correction block is to drive the output, x_e, to zero by subtracting the estimate of the

sync parameter from the input value. This parameter estimate is created by the estimation block; the estimation block, in turn, uses the output from the ADC to generate the parameter estimate, $\hat{x}_i$, thus completing the feedback loop. The estimation block typically consists of a noisy detector, which creates a low-SNR estimate of the residual parameter error, x_e, and a filter, which bandlimits the detector output to create the final parameter estimate. The filter typically has one or more poles at zero; as will be shown below, the infinite DC gain of this filter is used to drive the DC value of the residual error to zero.

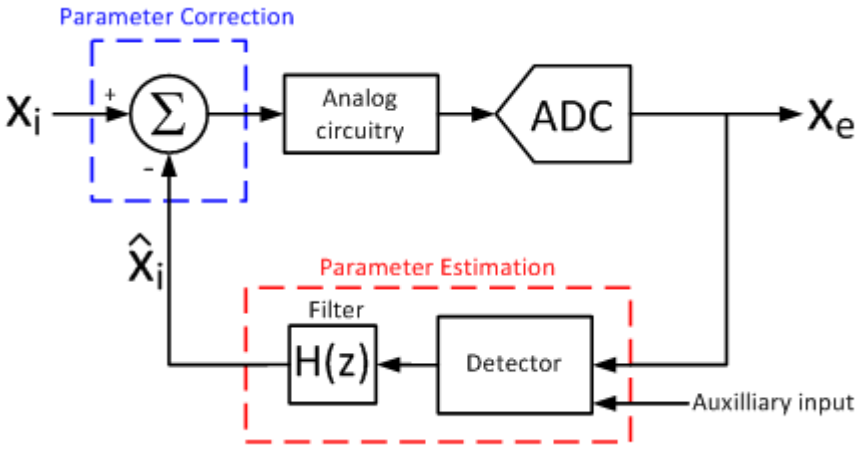

Figure 3.22: Simplified model of feedback synchronization loop

The detector also has an auxiliary input that it uses to create better estimate of the synchronization parameter. In general, the type auxiliary input used further classifies the synchronization algorithm into one of three types [56]: if the auxiliary input is fixed pattern based upon a predetermined training sequence that is used to aid initial synchronization, the synchronization is called *data-aided* (DA); if the auxiliary input is the sliced data decision from the detector (such as in the case of a DFE), the process is

referred to as *decision-directed* (DD); and if no additional data is used, the algorithm is called *non-data-aided* (NDA) or *blind*.

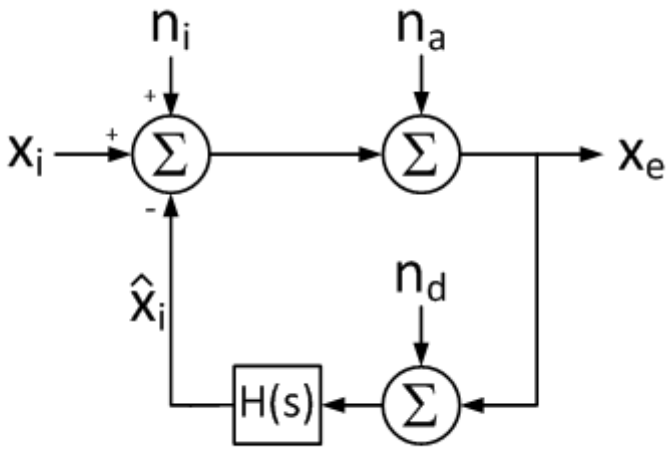

Figure 3.23: Linearized model of feedback synchronization loop

A linearized model of the feedback synchronization loop is presented in Figure 3.23. Here, the feedback loop is fully linearized in order to facilitate analysis, and noise sources are added to indicate all the sources of noise: a noise term, n_i, is used to represent the noise that is present at the input to the system; n_a is used to represent the additional noise that is added by the analog circuits (including the ADC) in the feedforward path of the loop; and n_d represents the noise added by the detector circuit within the estimation block. To simplify the analysis, two assumptions are made: the first is that the filter in the feedforward path of the circuit is sufficiently wideband that it can be ignored; the second assumption is that the digital filter in the estimation block, H(z), can be approximated by an s-domain filter, H′(s). Both these assumptions have minimal effect on the analysis,

given that the bandwidth of the loop is much smaller than the 1GHz bandwidth of the analog circuits and the 2Gs/s clock rate of the digital backend.

Given these assumptions, the residual parameter error, $X_e(s)$, can be expressed as:

$$X_e(s) = \frac{X_i(s)}{1 + H'(s)} + \frac{N_i(s) + N_a(s)}{1 + H'(s)} + N_d(s) \cdot \frac{H'(s)}{1 + H'(s)} \tag{2.9}$$

For first-order systems, H′ (s) is typically set to an integrator with some unity-gain frequency, (i.e. $H'(s) = \omega_u/s$). Under this condition, $X_e(s)$ becomes:

$$X_e(s) = [X_i(s) + N_i(s) + N_a(s)] \cdot \frac{\frac{s}{\omega_u}}{1 + \frac{s}{\omega_u}} + N_d(s) \cdot \frac{1}{1 + \frac{s}{\omega_u}} \tag{2.10}$$

$$= [X_i(s) + N_i(s) + N_a(s)] \cdot A_{HPF}(s) + N_d(s) \cdot A_{LPF}(s)$$

where $A_{LPF}(s)$ is a first-order low-pass filter and $A_{HPF}(s)$ is a first-order high-pass filter, both with a 3dB frequency of ω_u. With this formulation of the feedback loop, one can observe a basic tradeoff in the design of the filter in the estimation block, H′(s). In order to have a fast acquisition time and good tracking bandwidth, we desire the loop bandwidth, ω_u to be large. However, to minimize the effect of the detector noise, the loop bandwidth needs to be small. Finding the optimal value for the loop bandwidth then becomes a system engineering exercise, balancing the acquisition and tracking requirements with the requirements on the estimation accuracy.

The performance of the detector within the estimation block can be characterized by examining its transfer function. The input to the estimation block is the residual estimation error, x_e, and the output is $\hat{x}_e$. Since the input to the estimation block typically is noisy or corrupted in some way and the modulation of the data signal is effectively

random, it is more helpful to look at the expected value and variance of the detector output. An example plot is shown in Figure 3.24.

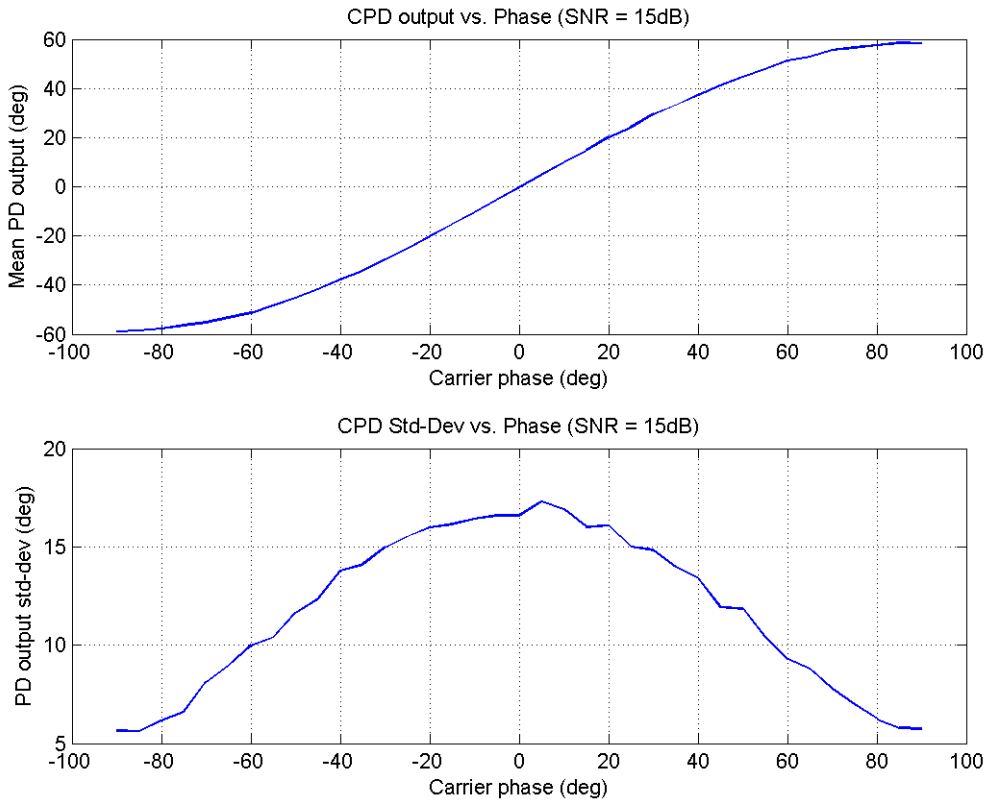

Figure 3.24: Sample detector output. (Top plot is expected value of the detector output. Bottom plot is the standard deviation of the detector output.)

A few desirable characteristics can be discerned from the plot above. First, it is crucially important that the expected value of the detector output be zero when the input residual error is zero. Otherwise, this offset will directly map to estimation error in the loop. However, the transfer function need not be linear and the error of the estimator output at large non-zero inputs is largely inconsequential. This is due to the fact that the estimator input is always driven towards the vicinity of zero during proper operation; the

synchronization error will only be large during initial acquisition, and during acquisition, coarse errors in the detector output slow the acquisition process only slightly. A constant slope to the detector output at small inputs is still desirable in order to keep the loop dynamics constant, however. Another desirable characteristic is for the detector to be low-noise; this is manifested by an output with a small standard deviation. The noise at the output of the detector can come from a variety of sources: noise in the input signal may cause noise in the output signal. Errors in other parameter estimate may degrade the quality of the estimate. Also, some detectors create noisy estimates themselves; in particular, NDA-based algorithms that have no knowledge of the underlying data bits transmitted typically require techniques that produce noisy estimates of the synchronization parameter, even in perfect SNR scenarios [57].

Additionally, there are some desirable characteristics for a detector that are not immediately obvious from an inspection of Figure 3.24. In particular, it is very important that the expected value of the detector output does not change significantly under different system operating conditions. For instance, Figure 3.25 shows the output of one particular symbol timing error detector with varying received signal SNR. As shown in the figure, as the input SNR reduces, the slope of the detector output also reduces; the varying gain of the detector would cause an SNR dependent variation in the synchronization loop dynamics, which could cause either instability or slow performance under different conditions.

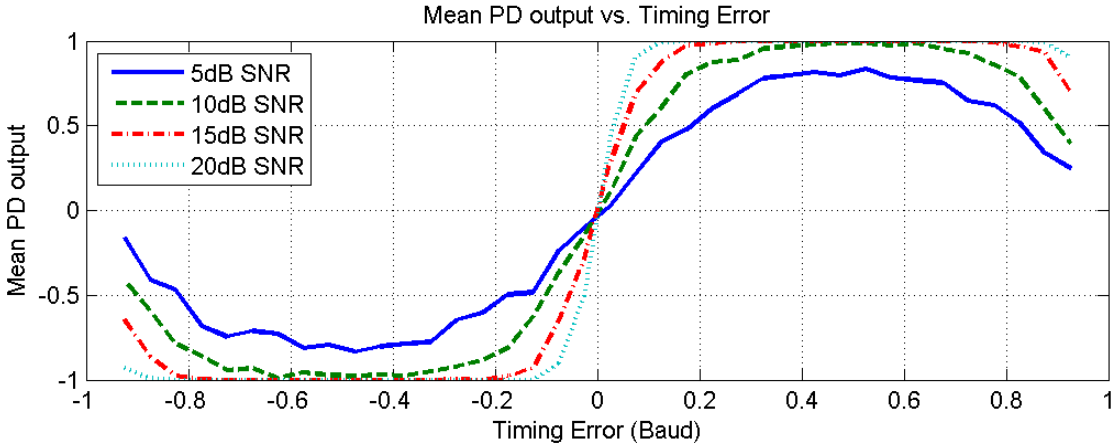

Figure 3.25: Symbol timing detector with undesirable SNR-dependent gain

Even worse than the above example is a case where the detector for a particular synchronization parameter is affected by the value of a different sync parameter. For instance, Figure 3.26(a) shows the output of a symbol timing detector in the presence of varying amount of carrier phase offset. As can be seen, the gain of this symbol timing detector is dependent on the carrier phase offset; the carrier phase-dependent gain of the detector would cause the dynamics of the symbol timing loop to be dependent on the carrier phase as well, which would likely cause instability. Figure 3.26(b) shows the output of another symbol timing detector in the presence of carrier phase offset. In this example, the detector output has an offset that is dependent on the carrier phase. In either case, the carrier phase error must be kept small so that it doesn't interfere with the proper operation of the symbol timing loop. Furthermore, if either of these detectors were used in conjunction with a carrier phase detector that had a dependence on symbol timing, the two loops would couple into one another and instability would likely result. Therefore, in

becomes incredibly important for the detectors to be designed in such a way that coupling cannot between two or more loops cannot occur.

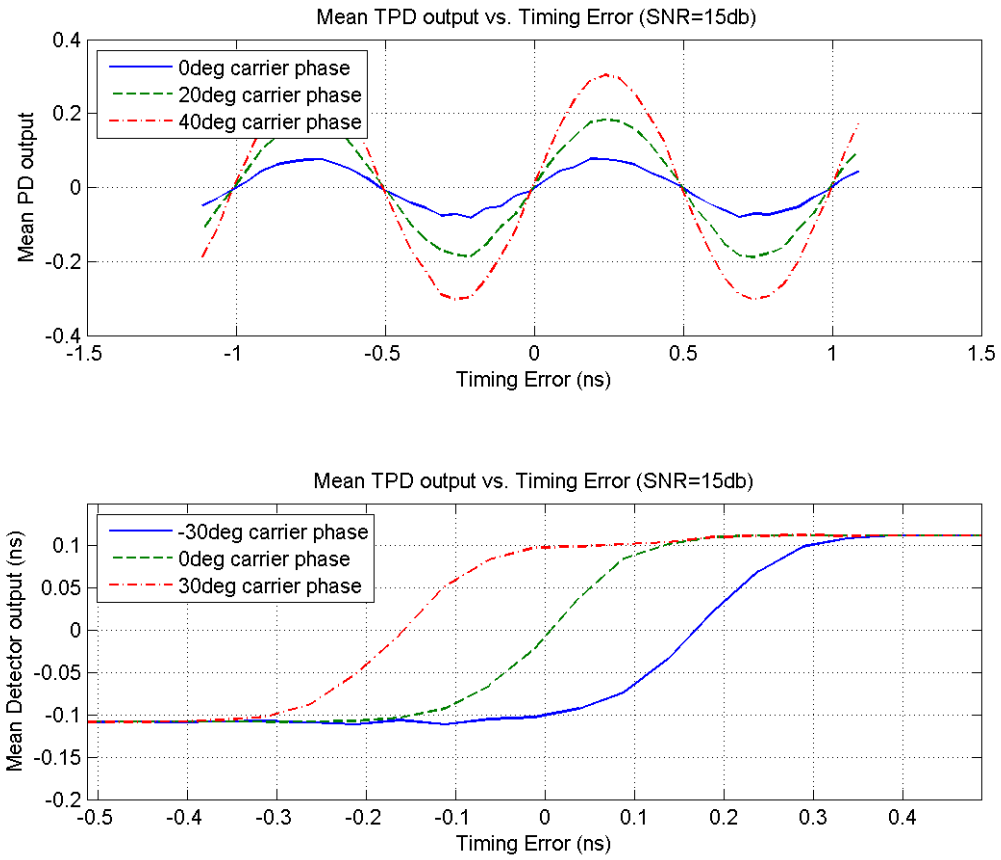

Figure 3.26: Two symbol timing detectors that would cause instability

3.4.2 Carrier Phase Estimation

The carrier phase synchronization loop uses a decision-directed detector that uses the sliced decision bits in its detection loop. As a result, this loop is loosely dependent on accurate symbol timing, since accurate symbol timing is a prerequisite of low-BER bit decisions. The exact algorithm used by the detector is:

$$\hat{\theta}_e = MF_I \cdot b_Q - MF_Q \cdot b_I \tag{2.11}$$

where $MF_{I/Q}$ is the matched output on the I/Q channel and $b_{I/Q}$ is the sliced bit decision on the I/Q channel. Figure 3.27 shows the expected transfer function of the detector output for various symbol timing offsets. Provided that the symbol timing offset is sufficiently small to prevent large numbers of decision errors, the expected behavior of the detector does not vary significantly with symbol timing offset. The figure also shows the standard deviation of the detector output; it can be observed that the standard deviation of the detector output increases significantly with increased symbol timing error. With perfect symbol timing, the carrier phase error out of the detector is 17° RMS, while at 300ps symbol timing error, the carrier phase error increases to 24° RMS.

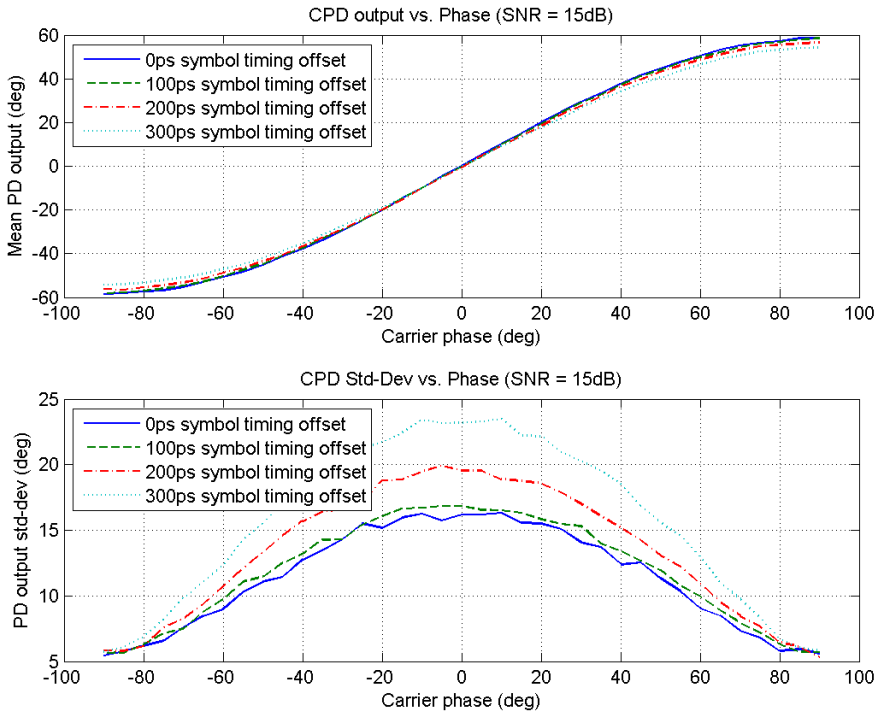

Figure 3.27: Decision-directed carrier phase detector output. Expected value of output (top) and standard deviation of output (bottom)

As mentioned previously, this detector uses a decision-directed approach for estimating the carrier phase error. Since proper decisions require reasonable symbol timing, the behavior of this detector during acquisition can be problematic. In order to circumvent this issue, a known training sequence is used to allow the loop to properly converge before actual data is transmitted. It is possible to use a non-data-aided detector in this loop; however, as is well-known, NDA loops typically are much noisier than their DA or DD counterparts [57]. For instance, the following algorithm is an NDA carrier phase detector:

$$\hat{\theta}_e(nT_s) = -Im\left[d^2(nT_s) \cdot d^2\left((n-1)T_s\right)\right]$$
$$d(nT_s) = d_I(nT_s) + j \cdot d_Q(nT_s) \tag{2.12}$$

where T_s is the symbol period (=1ns) and $d_{I/Q}$ are the I-channel and Q-channel ADC outputs. Note that this detector is insensitive to symbol timing, as the mathematics of (2.12) do not take the symbol boundaries into account in any manner. The expected value and standard deviation of the phase detector output are shown in Figure 3.28. The plot of the expected value of the output confirms the fact that symbol timing error does not affect the operation of the detector. However, the standard deviation of the detector output is about 6dB worse than the detector described in (2.11), as NDA detectors suffer a noise performance penalty.

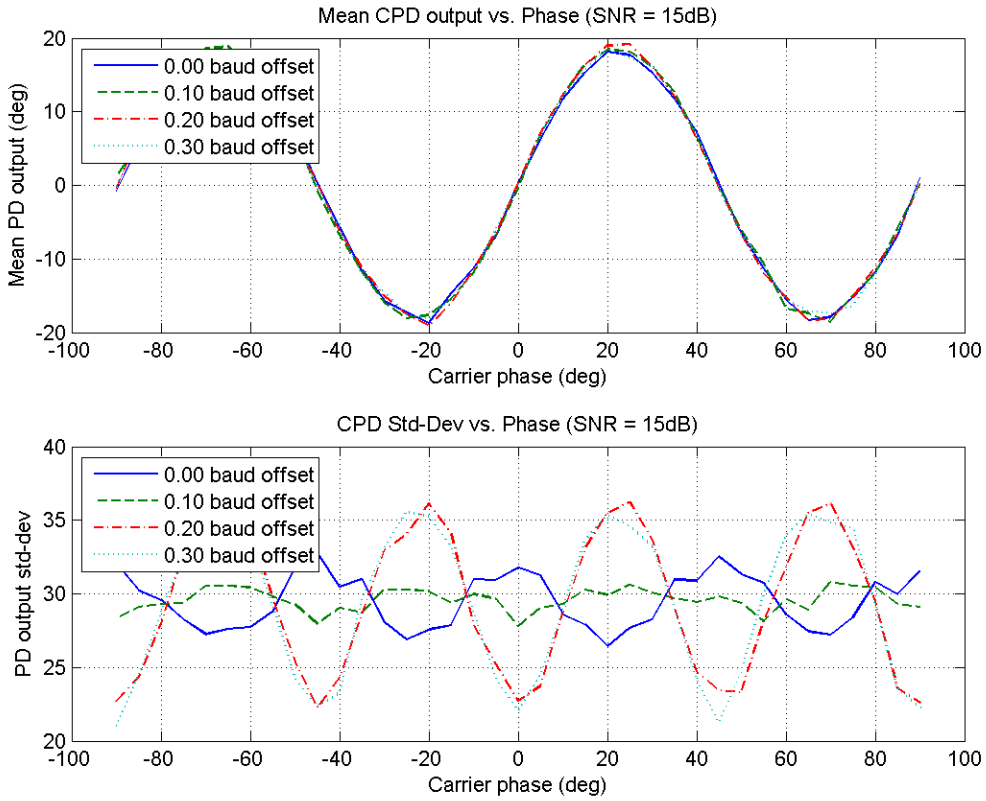

Figure 3.28: Non-data-aided carrier phase detector output. Expected value of output (top) and standard deviation of output (bottom)

The carrier phase synchronization must be able to accommodate a carrier frequency offset in addition to the carrier phase offset. A carrier frequency offset can be viewed as a phase ramp. A phase ramp can be expressed as:

$$\theta_{ramp}(t) = f_{offset} \cdot t \cdot u(t)$$
$$\Theta_{ramp}(s) = \frac{f_{offset}}{s^2} \tag{2.13}$$

where f_{offset} is the carrier frequency offset and u(t) is the unit step function. A phase step is:

$$\theta_{step}(t) = \theta_{in} \cdot u(t)$$
$$\Theta_{ramp}(s) = \frac{\theta_{in}}{s} \tag{2.14}$$

If one substitutes $\Theta_{ramp}(s)$ for $X_i(s)$ in equation (2.9), it becomes apparent that a second-order filter is required to obtain synchronization without a steady-state phase error. This is similar to the second-order filter required in PLL's to acquire frequency lock. In order to maintain reasonable phase margin at the unity-gain frequency of the loop, the second-order loop filter must have a LHP zero in its transfer function. The loop filter can be expressed as its continuous-time equivalent, H′(s):

$$H'(s) = \frac{\omega_u}{s^2}(s + \omega_z) \tag{2.15}$$

where ω_u is the unity-gain frequency of the loop and ω_z is the LHP zero. Solving for the closed loop transfer function, A(s), we obtain:

$$F(s) = \frac{1}{1 + H'(s)} = \frac{s^2}{s^2 + \omega_u s + \omega_z \omega_u} \tag{2.16}$$

which is a second-order high-pass filter, as expected. Therefore, the response of this filter to the phase step in (2.14)and the phase ramp in (2.13) is:

$$\Theta_{step,out}(s) = F(s) \cdot \Theta_{step}(s) = \frac{s \cdot \theta_{in}}{s^2 + \omega_u s + \omega_z \omega_u}$$
$$\Theta_{ramp,out}(s) = F(s) \cdot \Theta_{ramp}(s) = \frac{f_{offset}}{s^2 + \omega_u s + \omega_z \omega_u} \tag{2.17}$$

The design of the filter therefore simplifies down to the task of choosing the appropriate ω_u and ω_z. The loop unity-gain frequency, ω_u, is chosen by trading off the overall settling speed of the loop against the noise-filtering of the loop. As discussed above for the single-order case, increasing ω_u speeds up loop acquisition and tracking at

the expense of an increased noise bandwidth. The zero frequency, ω_z, presents another degree of freedom in modifying the dynamics of the loop. The step response and the ramp response for several filters with different ω_z is shown in Figure 3.29. In this figure, the unity gain frequency of the filter, ω_u, is normalized to one. As can be seen in the figure, increasing the zero frequency increases the Q of the filter response; when $\omega_z \ll \omega_u$, the filter response degenerates into a simple first-order filter. As a result, the step response does not have any overshoot, but the filter cannot accurately cancel out the error from a constant phase ramp. As the zero frequency is increased, the initial speed of the step response increases, but overshoot occurs, limiting the overall response speed.

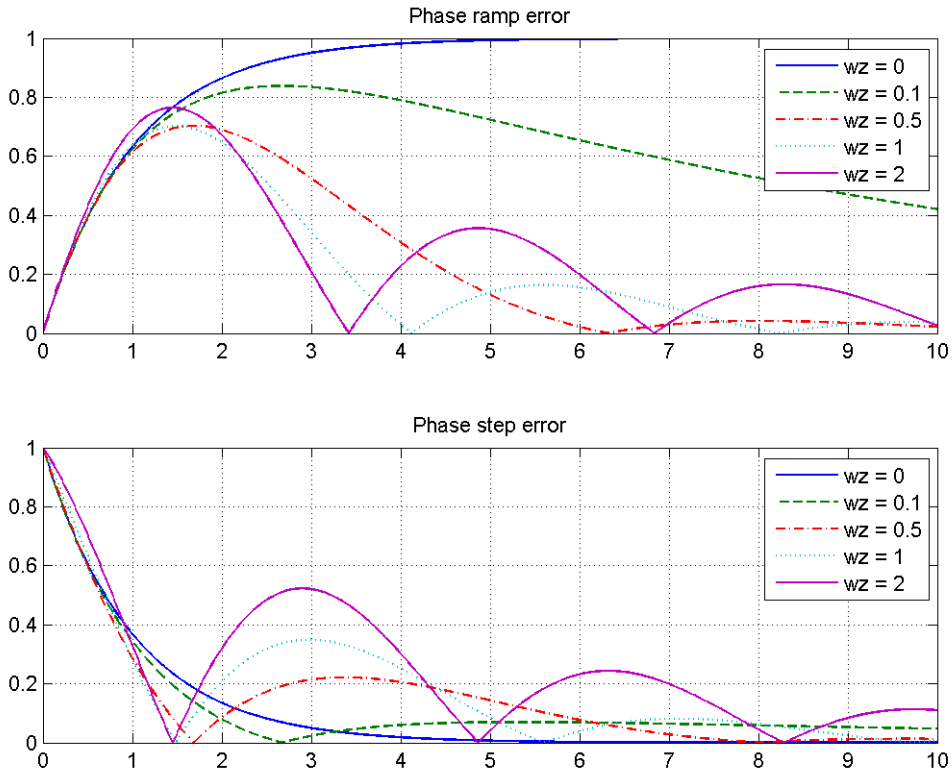

Figure 3.29: Ramp response and step response of second-order filter with fixed ω_u and varying ω_z

The carrier phase recovery loop uses a “gearshifting” technique for its filter. During initial carrier phase and frequency acquisition, the unity gain frequency of the loop is increased in order to speed acquisition; noise performance of the loop is suboptimal during the acquisition phase, however. Once the carrier is acquired, the unity gain frequency of the loop is gearshifted down to a lower value in order to limit the noise in the loop. The gearshifting ratio used in this design is a factor of 5. During acquisition the unity-gain frequency is set at roughly 8MHz (50Mrad/s) in order to pull in a 6MHz carrier frequency offset that can arise from 100ppm crystals. During tracking mode, the unity-gain frequency is reduced to roughly 1.6MHz (10Mrad/s) in order to bandlimit the noise from the phase detector. A loop bandwidth less than the carrier frequency offset is allowable because the carrier frequency offset is static and has very little frequency drift due to the high frequency stability of the source crystal. The zero frequency is set to be half the unity gain frequency, as it presents a reasonable compromise between transient response speed and overshoot. Both the unity gain frequency and the zero frequency can be easily modified by changing the coefficients in the digital loop filter.

3.4.3 Symbol Timing Estimation

The timing synchronization loop uses a non-data-aided (NDA) detector to ensure timing acquisition regardless of accurate bit decisions or other system parameters like carrier phase/frequency offset. By using an NDA detector for symbol timing, coupling between the symbol timing loop and other synchronization loops is prevented. The estimation block uses a fourth-power detector to make the estimate independent of carrier phase or data bit:

$$\hat{\tau}_e(nT_s) = \left[\left(d_I^2(nT_s) - d_Q^2(nT_s)\right) \cdot \left(d_Q^2\left((n-1)T_s\right) - d_I^2\left((n-1)T_s\right)\right)\right] \cdot (-1)^n \quad (2.18)$$

where T_s is the symbol period (=1ns) and $d_{I/Q}$ are the I-channel and Q-channel ADC outputs. The expected value and standard deviation of the timing detector output are shown in Figure 3.30. The plot of the expected value of the output confirms the fact that carrier phase error does not affect the operation of the detector. As is shown in figure, the noise performance of this detector rather poor; the RMS error of the symbol timing detector is as big as 360 ps. However, there are two factors that enable the use of such a noisy detector within the system. First, as with all estimation loops, the loop filter will significantly bandlimit the noise produced by the detector. Therefore, as discussed above, the noise performance of the symbol timing loop can be traded off against the acquisition and tracking speed of the loop. Secondly, as was shown in Figure 3.12 (on page 70), the overall receiver is robust to symbol timing errors as large as 100ps (0.1baud), as they cause less than 0.5dB sensitivity degradation. Therefore, a noisy NDA timing detector can be used in order to minimize the possibility of any loop-coupling issues described in section 3.4.1.

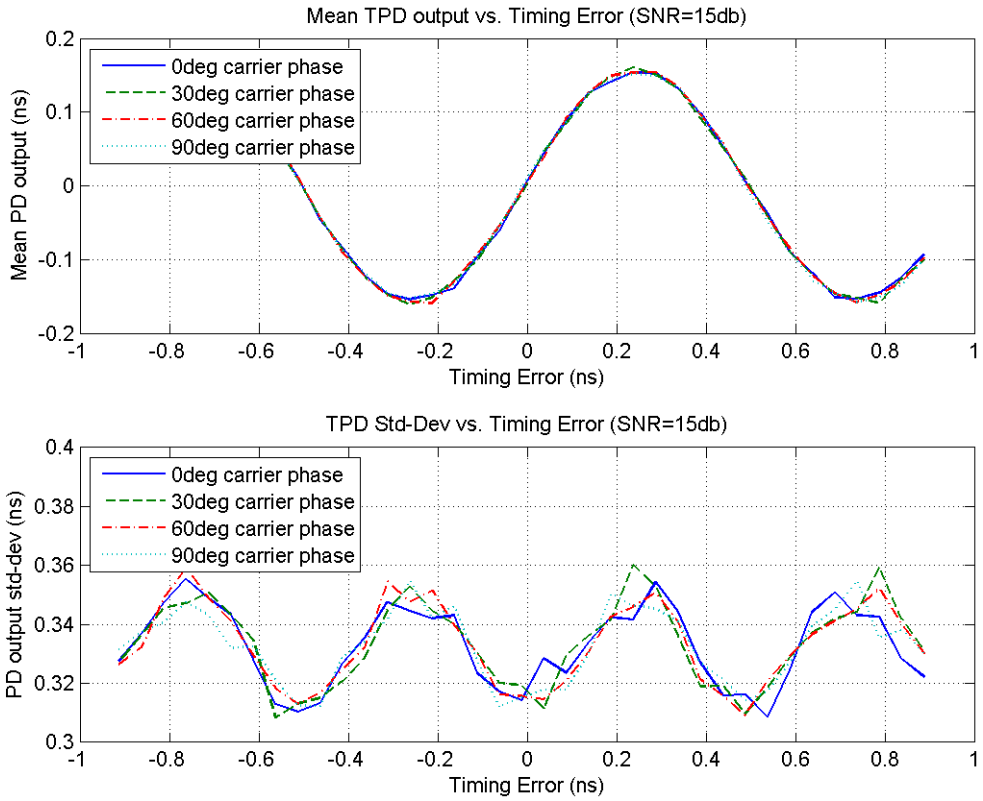

Figure 3.30: Non-data-aided symbol timing detector output. Expected value of output (top) and standard deviation of output (bottom)

The filter used in the timing synchronization loop is a simple first-order integrator, as it is capable of driving the steady-state symbol timing error to zero while guaranteeing loop stability. The bandwidth of the filter is approximately 1.5MHz, and the update rate of the detector is 1GHz, so the filter reduces the noise from the detector by roughly 25dB ($=10\log_{10}[1\text{GHz}/(2\times1.5\text{MHz})]$, reducing the noise from the detector to roughly 16ps RMS after filtering.

3.4.4 Channel Estimation

The channel estimation loop is inherently built into the adaptation loop of the decision feedback equalizer. A simplified diagram of a multipath channel and a DFE is shown in Figure 3.31. The multipath channel is represented by an FIR filter with coefficients, a_m. The steady-state DFE tap coefficients, c_m, are the direct estimation of the multipath coefficients; in other words, $c_m = \hat{a}_m$, where $\hat{a}_m$ is the estimate of a_m . The adaption of each DFE tap coefficient is governed by the well-known LMS algorithm:

$$c_m\big((n+1)T_s\big) = c_m(nTs) + \mu \cdot \Delta_m(nTs)$$
$$\Delta_m = \big(b_0 - \hat{b}_0\big) \cdot b_m \tag{2.19}$$

Where b_0 is the current slicer decision, b_m is the slicer decision from the m periods ago, $\hat{b}_0$ is the current slicer input, and μ is the LMS adaptation coefficient.

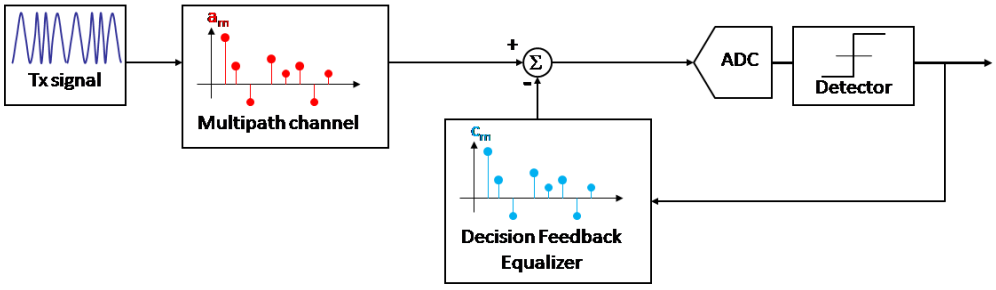

Figure 3.31: Simplified diagram of multipath channel and DFE

The DFE used in this receiver is a complex DFE, with a quadrature input, quadrature output, and complex-valued DFE taps. Taking this into account, (2.19) can be expressed as follows:

$$\begin{aligned} c_m\big((n+1)T_s\big) &= c_m(nTs) + \mu\big[\Delta_{I,m}(nTs) + j\cdot\Delta_{Q,m}(nTs)\big] \\ \Delta_{I,m} &= \big(b_{I,0} - \hat{b}_{I,0}\big)\cdot b_{I,m} + \big(b_{Q,0} - \hat{b}_{Q,0}\big)\cdot b_{Q,m} \\ \Delta_{\mathrm{Q},m} &= \big(b_{Q,0} - \hat{b}_{Q,0}\big)\cdot b_{I,m} - \big(b_{I,0} - \hat{b}_{I,0}\big)\cdot b_{Q,m} \end{aligned} \tag{2.20}$$

For a multipath channel with 15ns RMS delay spread, the multipath is expected to die off in about 3σ, or roughly 45ns. Since the symbol time of the MSK system implemented is 1ns, a full DFE would require roughly 45 taps to fully equalize the channel. While the DFE estimation engine is implemented in the digital domain, the DFE correction subcircuit (i.e. the equalizer circuit itself) is implemented in the analog domain. Due to the parasitic capacitance that is associated with each DFE tap, the analog portion of the DFE can only have a limited number of taps, roughly 16. Therefore, in order to make optimal use of the limited number of analog equalizer taps available, the DFE estimation engine must not only estimate the magnitude of the DFE taps, but it must also determine where along the DFE delay line the 16 taps should be positioned. Another way to envision this problem is to assume that you have a 45-tap DFE subject to the constraint that only 16 of the taps can have a non-zero magnitude at any given time. When posited in this form, the obvious solution is to adaptively allocate the DFE taps to the positions in time that correspond to the largest absolute values of the channel impulse response. This process is dubbed *adaptive tap allocation* in this work.

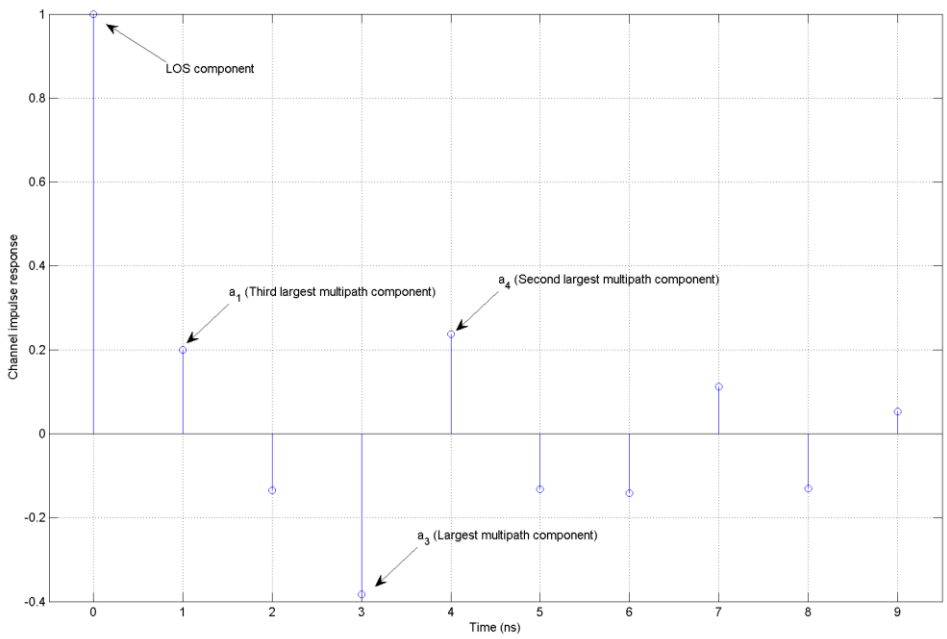

Figure 3.32: Sample channel impulse response. Adaptive tap allocation algorithm would identify 3ns delay, 4ns delay and 1ns delay (corresponding to a_3, a_4, and a_1) as the optimal tap placement

For instance Figure 3.32, shows an example channel impulse response with a length equal to 9 symbol periods. The first tap of the channel impulse response corresponds to the line-of-sight transmission path and subsequent taps are caused by unwanted multipath reflections. Let us further assume in this example that the analog equalizer circuit was implemented with only 3 taps. The role of the adaptive tap allocation process is to determine where to position the 3 available taps; or in the parlance of (2.19), which 3 values of $c_1, c_2, \dots, c_9$ should be set to nonzero values. For the example of Figure 3.32, the optimal values of the constrained c_n are just the 3 largest taps of $a_1, a_2, \dots, a_9$, as shown in the figure. If the channel impulse response were to change

to Figure 3.33, then the adaptive tap allocation process would change to select the new largest 3 taps, as shown in the figure.

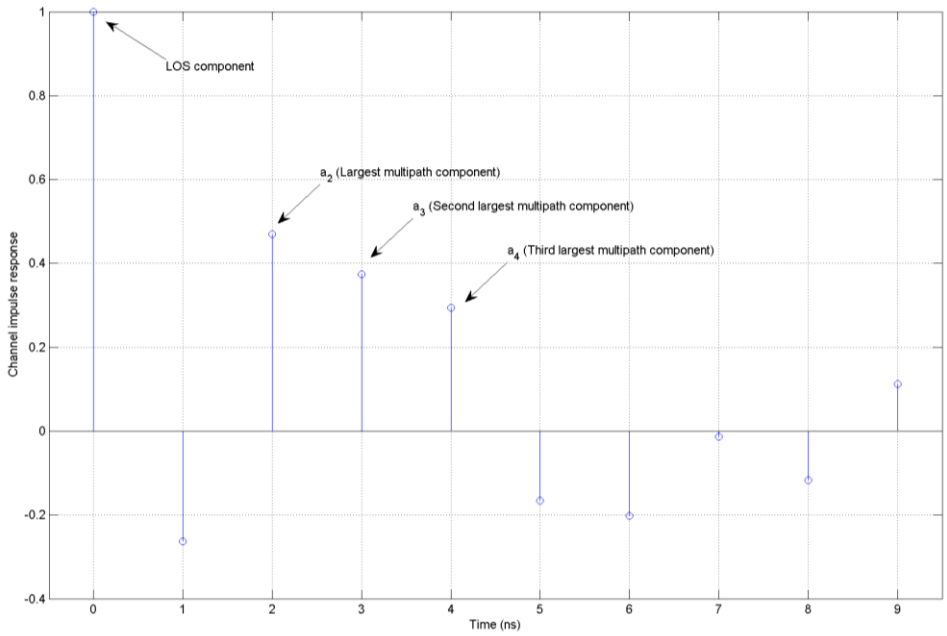

Figure 3.33: Sample channel impulse response. Adaptive tap allocation algorithm would identify 2ns delay, 3ns delay and 4ns delay (corresponding to a_2, a_3, and a_4) as the optimal tap placement

In the case of the MSK receiver, in order to determine the 16 largest values of the multipath channel impulse response, every tap value in the 45ns multipath profile must be estimated. However, the taps that are not "activated" via the adaptive tap allocation algorithm cannot use the tap magnitude estimation algorithm described in (2.19), because $c_{m,inactive}$ is forced to be zero thus breaking magnitude adaptation loop. However, even with a DFE tap forced to be zero, an estimate of the channel impulse response can still be derived via:

$$\hat{a}_{m,inactive}\,(nT_s) = \frac{1}{N}\sum_{k=1}^{N}\Delta_m\big((n-k)T_s\big) = \frac{1}{N}\cdot\frac{1-z^{-N}}{1-z^{-1}}\cdot\Delta_m(nT_s)$$
$$\hat{a}_{m,active}\,(nTs) = c_m(nTs) = \mu\frac{z^{-1}}{1-z^{-1}}\Delta_m(nT_s) \tag{2.21}$$

where $\Delta_m(nTs)$ is defined in (2.19) and *N* is the number of samples over which to take the mean. Once all the values of the channel impulse response are estimated, a simple sorting operation (by absolute value) can identify the largest magnitude channel taps to equalize.

It turns out that the hardware necessary to implement the adaptive tap magnitude calculation as shown in (2.19) can be easily repurposed to estimate the channel impulse response of an inactive tap, as shown in (2.21). If the values of $\hat{a}_{m,inactive}$ are only updated every *N* symbols, then only a single multiply-accumulate (MAC) block is needed for each tap location. When a particular tap is selected by the adaptive tap allocation algorithm (i.e. the tap is "active"), then the MAC is configured like an integrator in order to implement the LMS adaptation routine in (2.19). When that tap is not selected (i.e. the tap is "inactive"), then the averaging operation of (2.21) can be achieved by simply resetting the same integrator every *N* cycles. Therefore, minimal additional hardware is required to implement the adaptive tap allocation algorithm, and it essentially comes for "free" with the standard DFE tap magnitude allocation block.

3.5 System Simulations

In order to test and verify the performance of the mixed-signal baseband receiver described here, a full simulation model of the 60GHz wireless link was built in MATLAB Simulink [58]. In order to speed up simulation time, a baseband-equivalent model of the mm-wave components was used [59], so the simulation timescale could be referenced to the lower frequency symbol period instead of the 60GHz carrier. A stochastic channel model was used to model the multipath behavior of the indoor channel; RMS delay spread for the channel was 15ns and the Rician K was 5dB. Random carrier phase offset was introduced and a carrier frequency mismatch of 6MHz (corresponding to 100ppm at 60GHz) between TX and RX was added. The nonideal mm-wave circuit models discussed in section 3.1 and detailed in Table 3.1 (page 54) were used to simulate PA nonlinearity and LO phase noise. The timing synchronization, carrier synchronization, and equalization loops as detailed in the section 3.4 were also included, so that a complete model of the baseband receiver could be simulated.

Figure 3.34 shows the BER vs. SNR waterfall curve for the baseband receiver. Each simulation data point is the ensemble average of 20 Monte Carlo runs of 10000 MSK symbols each. Each Monte Carlo run has a randomly generated, time-invariant channel impulse response, a randomly generated symbol timing offset, and a randomly generated carrier phase offset, and a 6MHz (100ppm on 60GHz) carrier frequency offset. An SNR of 10dB is required to achieve the 10^{-3} BER target, representing less than a 1dB implementation loss from the ideal MSK receiver with equalization shown in Figure 3.10.

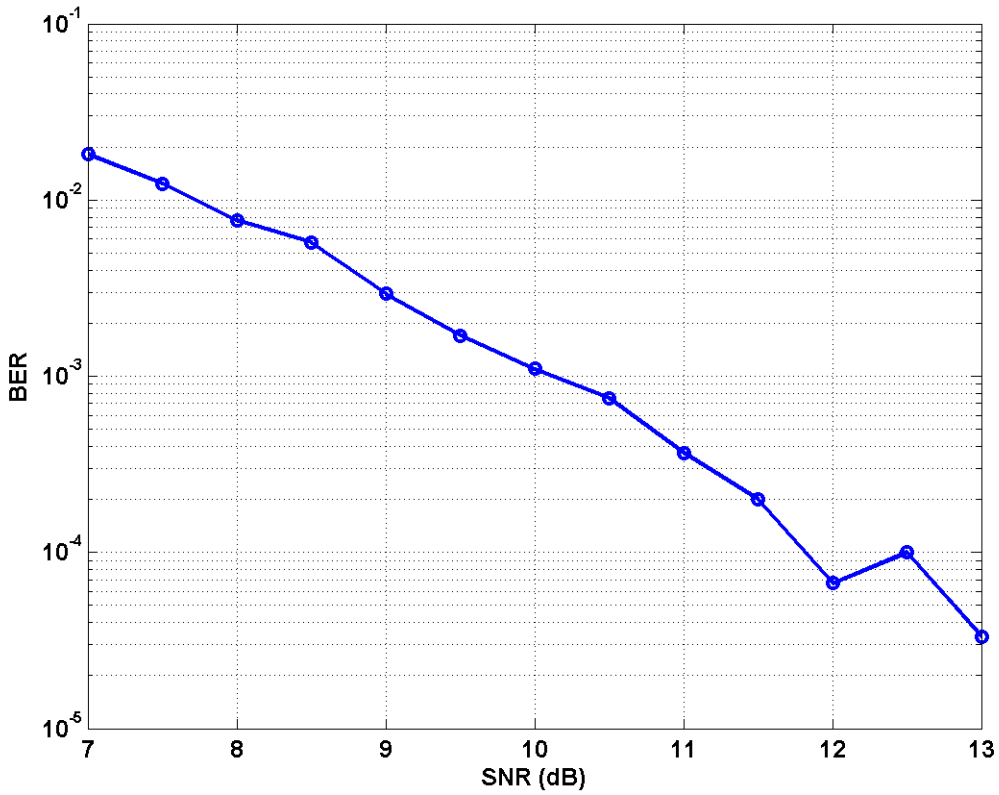

Figure 3.34: BER vs. SNR curve of full Simulink receiver simulation

Figure 3.35(a) shows the acquisition of carrier phase and frequency performed by the carrier phase recovery block. The solid line shows the actual carrier phase and the dotted line shows the carrier phase estimate provided by the CPR digital estimator and used by the analog CPR block to correct for phase offset. The input phase has is a steady ramp, due to the 6MHz frequency offset between transmitter and receiver. Figure 3.35(b) shows the phase error between the estimate and the actual carrier phase. The CPR block quickly acquires carrier frequency and phase in less than 100ns, and the subsequent phase tracking error from 100ns to 1000ns has a standard deviation of 6.7°. As discussed in section 3.4.2, the carrier phase recovery block uses a gear-shifting approach, where the

bandwidth of the recovery loop is higher during acquisition and lower during tracking. As can be observed in Figure 3.35(b), the gear-shift occurs at T=1000ns. From this point on, the carrier phase recovery block track the phase error with a narrower bandwidth, resulting in less noise in the carrier phase estimate. After gear-shifting, the RMS phase error is only 2.1°. Most of this residual phase error is due to the limited update rate of the CPR block; since the CPR block only updates at 2ns intervals, the 6MHz carrier offset with result in a 4° ripple on the phase error which is clearly discernable in the figure.

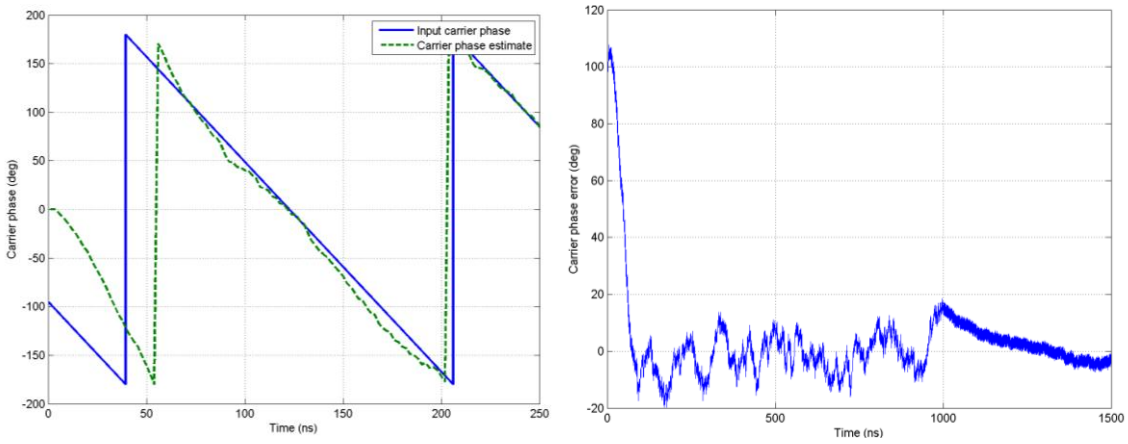

Figure 3.35: (a) Acquisition of carrier phase and frequency and (b) Residual carrier phase error

Figure 3.36 shows the acquisition of symbol timing. In comparison to carrier acquisition, symbol timing acquisition occurs much more slowly, taking about 1200ns to acquire. This is due to several factors: first, symbol timing is acquired using a non-data-aided (NDA) detector, as opposed to the decision-directed (DD) detector used in carrier phase tracking. In order to keep the high noise levels of the NDA detector in check, the symbol timing loop has a significantly lower bandwidth that the carrier phase loop, resulting in slower timing acquisition. Secondly, the symbol timing detector is very

sensitive to errors in the equalization loop. Comparing the initial acquisition of timing (Figure 3.36) to the initial acquisition of phase (Figure 3.35(b)), one can observe that the timing acquisition makes significantly more incorrect decisions en route to acquisition, limiting the acquisition speed. Lastly, the acquisition of carrier phase was sped up by use of the gearshifting technique discussed above; the timing acquisition loop did not have a gearshifting loop included. Nonetheless, once symbol timing is acquired, the residual timing error is 15ps, RMS. As discussed in section 3.3.1, the MSK receiver is quite robust to symbol timing errors; therefore, this level is sufficiently low so as to not affect overall receiver sensitivity.

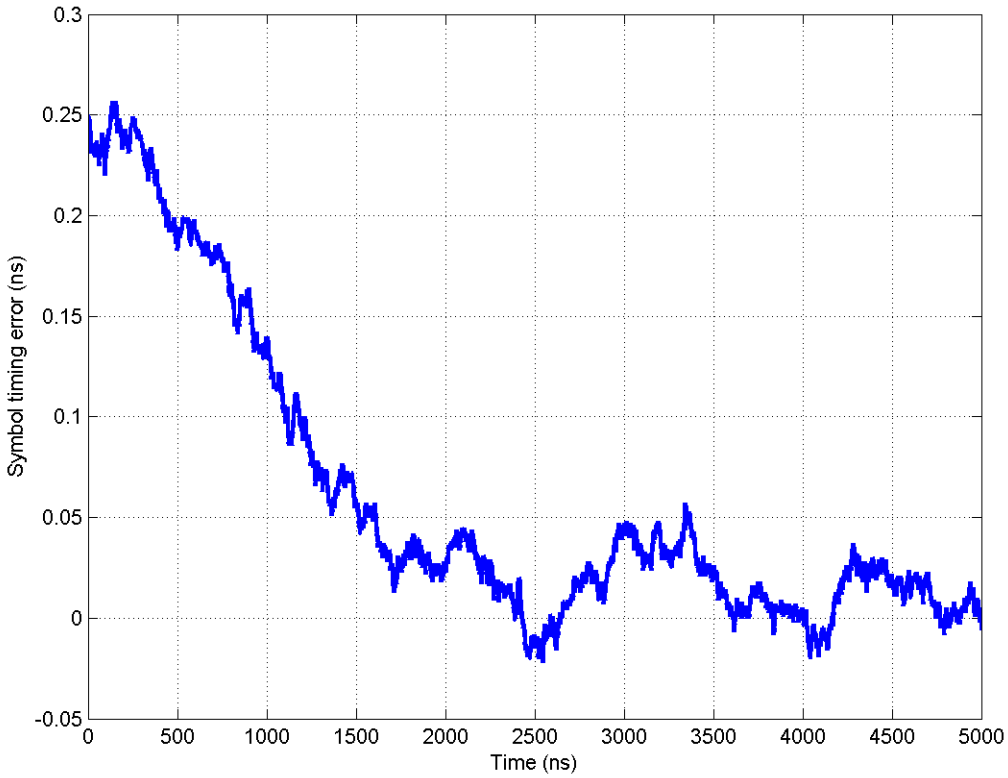

Figure 3.36: Symbol timing acquisition

Figure 3.37(a) shows the real part of a sample channel impulse response used in one of the Monte Carlo simulations. The Rician K of this particular channel was 4.3dB and the RMS delay spread was 12.7ns. Figure 3.37(b) shows the transient response of the adaptive tap allocation algorithm. This is the algorithm discussed in section 3.4.4 that determines the optimal placement in time for the 8 DFE taps. The presence of a bar at any vertical location in this plot indicates that that one of the 8 DFE taps is being used at that selected delay value. (Note that there are only 8 "bars" active at any given instant in time.) For instance, the cyan bar at tap #11 indicates that one of the DFE taps was allocated to the delay location of 11ns, corresponding to the channel impulse response value of -0.16. In this simulation, at T=1000ns, the adaptive tap allocation routine determines that tap #11 is one of the 8 largest tap values. Due to the fact that the 8-tap DFE can only cancel out a limited number of taps, this adaptation process tends to be fairly noisy in the presence of tap values that are of similar magnitudes. For instance, there are several channel impulse response values that possess an absolute of roughly +0.05. As a result, the tap allocation scheme periodically swaps between a subset of these taps. This can be most visibly seen in the somewhat periodic selection and deselection of the 17ns delay tap (the intermittent red bar in Figure 3.37(b)). While the intermittent activation of similar magnitude taps does not degrade overall equalizer performance, the "noisy" selection of the tap allocation routine can be suppressed by including some small amount of hysteresis in the tap selection process.

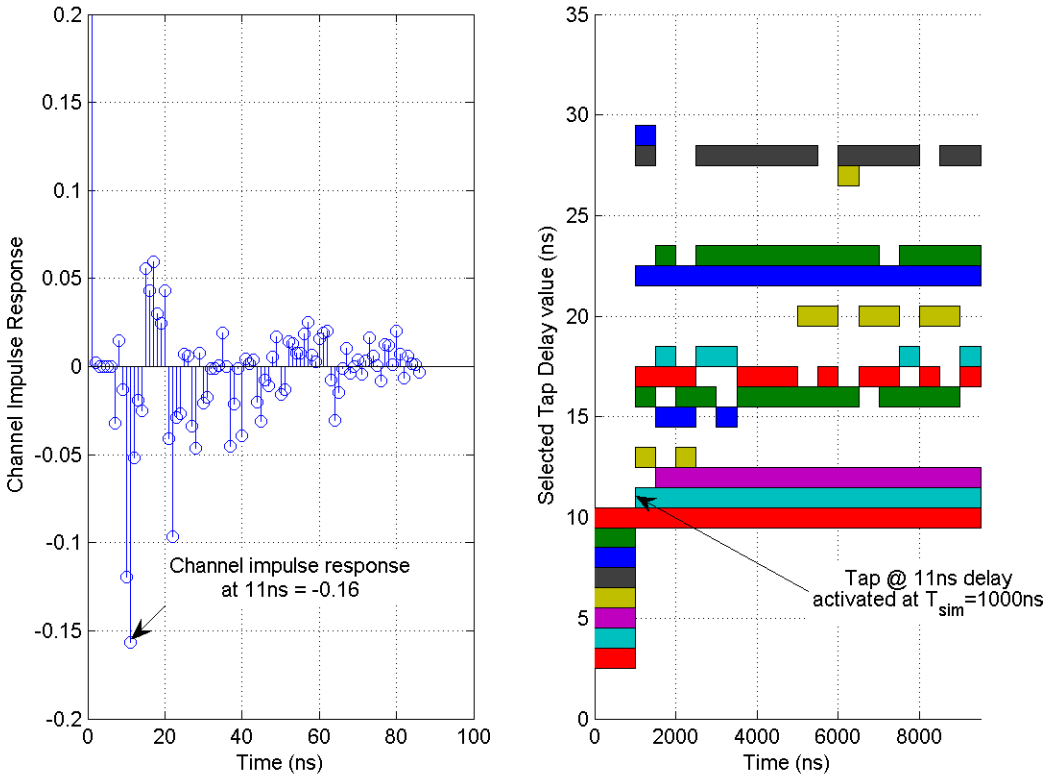

Figure 3.37: (a) Real portion of sample channel impulse response and (b) Tap activation from adaptive tap allocation routine

Figure 3.38 shows the tap magnitude adaptation of three taps within the DFE; the ideal adaptation values are also shown. No quantization in the actual tap value is shown in this plot, as the digital estimators are modeled with infinite precision. Quantization of these tap values is modeled in the analog block, in order to capture the finite precision of the analog DFE tap values.

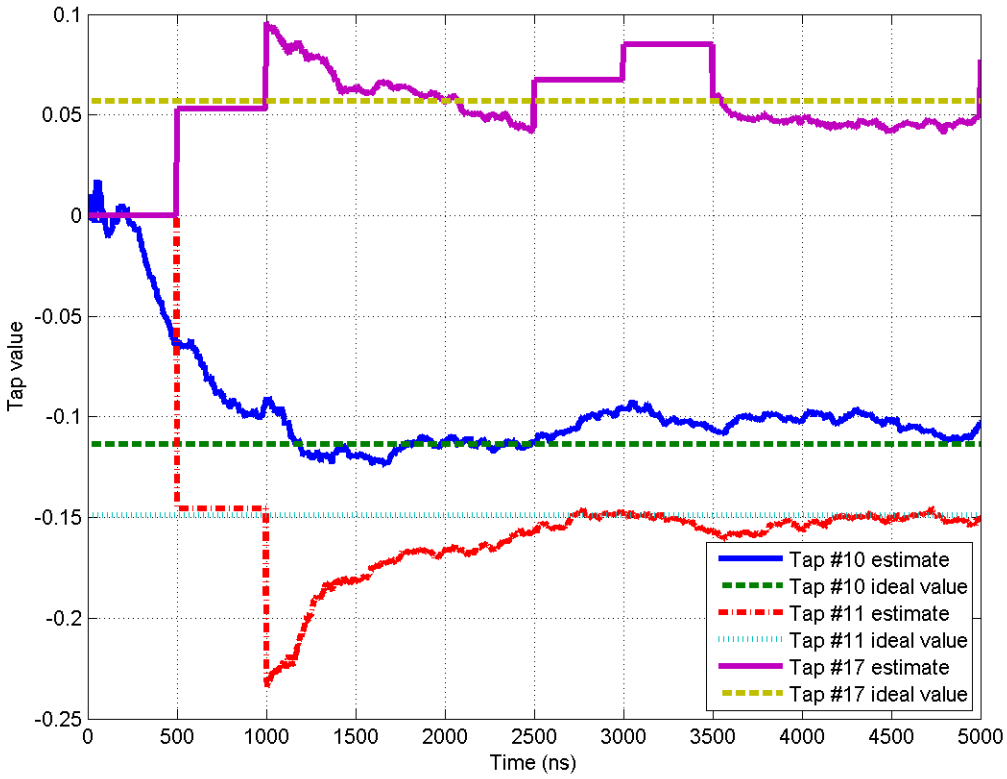

Figure 3.38: DFE tap magnitude adaption of a few taps

For visualization purposes, the plot also includes the running estimate of the tap value even when the tap is not activated by the tap allocation process described above; the tap estimate value for a deselected tap position is updated every 500ns, so the regions in Figure 3.38where the tap estimate is static indicate a time period when that particular tap position is not activated. For instance, in Figure 3.37(b), delay position #17 was not in use from T_{sim} = 2500ns to 3500ns; therefore, in Figure 3.38, the tap estimate for delay position #17 is static during that timespan. As can be observed in Figure 3.38, the

adaptation of the equalizer takes about 2000ns, after which time all tap values are close to their ideal values.

3.6 Conclusions

This chapter introduced a system architecture for a mixed-signal baseband receiver that is capable of functioning within an indoor 60GHz wireless channel. The modulation scheme was selected in order to perform robustly in the presence of nonideal CMOS mm-wave components. The partitioning of the receiver architecture between the analog and digital domain was chosen in order to minimize the overall system complexity and power dissipation of the receiver while still enabling the use of robust digital synchronization and estimation algorithms. Algorithms for synchronization were proposed that enabled the use of mixed-signal feedback loops to detect and correct for various synchronization parameters such as symbol timing, carrier phase and frequency, and channel multipath impulse response. Finally a full system simulation incorporating both digital and analog circuit models was implemented in order to verify the performance of the receiver architecture. The overall implementation loss compared to an MSK receiver with ideal equalization and synchronization was shown to be less than 1dB. Specifications for the analog circuit blocks, such as the ADC and DFE, were determined; a discussion of the implementation details of these circuits is presented in the next chapter.

4 Circuit Implementation

The previous chapter detailed the system design of the baseband mixed-signal receiver for a 60GHz wireless communications link. A simplified diagram of the proposed architecture is shown again in Figure 4.1. In this chapter, the circuit implementation of the mixed-signal receiver is discussed. In section 4.1, a more detailed view of the architecture is shown, and in subsequent sections, each of the analog circuit blocks are described.

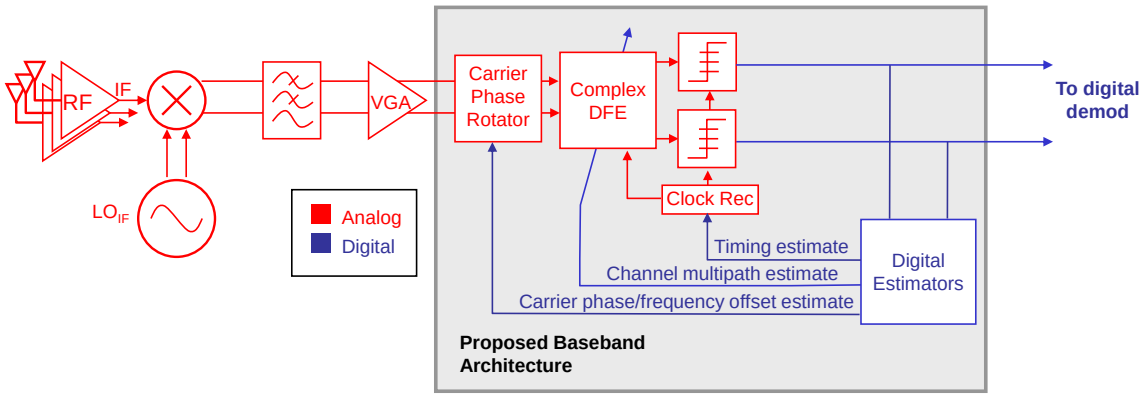

Figure 4.1: Simplified block diagram of proposed mixed-signal receiver architecture

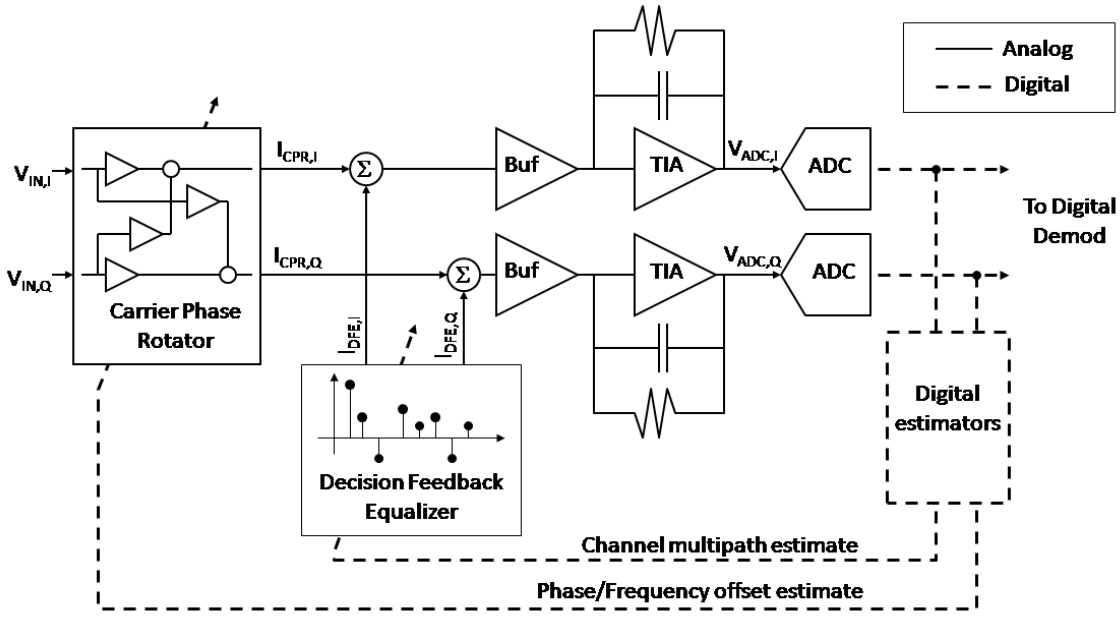

Figure 4.2: Circuit-level architecture of mixed-signal receiver

4.1 Block-level Architecture

Figure 4.2 shows a more detailed view of the baseband mixed-signal receiver circuit implementation. All analog circuits and signal paths are differential. The input signal ($V_{IN,I}$ and $V_{IN,Q}$) is a baseband quadrature voltage waveform . This signal passes through the carrier phase rotator block (CPR) which compensates for frequency and phase offset between the transmitter and receiver. The CPR output ($I_{CPR,I}$ and $I_{CPR,Q}$) is in the current domain. After the CPR, the decision-feedback equalizer (DFE) removes channel multipath. The DFE consists of 16 taps that can be allocated arbitrarily within a 32ns search range. Current-mode outputs for the CPR and DFE blocks are desirable, as they lend themselves easily to the summation operation that follows; because both the

CPR and DFE outputs are in the current-domain, the summation of these waveforms can be achieved by simply connecting the two signals together in parallel.

In order to quantize the received signal, it must be converted back into the voltage domain. The transimpedance amplifiers (TIAs) perform that function. The TIAs serves a secondary purpose as track-and-hold amplifiers. The track-and-hold functionality is required to provide precise sampling instants for the 4-bit, 2Gs/s ADCs.

The summation nodes where the currents from the CPR and DFE blocks are combined has large parasitic capacitances due to the significant number of devices connected to these nodes and the parasitic wiring capacitance required to connect all these blocks together. The input to the TIA is a summing junction node, and as such, is sensitive to excess capacitance. In order to decouple the CPR and DFE output capacitance from the TIA summing junction, a current-mode buffer (BUF) is inserted between these nodes.

4.2 Carrier Phase Rotator

As described in chapter 3, the role of the carrier phase rotator (CPR) is to correct for carrier phase and frequency error between the transmit and receive local oscillators. Given a quadrature input signal, $S_{I,in}$ and $S_{Q,in}$, and a desired phase rotation, θ, the desired output of the CPR block is:

$$\begin{aligned} S_{I,out} &= \cos\theta \cdot S_{I,in} - \sin\theta \cdot S_{Q,in} \\ S_{Q,out} &= \sin\theta \cdot S_{I,in} + \cos\theta \cdot S_{Q,in} \end{aligned} \tag{3.1}$$

The expressions in (3.1) can be expressed in a signal flow diagram, as shown in Figure 4.3, where the gain of each amplifier is the sin or cos of the rotation angle.

Therefore, it becomes clear that the carrier phase rotation operation is simply a set of gain and summation operations, as shown in the figure. The gain of each amplifier must be programmable, as it is dependent on θ; however, the gain range of each amplifier is from -1 to 1 on a linear scale and is inclusive of a gain of zero. Therefore, traditional CMOS VGA's are not applicable to this application, as they have a gain profile that is typically "linear in dB" with respect to the control voltage. Furthermore, the summation at the outputs of the VGA's favor a circuit topology that has a current-domain output.

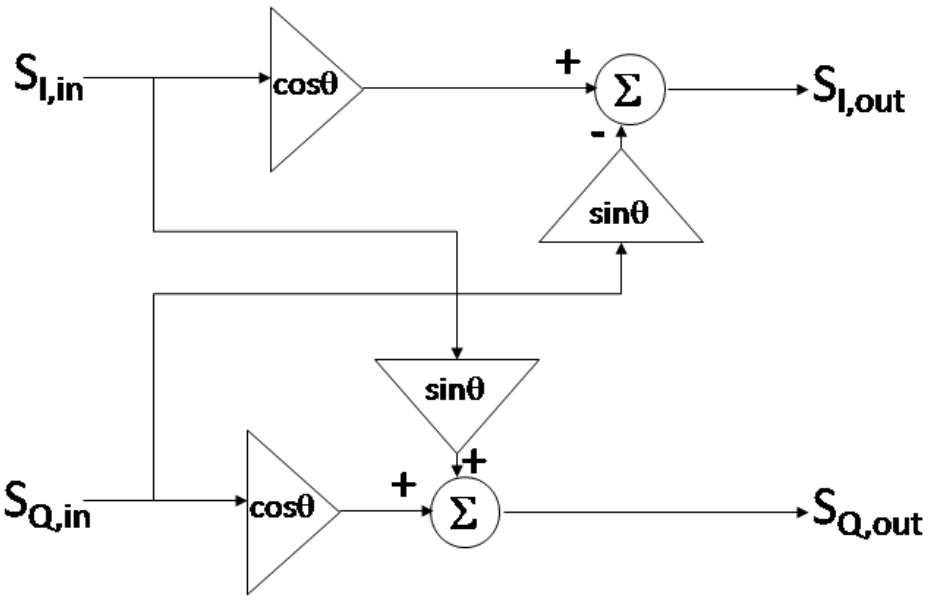

Figure 4.3: Signal-flow diagram of carrier phase rotator block

4.2.1 Gilbert Quad VGA

A doubly-balanced CMOS Gilbert quad is a good topology to implement the each of the four VGA's shown in Figure 4.3. Gilbert quads (see Figure 4.4) are typically used in RF switching mixers [60] or BJT translinear circuits [61], [62], but their characteristics

match the needs of the CPR quite well. The gain of the quad is a function of the control voltage, V_C, and the quad is capable of achieving positive, negative, and zero current gain. Also, with the addition of a differential transconductor at the input to the quad, the circuit has a voltage-mode input and a current-mode output, as required by the VGA's in the CPR block.

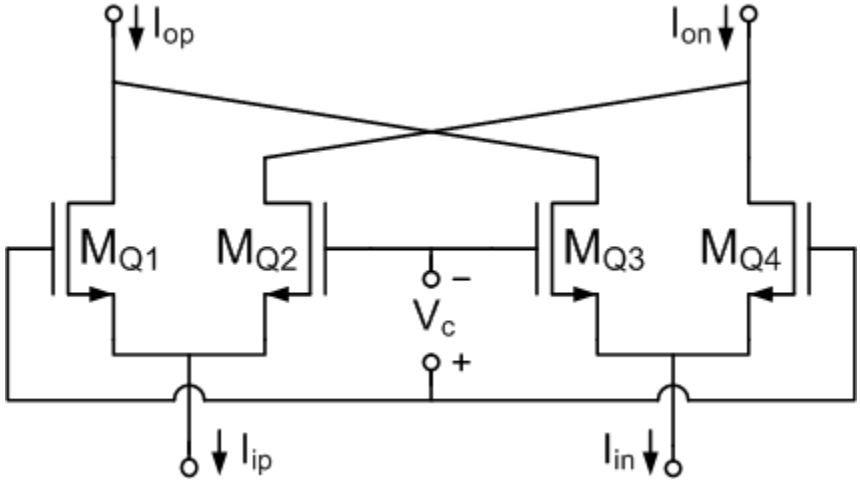

Figure 4.4: CMOS Gilbert Quad circuit

A full circuit diagram of the VGA, including both the input transconductor and the Gilbert quad, is shown in Figure 4.5. The transconductor topology [63] uses a local feedback loop of M_1, M_{2a} and M_3 to create an input device with a high effective g_m; therefore, the effective G_m of the input transconductor circuit is set by the resistor, R, that provides degeneration between the two input devices and by the mirroring ratio between M_{2a} and M_{2b}. This not only provides a fixed transconductance that is largely insensitive to process, but it also linearizes the transconductance of the circuit, as a passive resistor has significantly higher linearity than an active device. The input PMOS

devices, M_1, have their wells tied to their sources in order to prevent V_{TH} modulation from degrading circuit linearity.

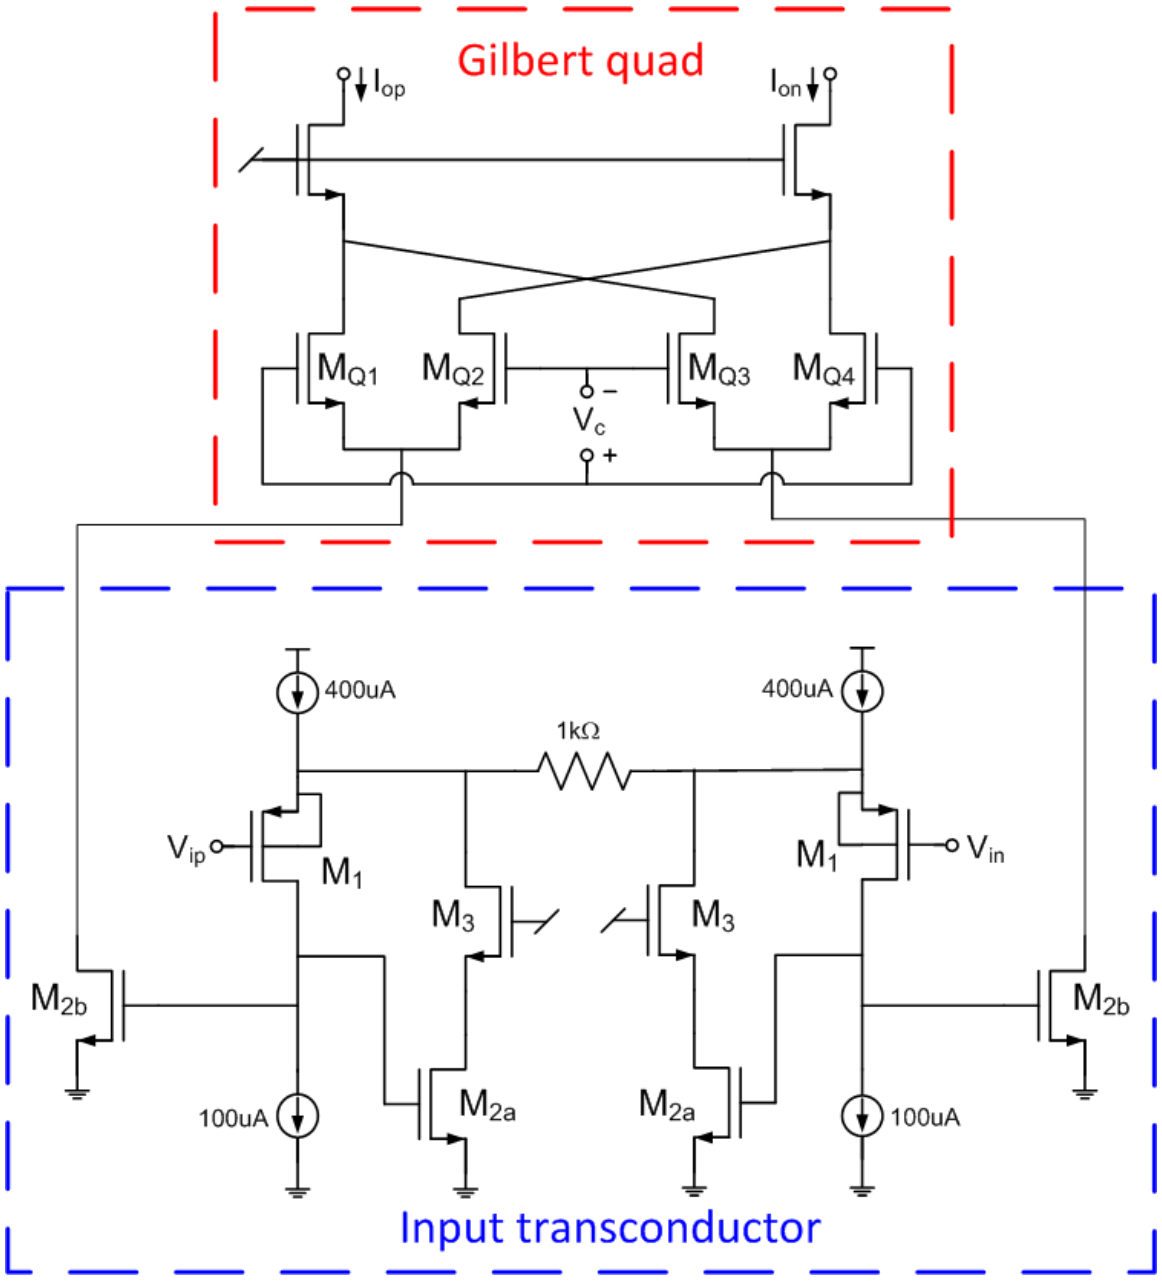

Figure 4.5: Input transconductor and Gilbert quad VGA

The VGA functionality of the circuit is implemented by the four-transistor Gilbert quad in Figure 4.6, M_{Q1} to M_{Q4}. If we assume that the transistors in the quad are biased in strong inversion and behave like square-law devices, the output current of the quad is:

$$I_{out,diff} = \frac{V_c}{4}\left(\sqrt{I_{bias}\left(8 - \frac{V_{ctrl}^2}{V_{dsat}^2}\right) + 4I_{in,diff}} - \sqrt{I_{bias}\left(8 - \frac{V_c^2}{V_{dsat}^2}\right) - 4I_{in,diff}}\right) \quad (3.2)$$

Plots of the transfer function in (3.2) are plotted in Figure 4.6, with V_{dsat} = 200mV, I_{bias} = 100uA, and varying V_c. As can be seen in the figure, the transfer function does not display the desired characteristic of translinearity. In order to get translinear operation in the Gilbert quad, the exponential voltage-to-current transfer function of a BJT is required [62]. However, biasing the CMOS transistors in weak-inversion will give a good enough approximation of the BJT transfer function. Figure 4.7 shows circuit simulation results of the Gilbert Quad transfer function for varying Vc with the devices biased in weak inversion. As can be seen in the figure, the transfer function of this weak-inversion quad has linearity that is over 20dB better than its square law counterpart.

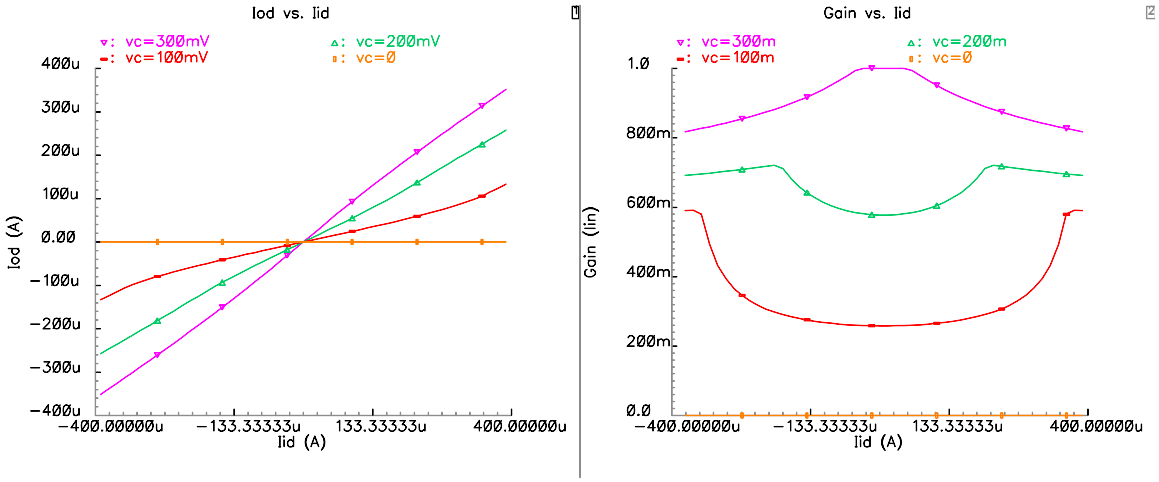

Figure 4.6: Gilbert Quad transfer function, assuming square-law devices in left plot. Derivative of transfer function (i.e. gain) in right plot shows significant nonlinearity

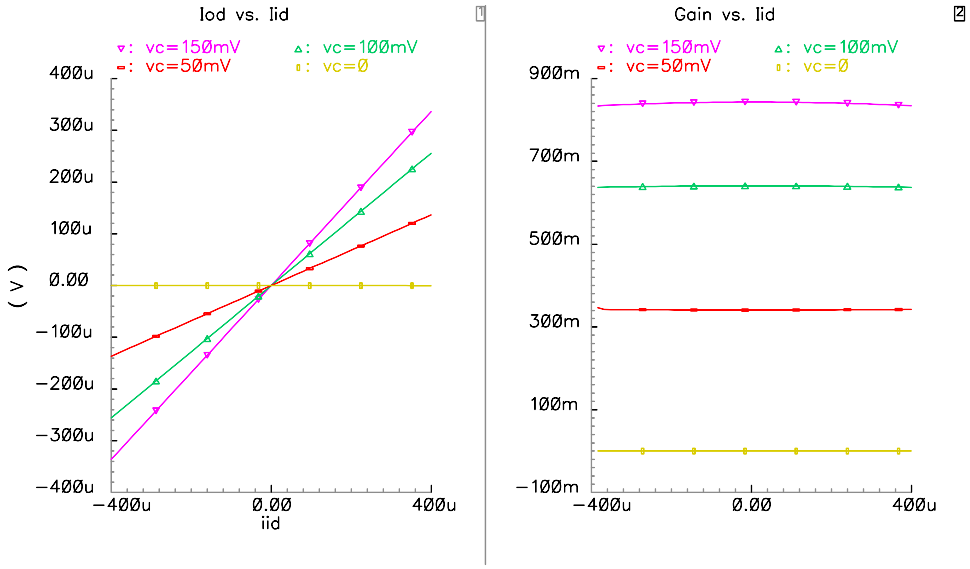

Figure 4.7: Gilbert Quad transfer function, with weak-inversion devices in left plot. Derivative of transfer function (i.e. gain) in right plot shows little nonlinearity in the signal path

There are two main limitations to using a weak-inversion Gilbert quad as the VGA device. First, in order to bias the quad transistors in their exponential, weak-inversion

regime, the transistors must be very wide. Therefore, the parasitic capacitances at the outputs of the quad are significant and can possibly reduce the bandwidth of the circuit. This issue is addressed by the inclusion of the high-speed, low-input-impedance current buffer shown in Figure 4.2 and described later in section 4.4.

Secondly, due to many second-order effects such as V_{TH} modulation and finite output impedance, the gain of the Gilbert quad is a non-linear function of the control voltage, V_c. A plot of the VGA gain vs. control voltage is shown in Figure 4.8, and nonlinearity in the gain profile is evident. In order to implement the desired carrier phase rotation operation, the gain of each component VGA must be precisely controlled; therefore, some mechanism to overcome the nonlinear voltage-to-gain relationship must be devised. This is described in the following section.

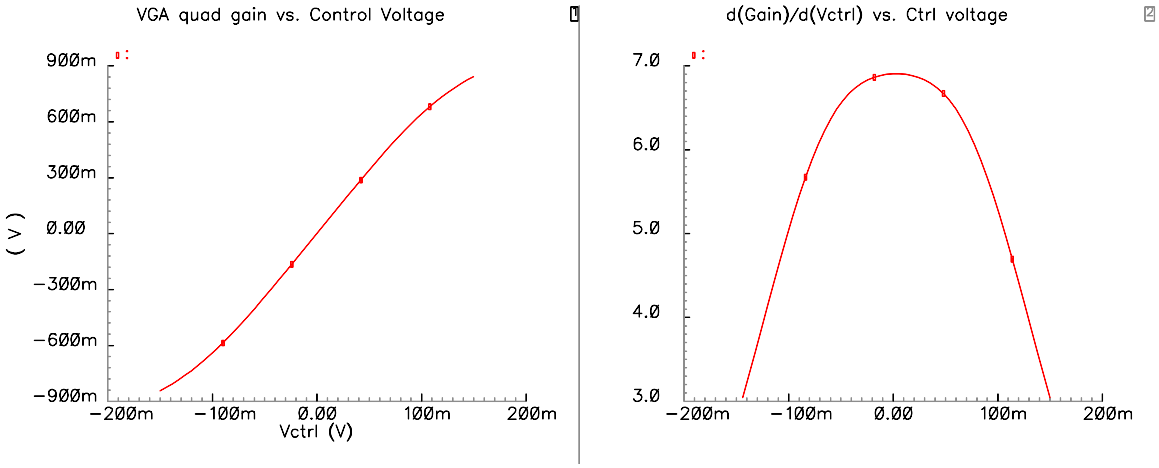

Figure 4.8: Gain of the Gilbert Quad versus control voltage (left). The derivative of the gain in the right plot shows significant nonlinearity in the gain transfer function

4.2.2 Feedback Tuning Circuit

In order to get the precise gain of cosθ or sinθ in the VGA blocks in the CPR, closed-loop techniques are required to overcome the gain nonlinearity shown in Figure 4.8. A master/slave tuning technique is used to create this precision gain, and a diagram of the circuit is shown in Figure 4.9. In this technique, a replica of the Gilbert quad is used. The differential current into the replica quad is set by fixed current sources at ½× I_{DAC}, a switched-current DAC is connected to the output of the replica quad, and an amplifier is connected in feedback around the replica quad. The high gain of the amplifier enforces a differential virtual ground at the quad output; in order to maintain a zero differential voltage, the amplifier sets the differential voltage across the quad inputs so that the replica quad sinks the DAC current exactly. In this way, the gain of the quad is set to $I_{diff,DAC} / I_{diff,bias}$. The control voltage established on the replica quads is then used to control the primary quads in the CPR VGA. Therefore, the gain of each primary quad can be precisely controlled by the digital input to the DAC.

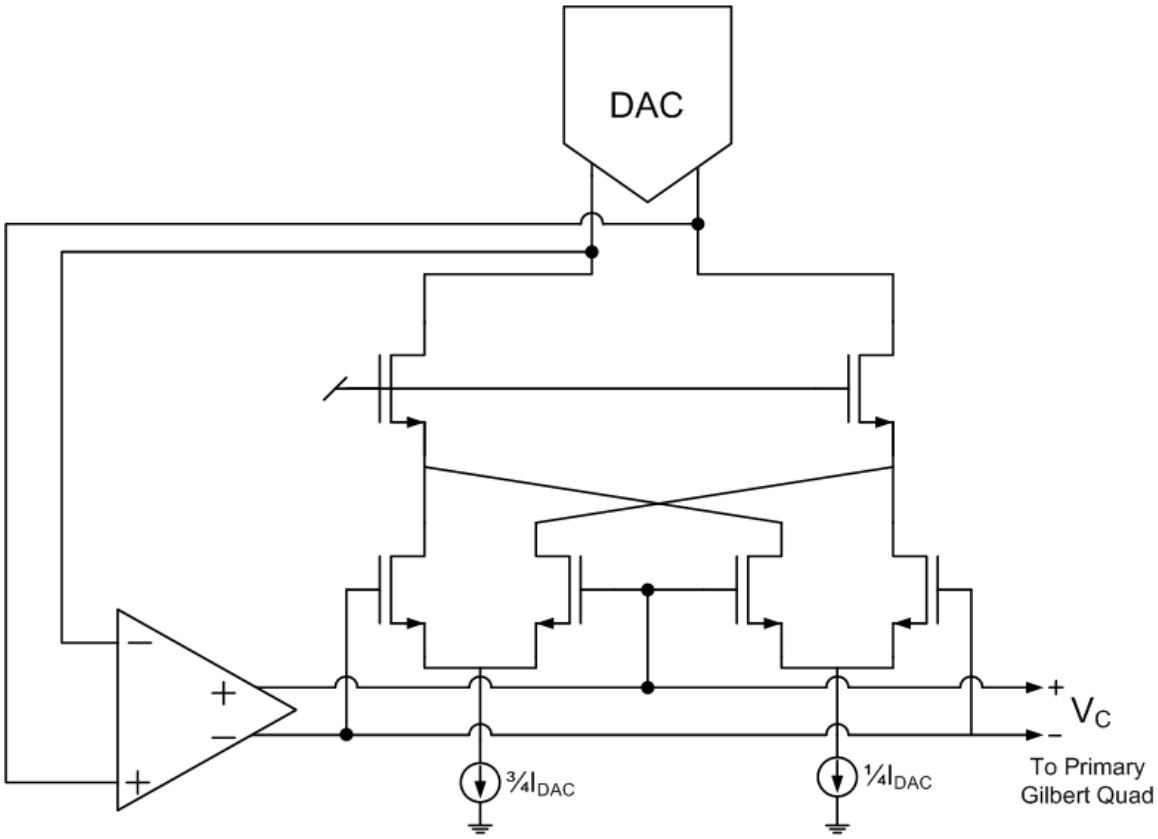

Figure 4.9: Replica tuning circuit for master/slave VGA tuning

Figure 4.10 shows an example where the gain of the VGA is set to zero. In this example, the DAC is configured to output a zero differential current. The principles of feedback and KCL enforce the current distribution in the replica quad as shown; the equal division of currents can only be achieved if the differential voltage on the gates of the quad is also zero. This quad control voltage is also forwarded to the primary VGAs, which would also have a gain of zero.

Figure 4.11 shows the transfer function of VGA gain vs. DAC current. When compared to Figure 4.8, the increased gain linearity of this approach is self-evident.

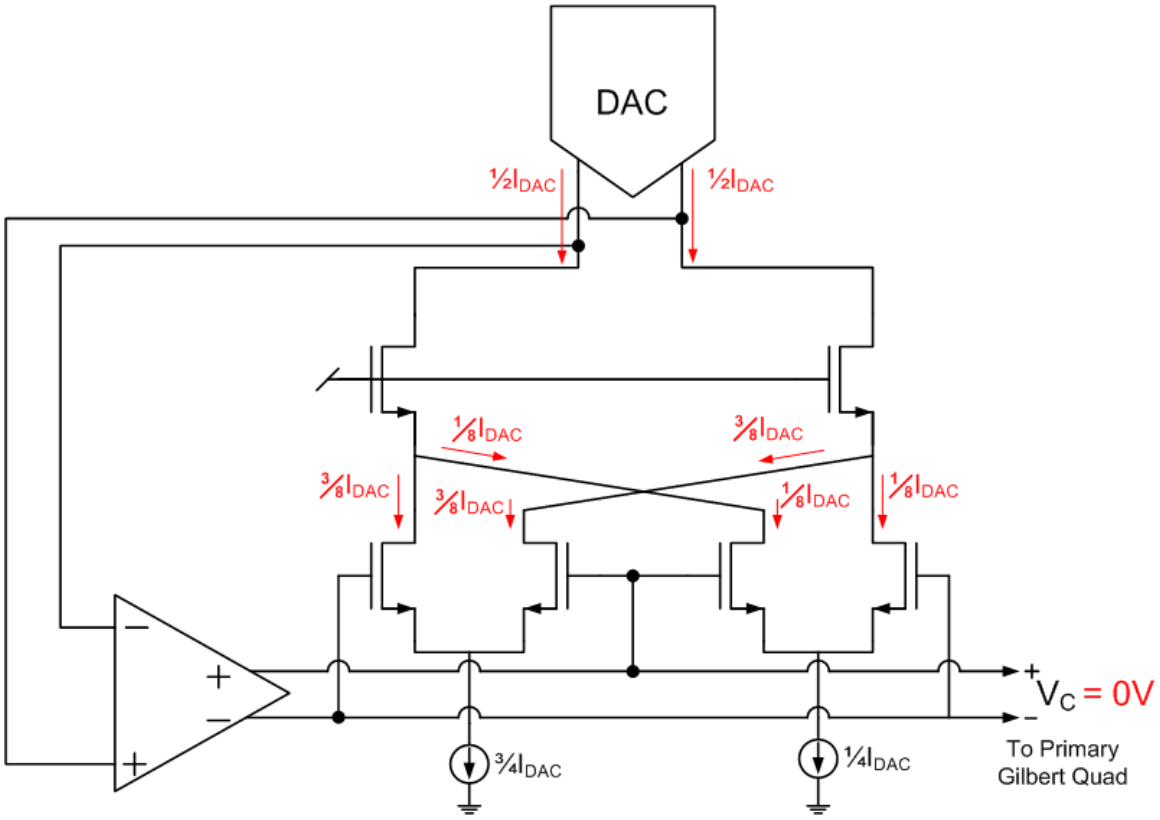

Figure 4.10: Example of replica tuning circuit with gain set to precisely zero

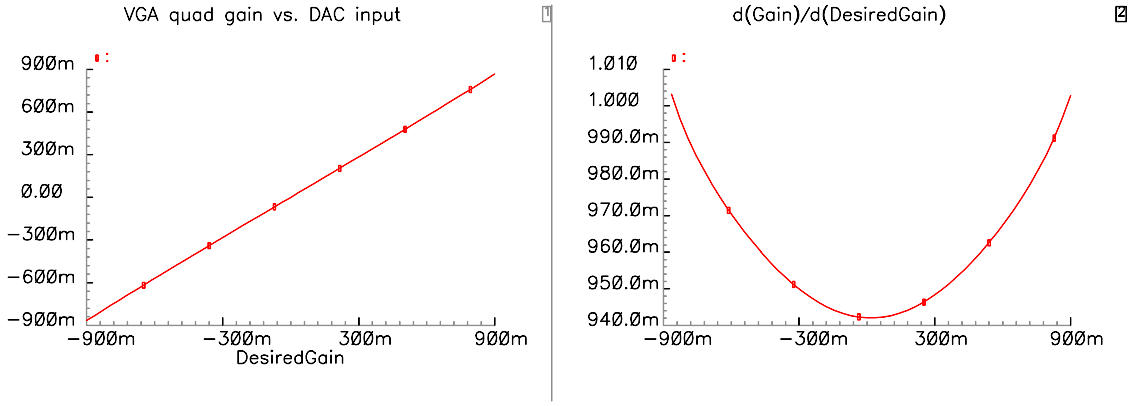

Figure 4.11: Gain of the Gilbert Quad versus desired gain as set by the DAC (left). The derivative of the gain in the right plot shows very little nonlinearity in the gain transfer function

4.2.3 DDFS Current-mode DAC

The previous section showed how the gain of the Gilbert quad VGA (and hence the phase rotation of the CPR block) was controlled by the switched-current DAC in Figure 4.9. However, as shown in (3.1), in order for the phase rotation operation to work, the output value of each DAC must be equal to either the sine or cosine of the desired rotation angle, θ. Switched-current DACs typically employ unit-sized current sources in their array in order to maintain good device matching and linearity. Therefore, some form of phase-to-amplitude conversion block, such as a DDFS, is required in the digital domain to generate the appropriate output.

In order to generate the speed and resolution requirements for the DDFS, the system performance in the presence of CPR error must be re-examined. Figure 3.11 (on page 70) showed that the overall system performance is unaffected by carrier phase error less than 5°. If we assume that 60GHz LO mismatch is less than 100ppm, then the frequency offset between TX and RX will be less than 6MHz. In order to maintain phase error less than 5°, the clock rate must be greater than: $6MHz \times \frac{360^\circ}{5^\circ} = 432MHz$. Therefore the clock rate of 500MHz was chosen for the DAC and DDFS blocks. The phase input to the DDFS block was truncated at 8 bits in order to keep phase quantization noise negligible as well.

Figure 4.12 shows the typical signal path for most DDFS DAC blocks. A control word for phase and a control word for frequency are the two data inputs. The frequency word gets integrated and added to the phase word to generate the total input phase. Most DDFS DACs utilize a ROM-based lookup table in order to perform the phase-to-

amplitude operation [64]. Subsequent to the ROM table lookup, the digital sine wave then is input into a standard linear DAC, where element-selection logic (typically row and column decode) is used to activate the desired number of DAC current elements. Lastly, the currents from these DAC elements are summed together at the output of the DAC to create the analog sine waveform. This traditional approach suffers from high levels of power dissipation, however. At clock frequencies in the hundreds of MHz, the ROM table can consume significant power, dominating overall system power, as in the case of [65].

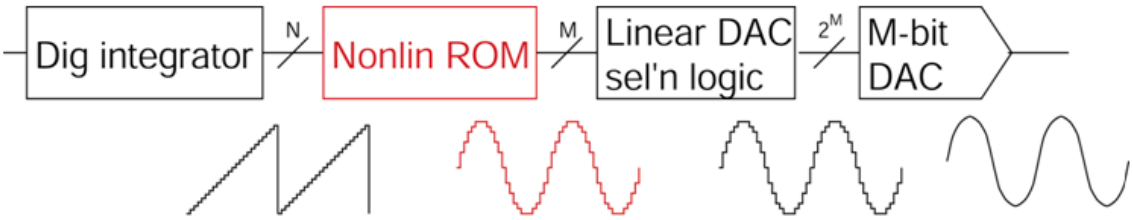

Figure 4.12: Traditional ROM-based DDFS architecture

Figure 4.13: DDFS architecture with trigonometric function embedded in DAC selection logic

An alternative approach to the ROM lookup architecture is shown in Figure 4.13, and is similar to the architecture reported in [66]. Here the ROM table is removed, and

instead the phase-to-amplitude operation is directly embedded within the element selection logic of the DAC. A. As with traditional ROM-based DDFS blocks, a digital integrator is used to create a phase ramp. Then, the phase ramp is converted into a digital sine wave by the nonlinear operation in the selection logic. This can be better understood by comparing the selection logic in this DAC to the selection logic in a traditional DAC. Figure 4.14 shows the transfer function for both types of DAC. In a traditional DAC, each additional LSB increase in the digital input code causes one additional LSB current source to be activated, as shown in the figure. However, with the nonlinear selection logic proposed here, each additional digital LSB does not activate an additional LSB current source. Instead, the selection logic is hard-wired to implement the phase-to-amplitude operation, as shown in Figure 4.14. As can be seen in the figures, each LSB current step is identically sized, but the transfer function of the selection logic is a nonlinear function. Furthermore, it can be noted that the DAC has fewer output amplitude levels than input phase levels. In this architecture, the 8-bit phase input generates a 6-bit DAC amplitude output, but the quantization in amplitude does not impact the carrier phase error and only creates minor levels of additional white noise.

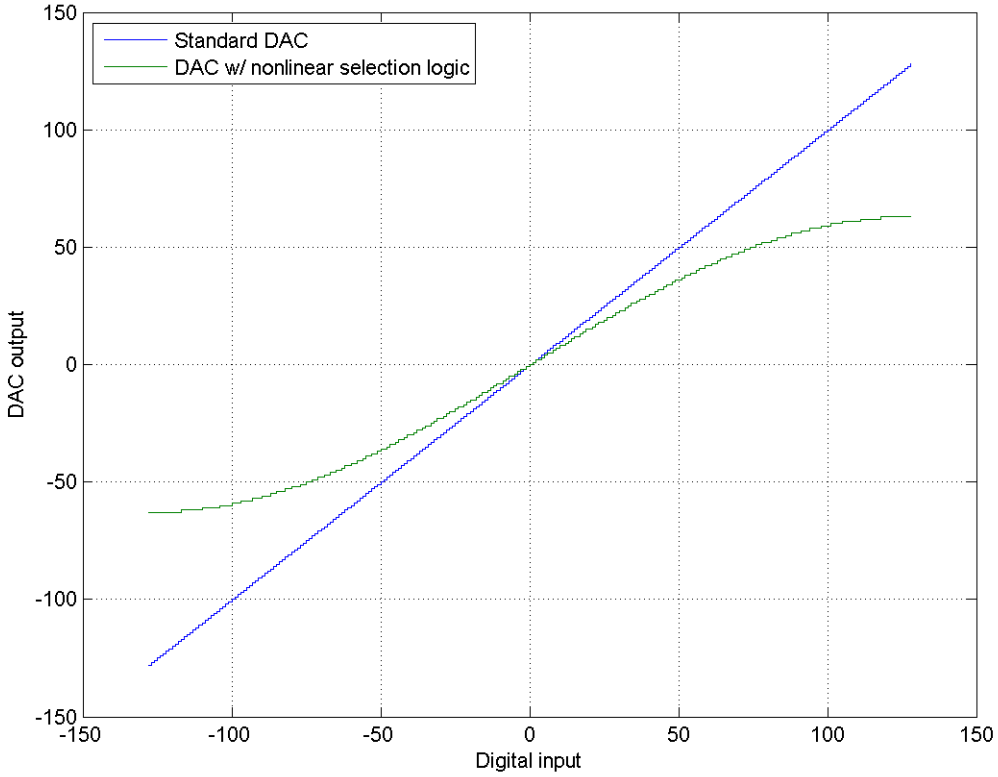

Figure 4.14: Transfer function of standard DAC and a DAC with nonlinear selection logic

4.3 Decision-feedback Equalizer

As described in chapter 3, the role of the decision-feedback equalizer (DFE) is to remove unwanted signal energy caused by channel multipath reflections and thus properly condition the received signal for demodulation and detection. The DFE consists of 16 complex-valued "taps" (8 in the I-channel and 8 in the Q-channel). Figure 4.15 shows the top-level architecture of a single channel of the DFE block. (i.e. there are two

instances of Figure 4.15; one for the I-channel and one for the Q-channel.) The DFE block runs at the symbol rate, so all the circuits within are clocked at 1GHz.

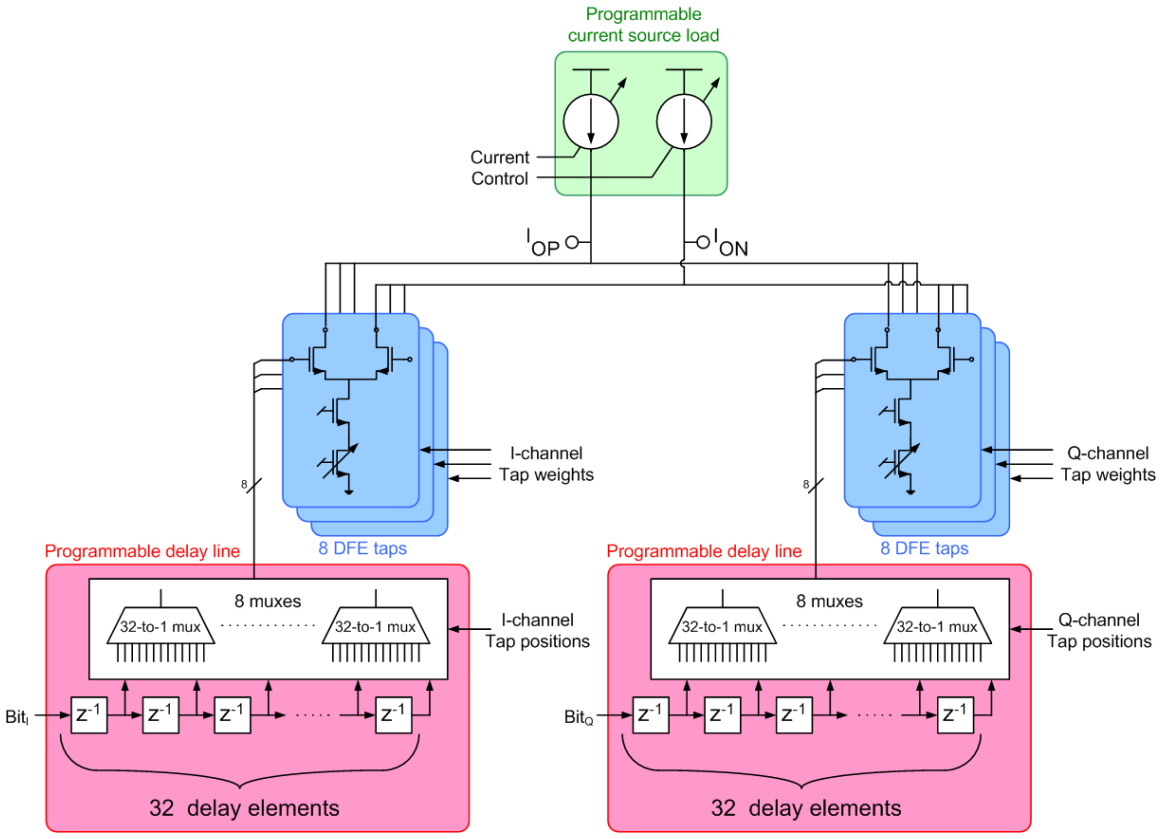

Figure 4.15: Top level block diagram of one channel of decision feedback equalizer

4.3.1 DFE Tap Structure with Programmable Current Bias

As shown in Figure 4.2, the DFE output is in the current domain in order to facilitate the summation operation that is inherent to the equalization process. A block diagram of the analog portion of a single DFE current-mode tap is shown in Figure 4.16. The magnitude of each tap is set by tying the desired number of gates of the NMOS current source to either a bias voltage or ground to activate or deactivate the transistor, respectively. Each tap has 7 current source transistors, so each tap can synthesize a DFE

value from $-7\times I_{LSB}$ to $+7\times I_{LSB}$ in steps of I_{LSB}. If a DFE tap value beyond that range is required, more than one DFE taps can be used together to synthesize a single larger DFE tap.

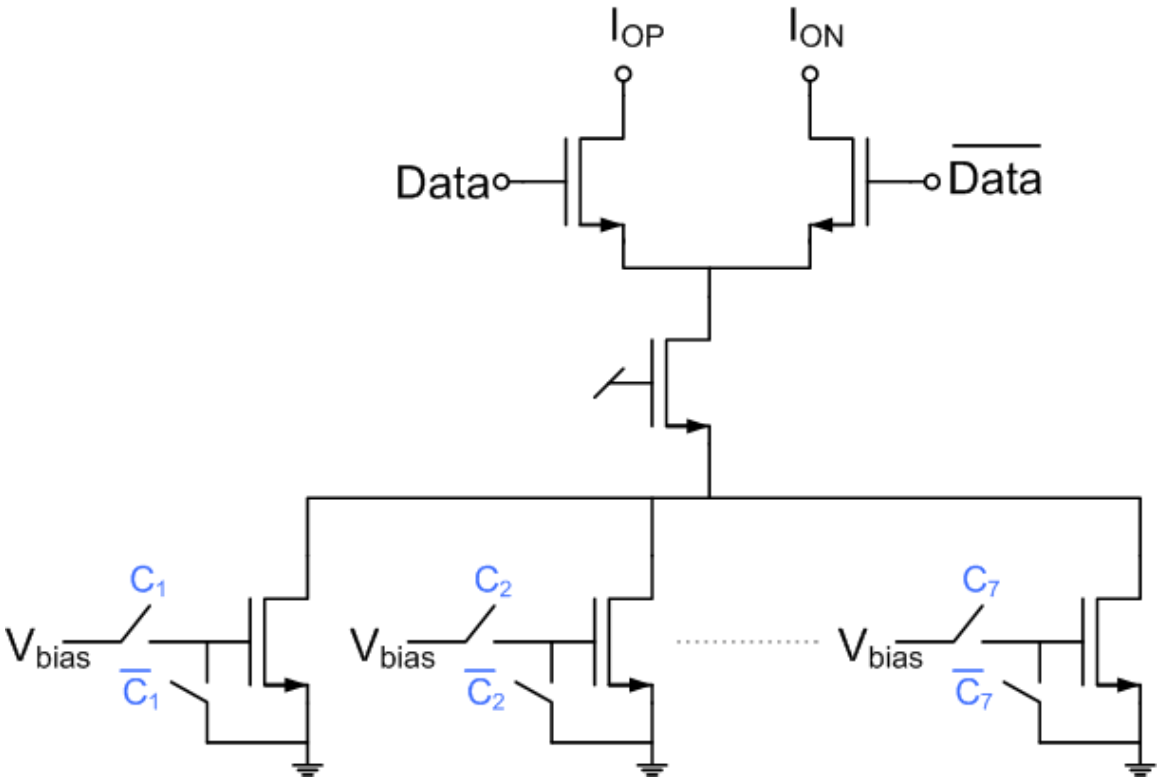

Figure 4.16: Analog portion of a single DFE tap

Figure 4.17 shows the full schematic for a tap, including the digital logic and clock circuitry that precede the analog DFE current tap. An input latch samples the delayed bit decision and generates a differential output bit signal. A crossbar circuit is used to either pass or invert the differential bit signal, depending on the sign of the DFE tap value. This signal then passes through a clock gating device (composed of an AND gate). Clock gating is used to disable the switching activity on the current tap in the case that DFE tap

value is equal to zero, (i.e. the DFE tap is not being used.) This technique reduces noise coupling to the output nodes from the digital switching in scenarios when some DFE taps are not needed.

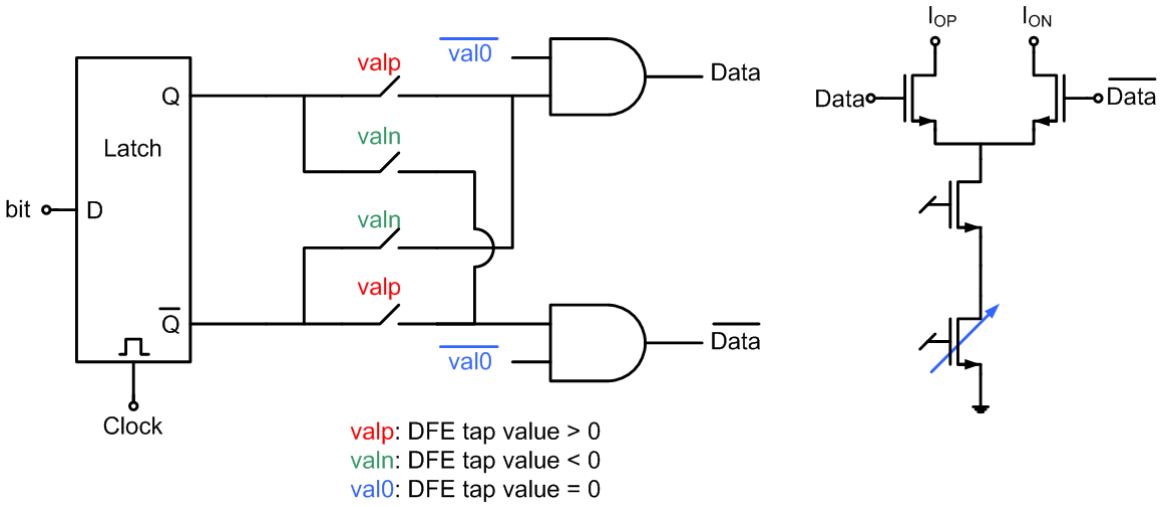

Figure 4.17: Full schematic of a single DFE tap

4.3.2 Programmable Current Source Load

As shown in Figure 4.15, the DFE contains a programmable current source load. This block is used so that DFE does not inject any common-mode current into the receive path. As can be seen in Figure 4.16, each DFE tap is composed of a switchable NMOS current source, with a tail current equal to $n \times I_{LSB}$. In order to prevent common mode current injection, two PMOS current sources are required, each with a current of $n \times I_{LSB}/2$. Rather than distributing these PMOS current sources into each DFE tap, it is more area-efficient to lump them all together in this single block. A schematic of the programmable current source load is shown in Figure 4.18. Also, since each side of this current source load is independently programmable, this block can be used to inject

and/or cancel differential offset at this node, which is generally useful during measurement and test.

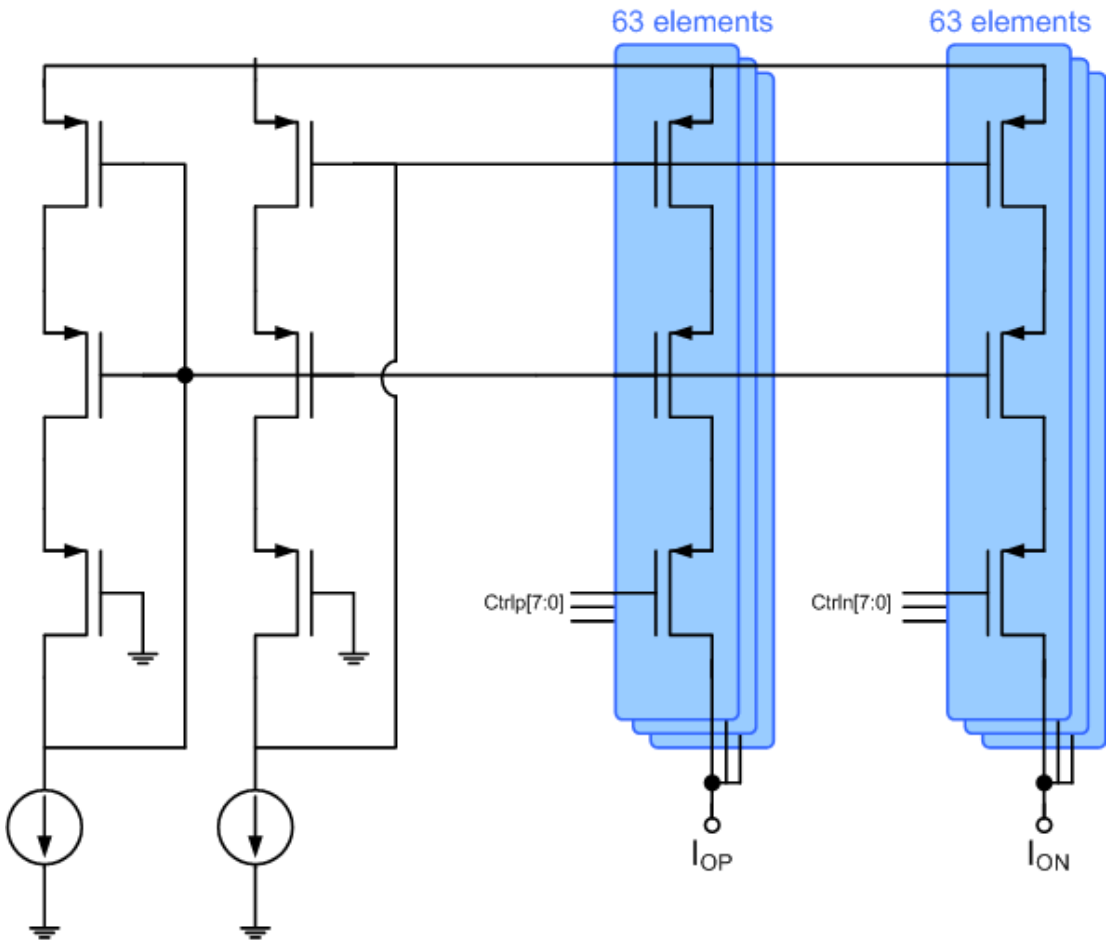

Figure 4.18: Programmable PMOS current source for DFE

4.3.3 Programmable Delay Line

The third block in the DFE is the programmable delay line. The delay line is necessary for the DFE in order to store the past bit decisions that are required in the equalization process. Unlike traditional DFEs, this block was made programmable in order to enable the adaptive tap allocation scheme described in section 3.4.4.

Conceptually, this block is quite simple: the outputs of the length-32 delay line are input into the switching block. Eight of these 32 inputs are selected and routed to the DFE for use in equalization. Hence, the switching block is composed of eight 32-to-1 muxes, each controlled by a separate 5-bit control word.

4.4 Summing Current Buffer

As shown in Figure 4.2 (page 114), the receiver has a current-domain summation node, where the currents from the CPR block and DFE blocks are combined. During design of the mixed-signal receiver, it became apparent that the parasitic capacitance at this node was quite high. The parasitic capacitance was due to two factors. First, these nodes had lots of devices connected to them; each summing node was connected to two Gilbert quad VGAs and 16 DFE current taps, and the resultant C_{DG} and C_{DB} from these devices were significant. Secondly, there was significant parasitic capacitance from all the routing required to hook these disparate blocks together. As will be discussed in section 4.5, the performance of the transimpedance amplifiers (TIA) that follow is very sensitive to input capacitance. As a result, it was determined that the large parasitic capacitance on these nodes could potentially degrade the bandwidth and performance of the subsequent blocks in the chain, and some form of buffering was required to decouple the DFE and CPR capacitance from the TIA inputs.

The input and output signals are in the current domain, so the buffer must be a current buffer. Desired characteristics include: a high circuit bandwidth so that the signal passes through the buffer undistorted; a low input impedance so that the parasitic

capacitances at the input do not degrade performance; a low output capacitance so that the operation of the TIA is not disturbed; good signal linearity; and a high input common mode voltage to accommodate the NMOS transistor stacks in the DFE and CPR blocks. Furthermore, since the TIA circuit does not have any common-mode control at the input, the buffer must be able to control its output common mode voltage.

Figure 4.19 shows the schematic of the current buffer. The core buffer is simply a cascoded PMOS current mirror. This simple circuit topology was chosen over other topologies due to the fact that it was the only topology found that could accommodate all the requirements mentioned above while still maintaining small-signal stability over process, temperature and supply variations. (Other schemes utilizing complex feedback techniques generally could not be stabilized at GHz-level bandwidths.) The worst-case bandwidth of this current mirror buffer was greater than 1.5GHz due to the use of thin PMOS devices biased at high current levels (~15uA/um). Common-mode output control was maintained by using NMOS current sources that were controlled in feedback by a common-mode sensing amplifier.

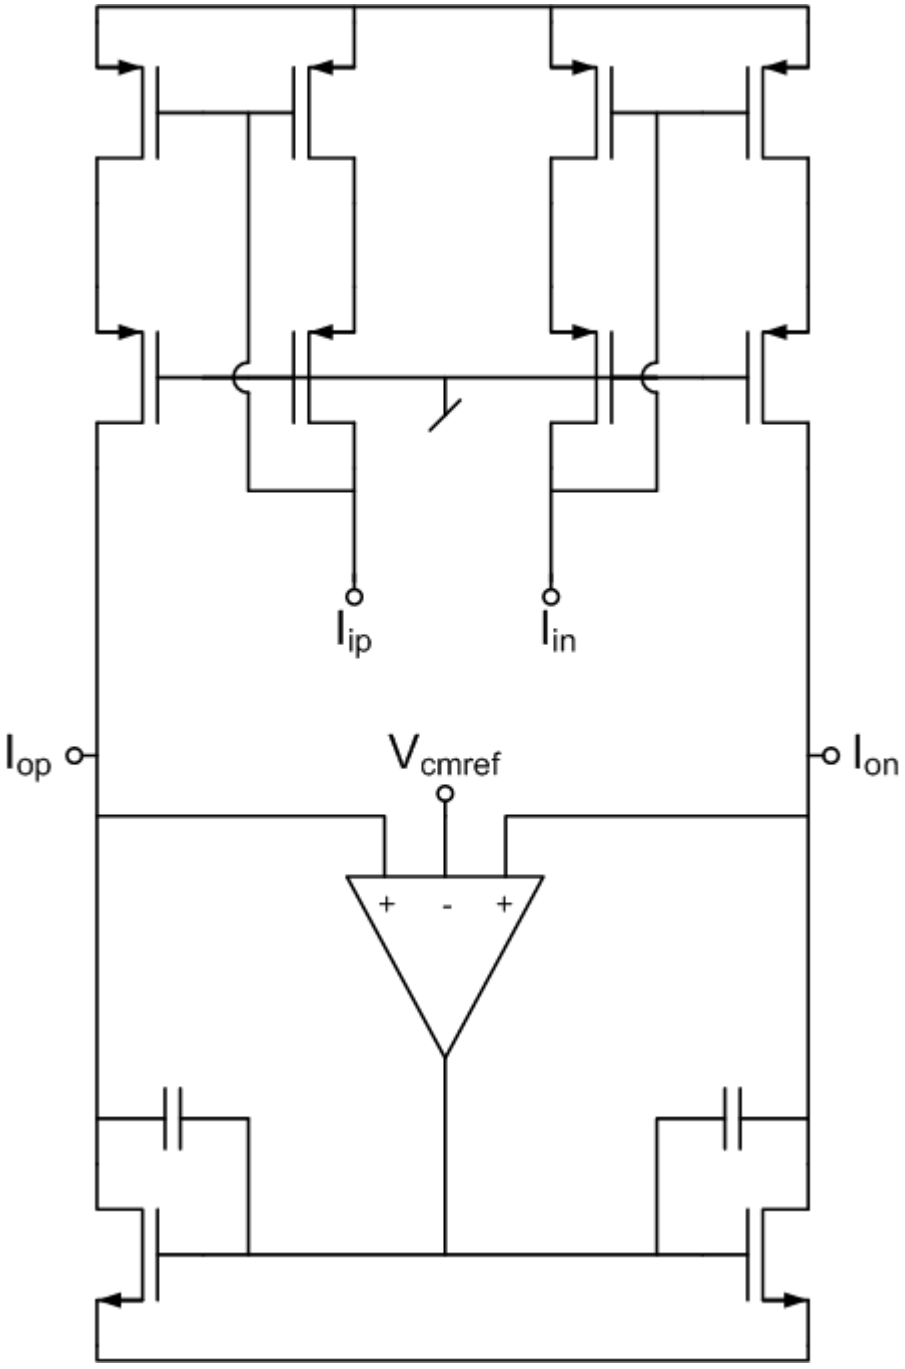

Figure 4.19: Unity-gain current buffer with output common mode control

4.5 Transimpedance Track-and-hold Amplifier

The role of the transimpedance track-and-hold amplifier (THTIA) is twofold: first, it needs to convert the incoming current-domain signal into a voltage domain signal for processing by the flash ADC. Also, it performs a track-and-hold function (also known as sample-and-hold) so that the input voltage to the flash ADC is held constant during the ADC evaluation mode.

The top-level schematic of the THTIA is shown in Figure 4.20. The THTIA is in track mode during ϕ1 and is in hold mode during ϕ2. The feedback resistor around the amplifier, R, provides the transimpedance functionality, and the feedback capacitor is used to store the output voltage during the THTIA's hold mode. If we assume the amplifier is ideal, the transfer function of the THTIA during ϕ1 is:

$$H_{THTIA}(s) = \frac{V_{out}}{I_{in}} = \frac{R}{1 + sRC} \tag{3.3}$$

In this design, R is 2kΩ and C is 50fF, so the output pole from the feedback network is at roughly 1.6GHz.

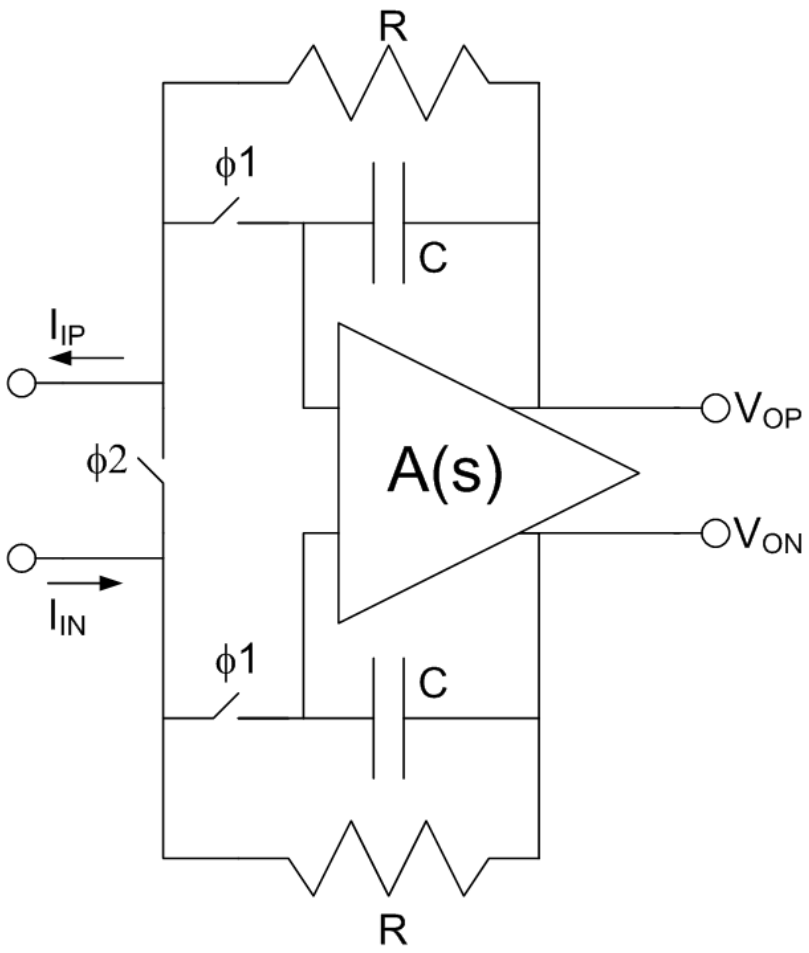

Figure 4.20: Top-level schematic of transimpedance track-and-hold amplifier. Φ1 is track mode. Φ2 is hold mode

There are two primary criterion for the design of the amplifier in the THTIA: first, the amp must have sufficient gain so that the closed-loop transfer function is predominantly determined by the feedback R and C; secondly, the amplifier must have sufficient bandwidth so that the output properly settles during ϕ1. Additionally, loop gain at bandwidth improves the linearity of the THTIA. If the amplifier has a DC gain of A_{DC} and a unity gain frequency of ω_u, the amplifier transfer function can be expressed as:

$$A(s) = \frac{A_{DC}}{1 + \frac{s}{A_{DC} \cdot \omega_u}} \tag{3.4}$$

Inserting this non-ideal amplifier into (3.3), we get:

$$H_{THTIA}(s) \approx \frac{R}{1 + sRC} \cdot \frac{A_{DC}}{1 + A_{DC}} \cdot \frac{1}{1 + \frac{s}{\omega_u}} \tag{3.5}$$

The first term of (3.5) is the ideal transfer function from (3.3). The second term is the gain error of the closed-loop transfer function due to finite DC gain, and the last term is an additional pole due to imperfect amplifier settling. In order to have gain error contribute $\ll$ 1 LSB of error in the ADC, we desire to have an amplifier with DC gain $\gg 2^N$, where N is the number of bits in the ADC. Also, in order for finite bandwidth to have minimal effect on the overall closed-loop transfer function, we need $\omega_u > 1/RC$ (= 1.6GHz). This bandwidth must be achieved while maintaining sufficient phase margin and gain margin to ensure loop stability.

4.5.1 THTIA amplifier

A schematic of the THTIA amplifier is shown in Figure 4.21. A modified two-stage amplifier topology was used, as the compensation network required for a traditional two-stage design would unnecessarily degrade the amplifier bandwidth. The first stage of the amplifier is a high-bandwidth circuit consisting of an NMOS differential pair (M1) and a PMOS diode load (M3a). The PMOS diodes form the input portion of output current stages (with current gain equal to 5) to drive the output PMOS devices (M3b). Additionally, PMOS current sources (M2) shunting the PMOS diodes are used to maintain a high current gain in the mirror with minimal current dissipation. The effective

G_M of the amplifier becomes the g_m of the input NMOS device multiplied by the current gain in the PMOS mirror, or $5\times gm_1$. Boosting the G_M in this manner is a power-efficient way to increase the unity gain frequency of the amplifier. The price paid for this approach is the nondominant pole of the PMOS current mirror; however in 90nm CMOS, it is relatively simple to push the nondominant pole into the 3-5GHz range. Simulations indicate that the DC gain of this amplifier is higher than 30dB and the unity gain bandwidth of the amplifier is at least 2GHz.

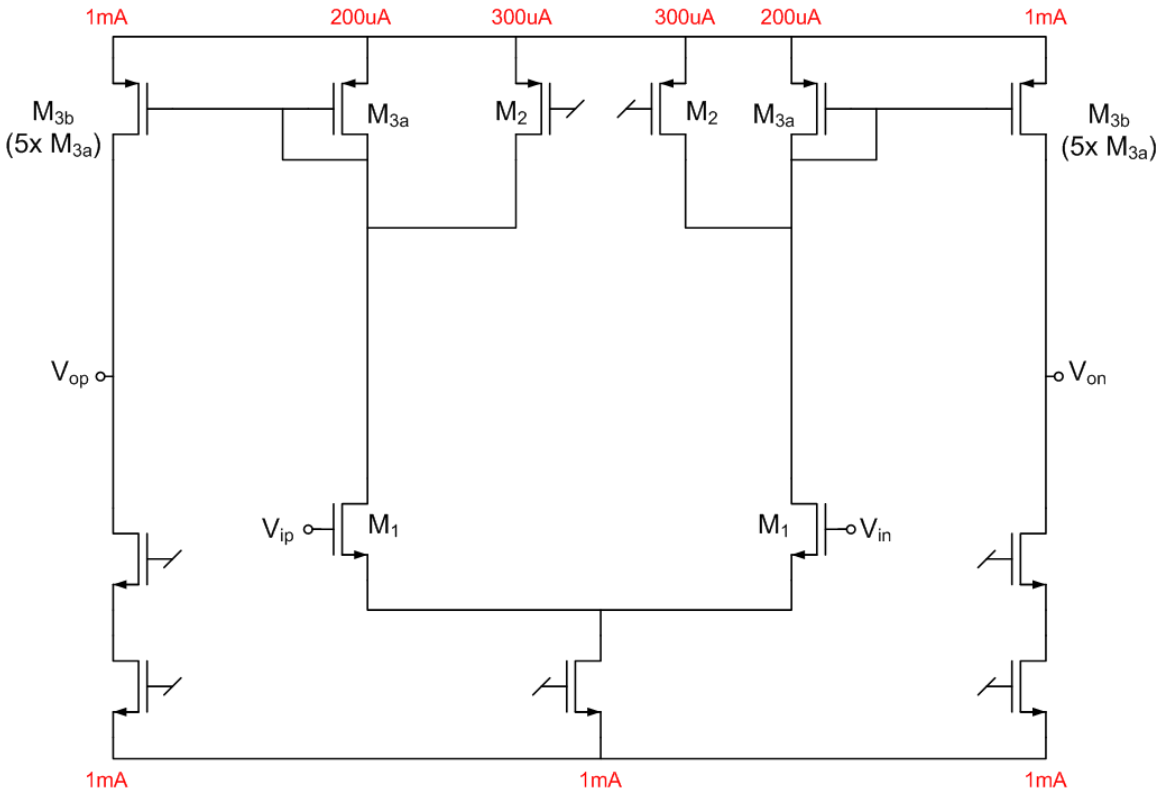

Figure 4.21: THTIA amplifier schematic

4.6 Flash ADC

The 2Gs/s, 4-bit flash ADC quantizes the voltage signal that is output by the THTIA circuit. As shown in Figure 4.2, there are two flash ADCs: one for the I-channel and one for the Q-channel. The top-level architecture of a single flash ADC is shown in Figure 4.22. All analog signals are differential, although the figure shows them as single-ended for simplicity.

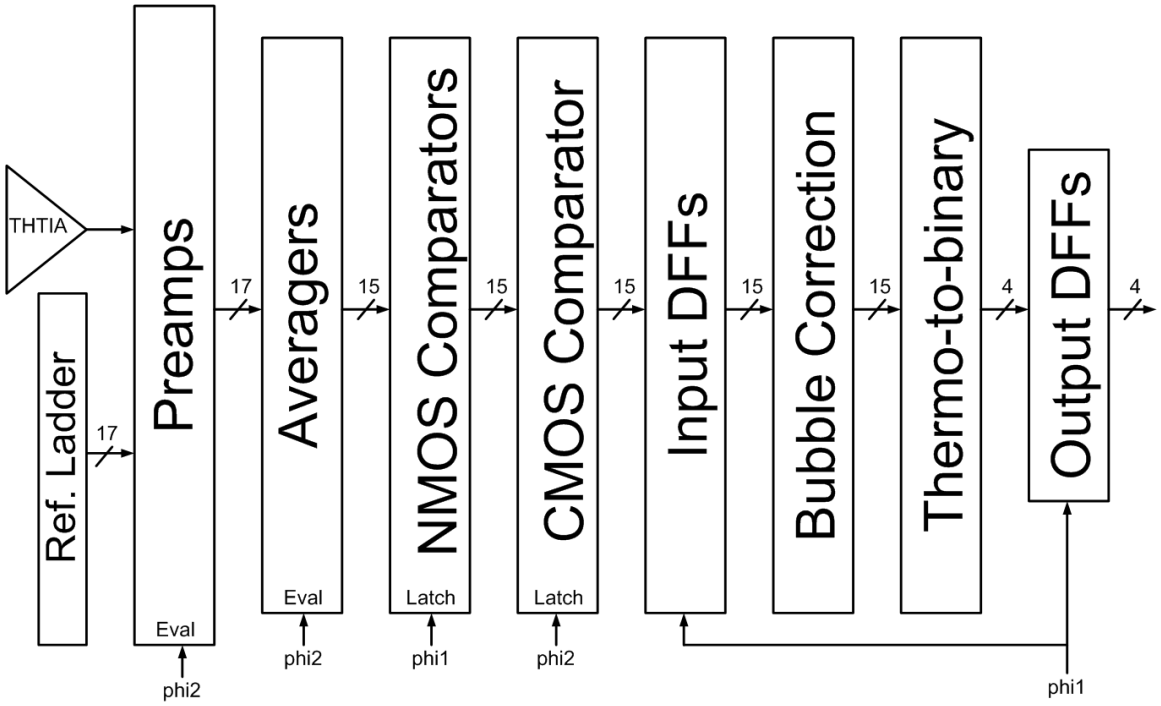

Figure 4.22: Top-level block diagram of 4-bit flash ADC

The THTIA is in track mode during phi1 and hold mode during phi2. During phi2, the preamps amplify the voltage difference between the THTIA output and the reference voltages supplied by the reference ladder. The outputs of the preamplifiers are then passed to the averaging circuits, which are also active during phi2. By the end of phi2,

the output voltage from the averaging circuits has fully settled, so at the conclusion of phi2 (the beginning of phi1), the NMOS comparator is released from reset and begins to regenerate. In order to facilitate low-power operation at 2Gs/s, the regeneration operation is pipelined over two stages of comparators. At the end of phi1, the NMOS comparators have completed most of the regeneration process; the CMOS comparators that follow enable the regeneration operation to complete during the subsequent half-period of the clock. After regeneration is complete, the 16 output levels are encoded in a 15-bit thermometer code, which is stored in the input D flip-flops (DFFs). A bubble-correction circuit is then used to remove any mismatch-induced sparkle codes in the thermometer code before use a ROM look-up table to perform the thermometer-to-binary conversion.

4.6.1 ADC Architecture and Mismatch Analysis

An ideal 4-bit ADC has 25.7dB of SNDR when excited by a full-scale sinusoid. However, offsets in the comparator threshold levels can degrade both the linearity of the ADC (measured in SFDR) as well as its SNDR. A typical rule of thumb for ADC design is that the 3σ value of the of comparator offset should be < ½ LSB, so the comparator offset should have a standard deviation of about 0.15 LSB. However, in order to test the efficacy of the ADC architecture proposed, a more stringent offset specification of $\sigma < 0.1$ LSB was used. As shown in Figure 4.23, this should result in a 3-sigma worst-case SNDR of about 24.5dB and a worst-case SFDR of about 34dB. The fullscale input level of the ADC is 640mV, peak-to-peak differential, corresponding to a 40mV differential LSB size. Therefore, with our 0.1 LSB target, the comparator offset should have a standard deviation less than 4mV.

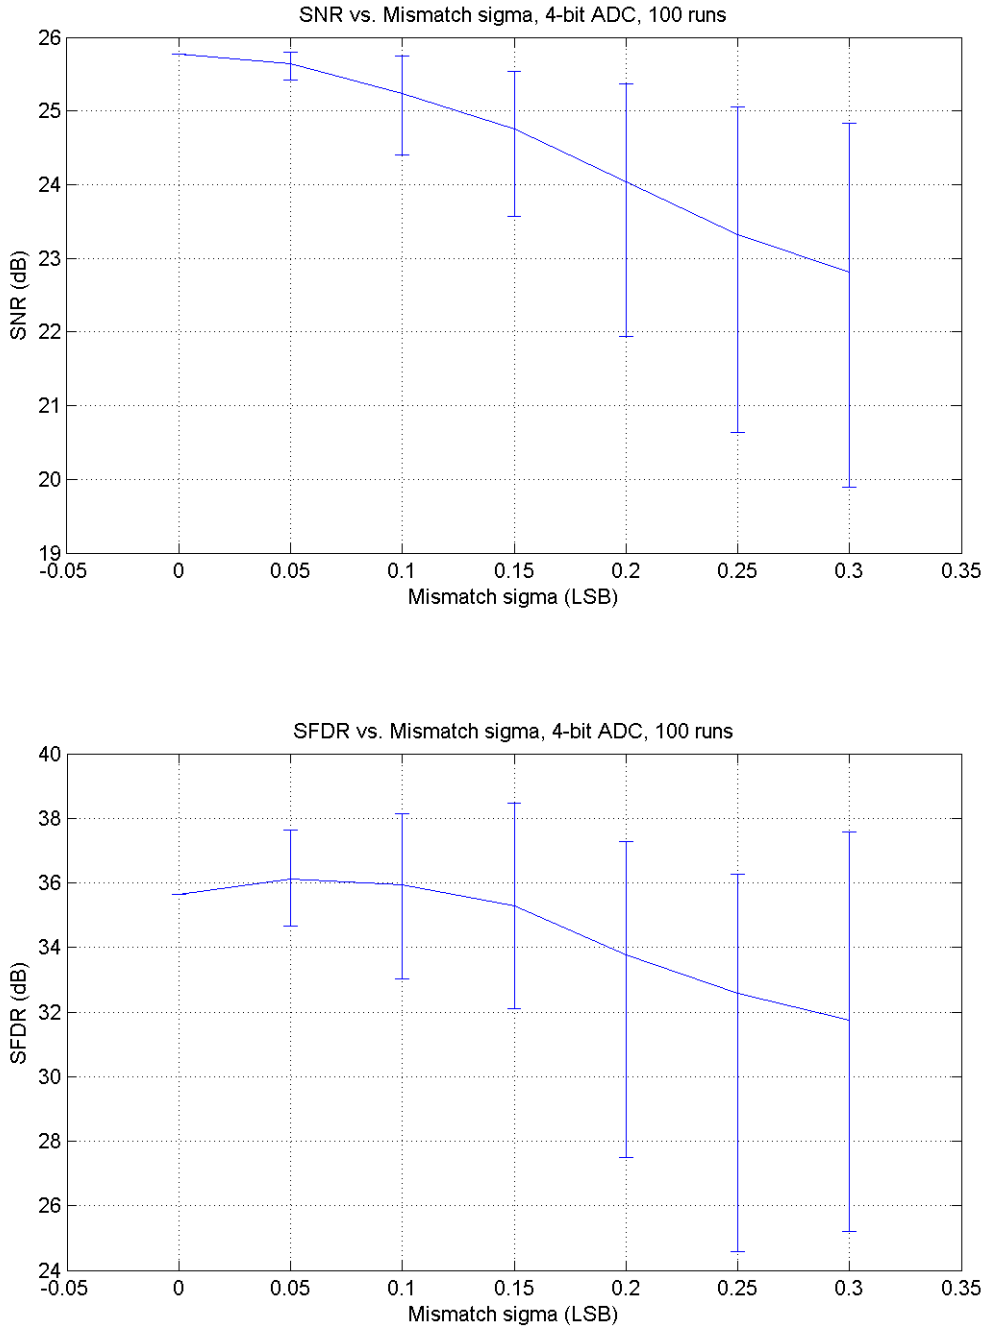

Figure 4.23: Results from a Monte Carlo run of a 4-bit ADC. The standard deviation of the comparator offset is swept in the X-axis. The midline represents the mean SNDR and SFDR. The error bars represent the best-case and worst-case measurements over 100 runs.

In order to meet this stringent offset requirement with minimum power dissipation, several techniques are required. As discussed in [67], preamps are frequently used to reduce comparator offset. As shown in Figure 4.24, the use of a preamp reduces the input-referred offset of the regenerative latch (which is often quite high) by the gain of the preamp. However, such a technique adds an additional offset from the preamp itself, and reducing the offset of the preamp through device sizing is a power-inefficient approach, as it both slows down the speed of the preamp and increases the load on the circuit that is driving it.

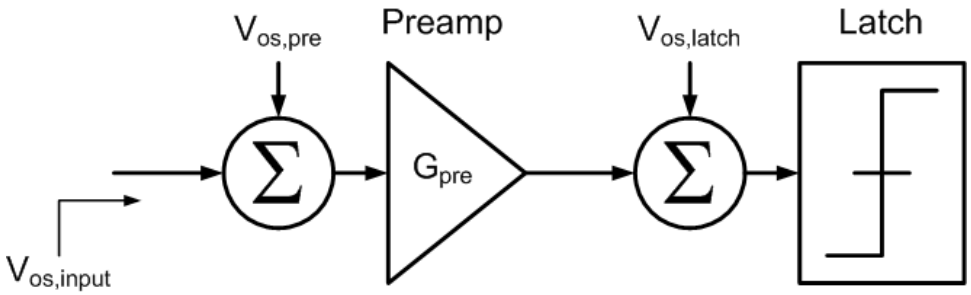

$$\sigma^2_{Vos,input} = \sigma^2_{Vos,pre} + \frac{\sigma^2_{Vos,latch}}{G^2_{pre}}$$

Figure 4.24: Offset-reduction using a preamp-based comparator

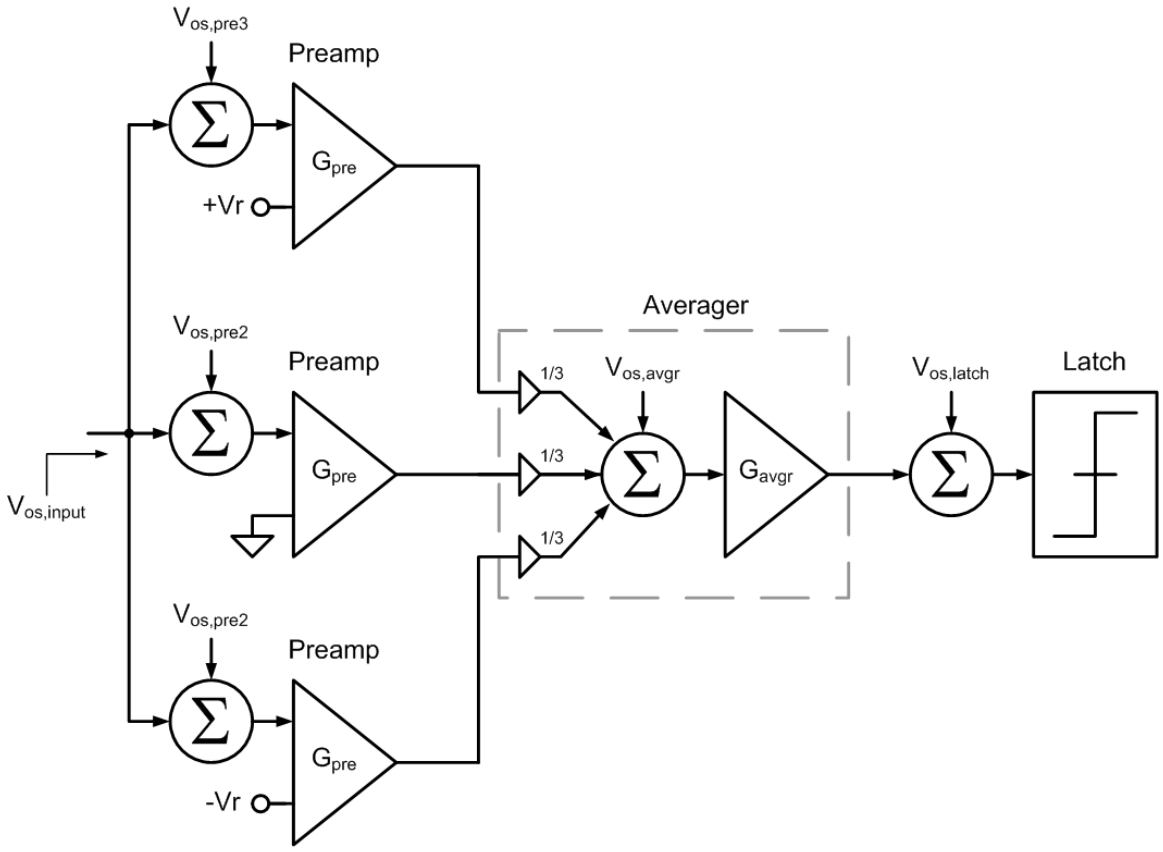

Figure 4.25: Offset-reduction using a preamp and averaging circuit

Another well-known technique for reducing comparator offset is the use of averaging [68]. In this technique, the averaging circuit is placed after the preamp, and the output of the averager is (a gained up version of) the average of several adjacent preamps. In Figure 4.25, a conceptual example circuit that averages 3 preamps is shown. In this example, the input-referred offset of the middle preamp is:

$$V_{os,input} = \frac{V_{os,pre\,1} + V_{os,pre\,2} + V_{os,pre\,3}}{3} + \frac{V_{os,avgr}}{G_{pre}} + \frac{V_{os,latch}}{G_{pre} \cdot G_{avgr}} \qquad (3.6)$$

Since each of the 3 preamps have an offset that is uncorrelated to one another, their variances add, resulting in a net reduction in the variance of the input-referred offset from the preamp by a factor of 3:

$$\sigma_{vos,input}^2 = \frac{\sigma_{vos,pre}^2}{3} + \frac{\sigma_{vos,avgr}^2}{G_{pre}^2} + \frac{\sigma_{vos,latch}^2}{G_{pre}^2 \cdot G_{avgr}^2} \tag{3.7}$$

There are several different ways to implement averaging, but the two most common are resistive averaging [69] and active averaging [67]. Resistive averaging is of limited utility in low-bit flash ADCs, as their edge effects tend to limit any potential gains when the number of comparators are low. A traditional active averaging circuit is shown in Figure 4.26. The averaging operation performed here uses three transconductor circuits to convert the voltage domain signal to current domain in order to facilitate easier summation. In this example, the input voltage is near the switching threshold of the middle preamp (Vi=0V), and the voltages of each node is listed in red. One can observe that the three transconductors that form the middle averaging circuit each have a different differential voltage applied to its input: ignoring any offset voltages, the middle transconductor sees a differential voltage of $A \cdot Vi$, and the top and bottom transconductor see differential inputs of $+A \cdot (Vi - Vr)$ and $-A \cdot (Vi - Vr)$, respectively.

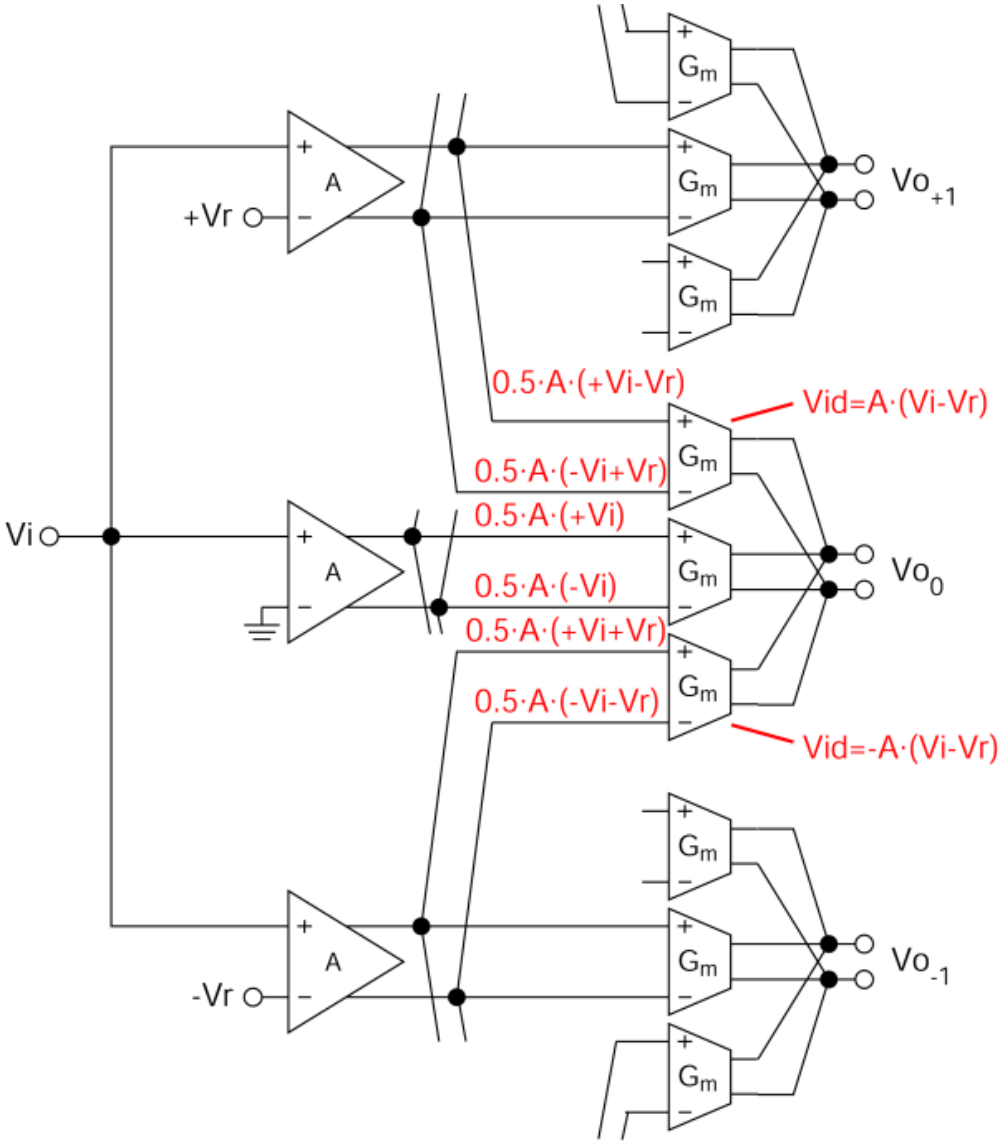

Figure 4.26: Averaging circuit using a traditional active averaging architecture. Voltages are shown for values of the input voltage near the comparator switching threshold.

This example highlights a major limitation of the traditional active averaging circuit topology: in reality, the G_M of the averaging circuit is not perfectly linear, but instead has a nonlinear dependence on the input differential voltage. In fact, for high-

speed operation, the averaging circuit transconductor is typically a NMOS differential pair. Assuming this diff pair has square-law behavior, the G_M of the structure can be calculated as:

$$G_M = g_{m0}\sqrt{1-\left(\frac{v_{id}}{v_{dsat\,0}}\right)^2} \tag{3.8}$$

where g_{m0} and v_{dsat0} are the quiescent small signal parameters of the transistors in the diff pair. As demonstrated in the Figure 4.26, two of the three averagers have a differential input voltage equal to $\pm A \cdot (Vi - Vr)$. Around the switching threshold, *Vi* will be 0V, and in the case of the 4-bit ADC discussed here, *Vr* = 40mV, and the gain of the preamp is around 2.5 . Also, the vdsat of the devices in the averager are about 125mV. Therefore, the transconductance of the outer two averaging transconductors reduces to:

$$G_M = g_{m0}\sqrt{1-\left(\frac{100mV}{125mV}\right)^2} = 0.36 \cdot g_{m0} \tag{3.9}$$

Plugging this value into the averager gain in Figure 4.26, we can recalculate the effect of the reduced averager gain on the input-referred offset:

$$\begin{aligned} V_{os,input} &= \frac{0.36 \cdot V_{os,pre\,1} + V_{os,pre\,2} + 0.36 \cdot V_{os,pre\,3}}{1 + (2 \cdot 0.36)} + \frac{V_{os,avgr}}{G_{pre}} \\ &+ \frac{V_{os,latch}}{G_{pre} \cdot \frac{G_{avgr} \cdot \left(1 + (2 \cdot 0.36)\right)}{3}} \end{aligned} \tag{3.10}$$

So, assuming that the preamp offsets are uncorrelated and of equal variance, the variance of the input-referred offset becomes:

$$\sigma^2_{vos,input} = \frac{\sigma^2_{vos,pre}}{2.45} + \frac{\sigma^2_{vos,avgr}}{G^2_{pre}} + \frac{2.9 \cdot \sigma^2_{vos,latch}}{G^2_{pre} \cdot G^2_{avgr}} \tag{3.11}$$

The input-referred offset in (3.11) is increased from that in (3.7) for two reasons. First, the averaging operation is less effective, since the adjacent preamps only get weighted with a reduced factor (0.36) instead of a fully weighted 3-way average. Therefore, the component of input-referred offset that is caused by offset in the preamp is increased, e.g. the first term in (3.11). Also the effective gain through the preamp is reduced due to the reduced G_M of the averaging circuit. As a result, the input-referred offset from the latch is also increased, e.g. the third term in (3.11).

Techniques to minimize the G_M reduction inferred from (3.8) are of limited utility. In essence, in order to minimize the G_M reduction effect, one must minimize the v_{id}/v_{dsat0} ratio by either reducing v_{id} or increasing v_{dsat0}. Reducing the v_{id} seen at the averager input is equivalent to reducing preamp gain, which is counterproductive to the reduction of input-referred offset. Since in a diff pair, $v_{dsat\,0} = 2I_{bias}/g_m$, increasing the v_{dsat0} of the diff pair in the averager can only be achieved by either reducing the g_m of the diff pair or by increasing its bias current; the first technique will cause the offset of the regenerative latch to dominate the total input-referred offset, and the second technique is simply an inefficient use of current.

Instead, modifications to the circuit architecture of Figure 4.26 are necessary to resolve the issue of G_M reduction in the averager circuit. Figure 4.27 shows a new averaging architecture, dubbed “staggered active averaging” by this dissertation. This architecture looks very similar to the traditional architecture shown in Figure 4.26, but the interconnect between the preamps and the averaging circuit has been modified. The wiring has been changed such the input differential voltage to all three G_M cells in the

averager is always $A \cdot Vi$. Instead of each G_M cell having a different differential "tilt" at its input, the G_M cells each have a different input common-mode voltage, as can be seen in Figure 4.27. (The common-mode voltage difference is $\pm A \cdot Vr = \pm 100mV$.) Simple NMOS diff pairs are very robust to common-mode voltage shifts of this magnitude, and provided that the tail current source of the diff pair has a suitably high impedance, the g_m variation of the diff pair with respect to input common-mode voltage is negligible. Therefore, the input-referred offset can be expressed by the formulation in (3.6) and (3.7), which represents roughly 15% reduction of $\sigma_{vos,input}$ compared to traditional active averaging techniques.

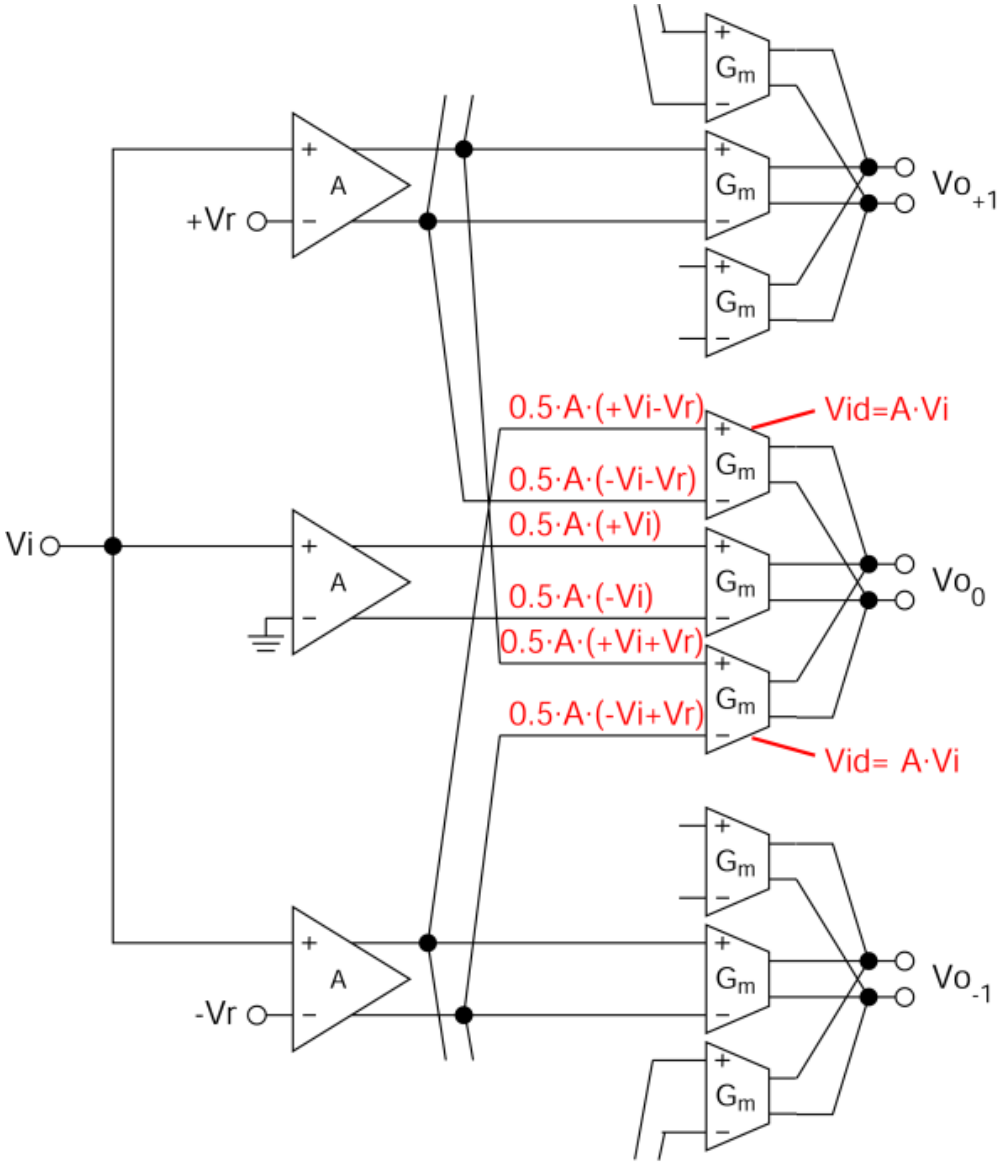

Figure 4.27: Averaging circuit using the new staggered active averaging architecture. Voltages are shown for values of the input voltage near the comparator switching threshold.

4.6.2 Preamplifier

A schematic of an individual preamp is shown in Figure 4.28. A reset switch is used in order to minimize the effect of comparator overloading and hysteresis.

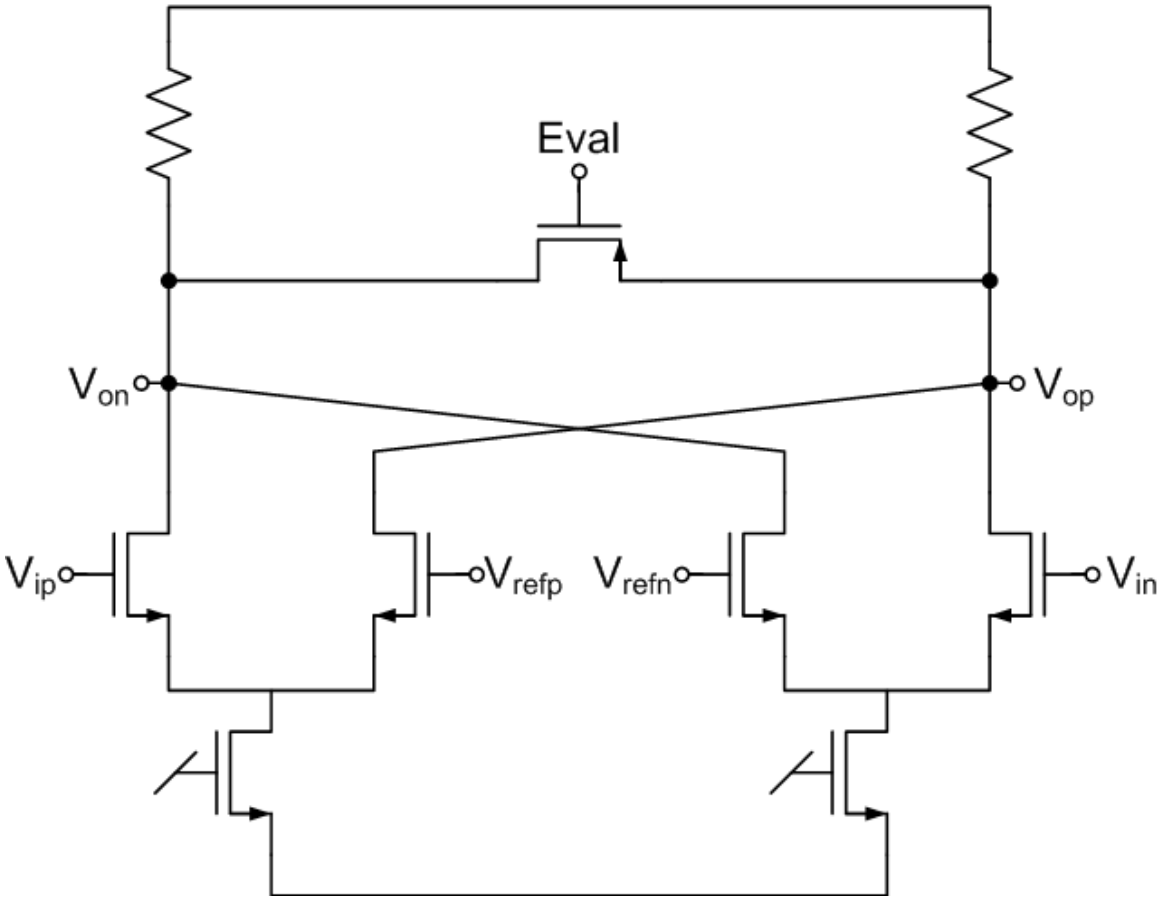

Figure 4.28: Preamplifier schematic

4.6.3 Averager

As depicted in Figure 4.27, each averager consists of three transconductance cells with outputs tied together. A schematic of an individual transconductance cell is shown in Figure 4.29. As with the preamp circuit, the averager also has a reset switch.

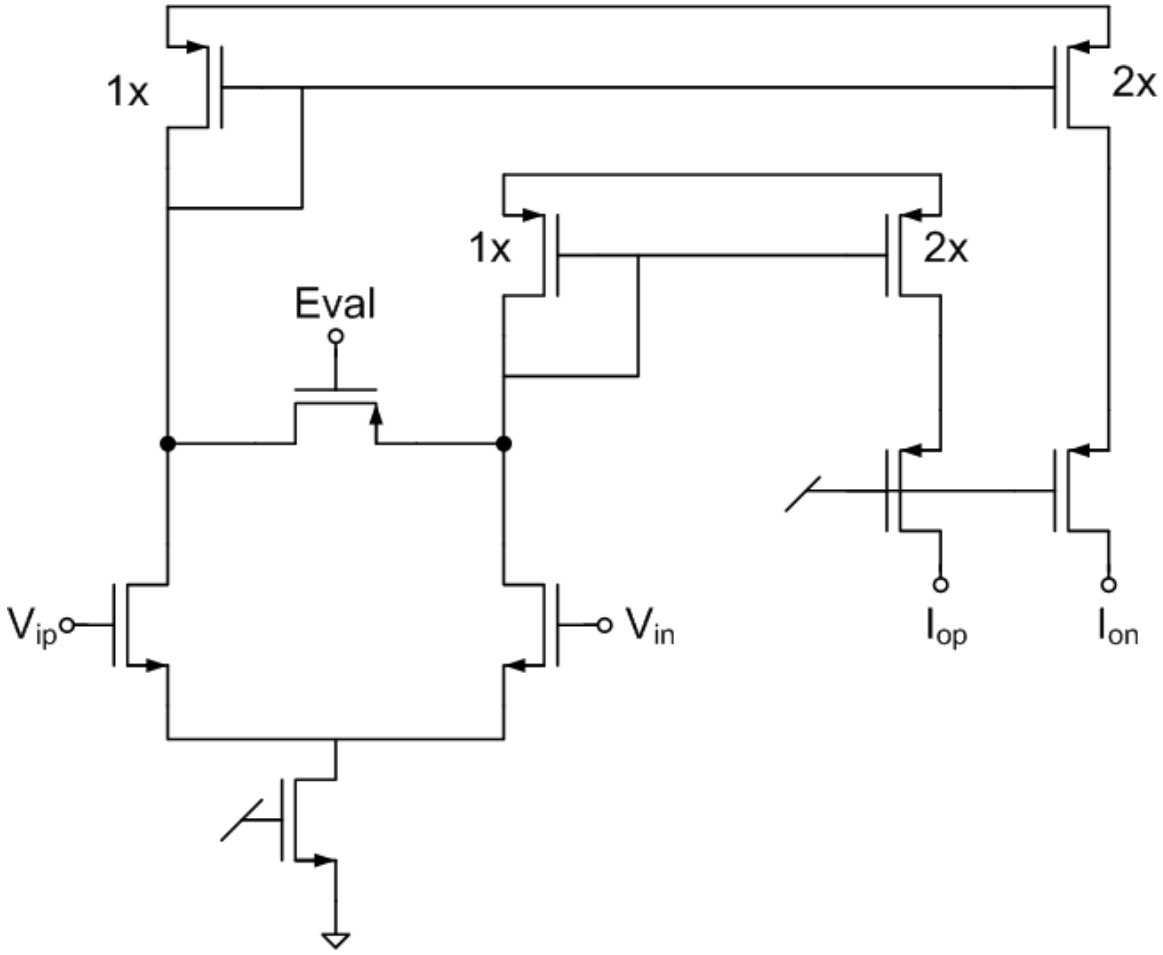

Figure 4.29: Schematic of single transconductor circuit within 3x averaging circuit

4.6.4 NMOS Comparator

The high-speed NMOS comparator is shown in Figure 4.30. (Note: The same nodes are used as the current input from the averager and the voltage output to the CMOS comparator.) It uses NMOS devices only in order to achieve optimal regeneration speed. A static latch is used here, because the dynamic offset characteristics of the process were

not well-characterized. Hence, power consumption was increased in order to guarantee robust offset performance.

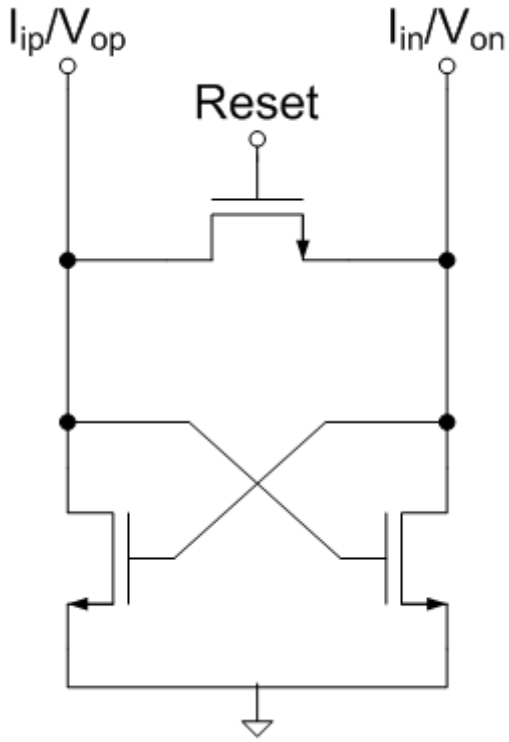

Figure 4.30: Static NMOS comparator

4.6.5 CMOS Comparator

As depicted in Figure 4.22, a second stage of regeneration was added in order to pipeline the regeneration process and allow more time for the digital data values to become valid. A schematic of the second-stage CMOS comparator is shown in Figure 4.31. A dynamic architecture is used for this comparator, as the inputs to this comparator should be large enough to overcome any dynamic offsets.

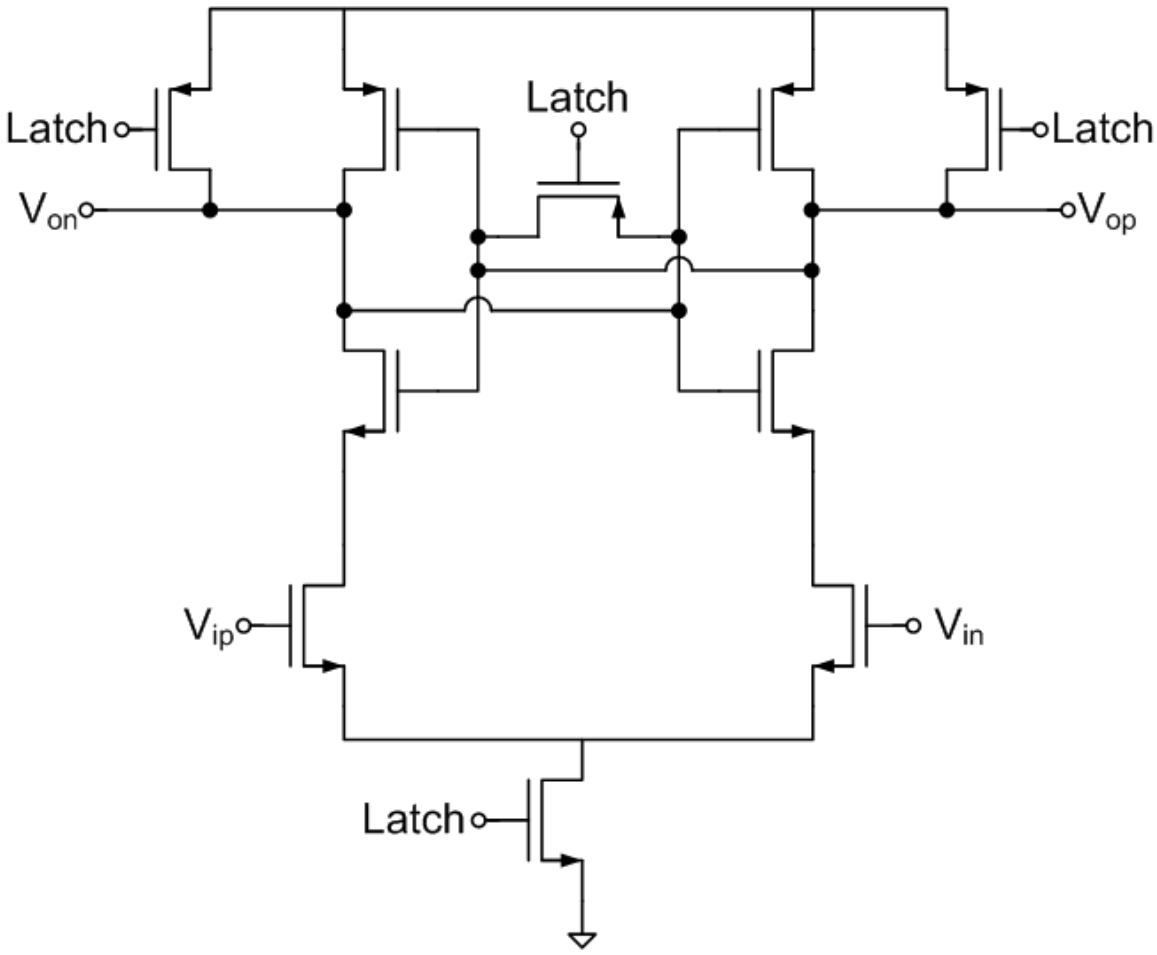

Figure 4.31: Dynamic CMOS comparator

4.6.6 Digital Logic

The digital logic at the backend of the ADC consists of D flip-flops (DFFs), bubble-correction circuitry, and a ROM look-up table for use in the thermometer-to-binary conversion. The bubble correction circuitry is used to correct for any errors in the thermometer coding that might result from static or dynamic offsets in the comparator array. It uses the two-out-of-three majority vote technique described in [70] to correct for sparkle codes of depth 1. The DFFs and ROM table use standard CMOS architectures.

5 Measurement Results

The prototype chip was fabricated in a 1.0V 90nm digital CMOS process provided by TSMC. The die measured 1.6mm x 1.55mm, although the chip is pad-limited. Active area is roughly 0.55mm^2. Figure 5.1 shows a photograph of the chip. A PCB was designed for test measurements, and chip-on-board packaging was used to bond the chip directly to the PCB board. Figure 5.2 shows the PCB and test setup.

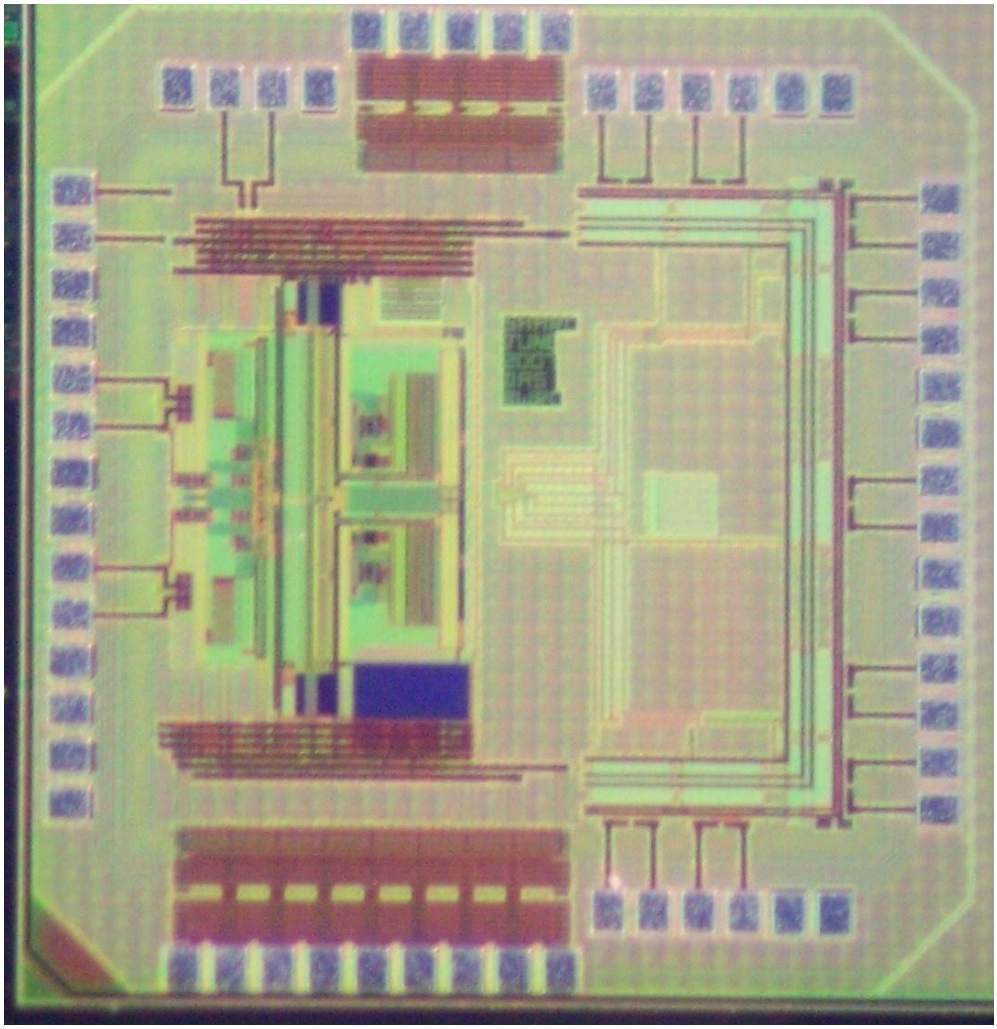

Figure 5.1: Die microphotograph

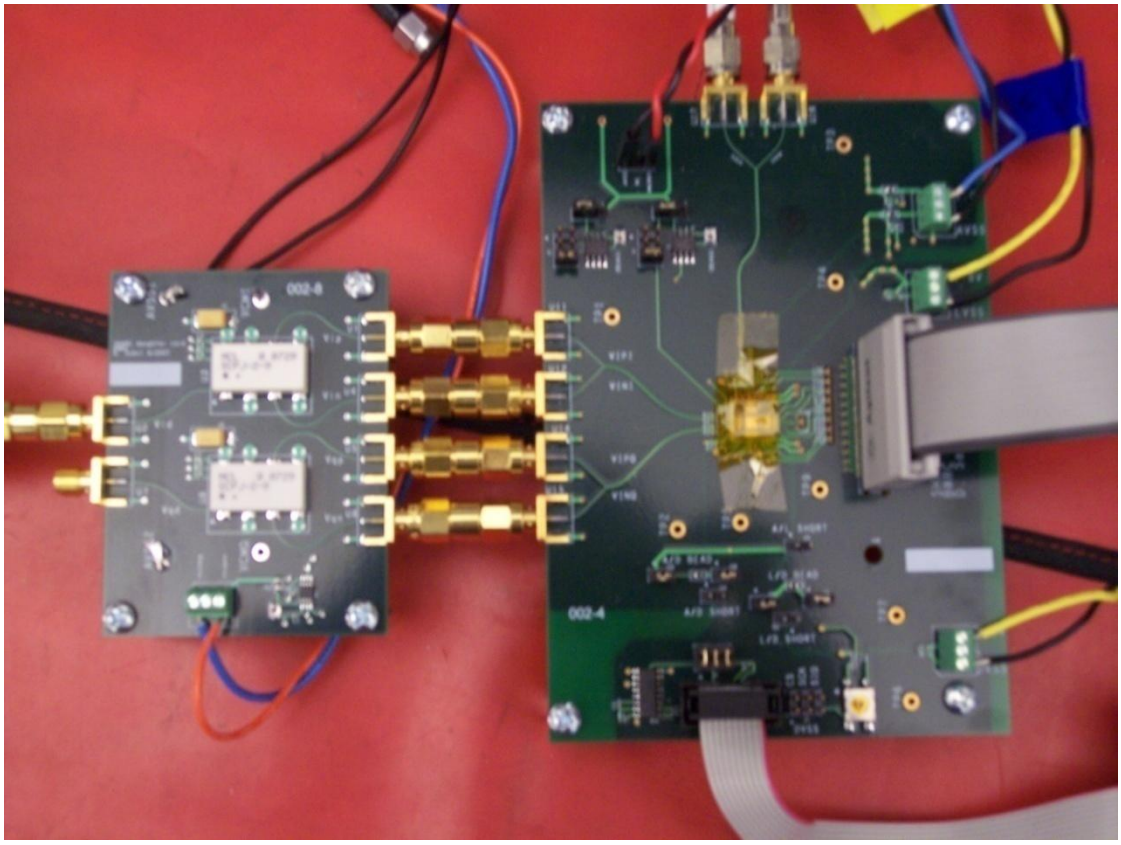

Figure 5.2: PCB configured for test

5.1 ADC Performance

The static performance of the ADC was characterized using INL and DNL tests. In these tests, the chip inputs were excited by slow sine waves, and the code density of the ADCs were observed. The signal frequency was roughly 300MHz and was chosen so that there were a prime number of periods during the 32kpt measurement record. The INL and DNL for both I-channel and Q-channel ADCs are shown in Figure 5.3. INL and DNL were both observed to be less than ±0.25 LSB, which indicates the efficacy of the staggered averaging approach discussed in section 4.6.1.

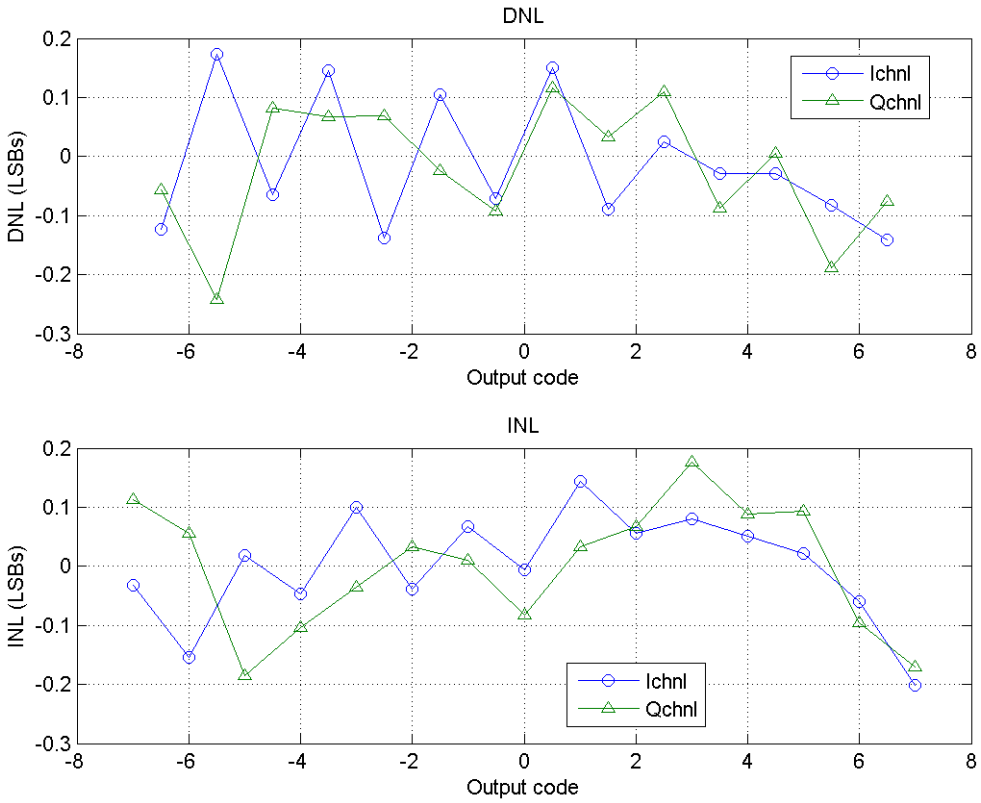

Figure 5.3: ADC DNL and INL

Figure 5.4 shows an FFT of a single ADC running at the full Nyquist rate. The clock rate is set at 2.048GHz and the input frequency is approximately 1GHz. The input signal amplitude is -0.5dBFS, which corresponds to 95% input loading. The output SNDR is 24.65dB, which corresponds to an 3.81 effective number of bits (ENOB). The THD is 35.5dB. The SFDR is 38.7dB and is limited by the sixth harmonic of the input signal; this is likely due to high-order harmonics being excited by the INL and DNL of the converters. It is worth noting that this measurement—and all the measurements described in this section—include the linearity of the transconductors and Gilbert quads

within the carrier phase rotator (CPR) block, as the input signal passes through the CPR en route to the ADC.

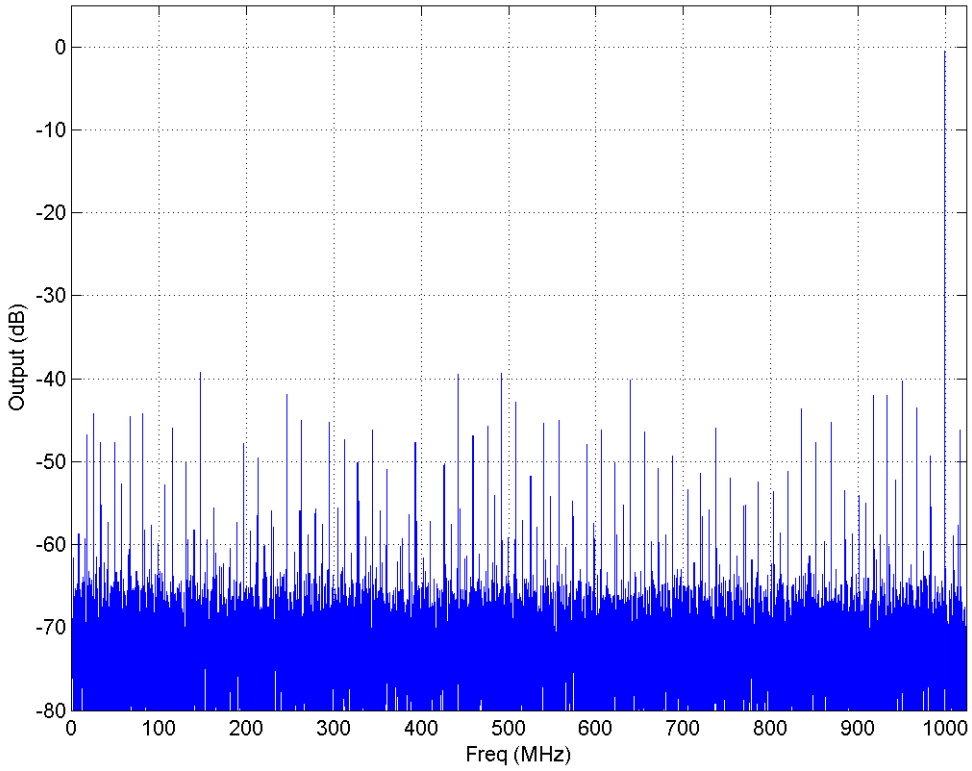

Figure 5.4: FFT of ADC output (F_{SIG} = 1GHz, F_{CLK} = 2.048GHz)

Figure 5.5 shows the dynamic performance of the ADC versus input amplitude. The input signal frequency and clock frequency remain at their previous values. The ADC shows a good linear relationship between input amplitude and SNDR for the entire valid-input range. A peak SNDR of 24.8dB is observed. As the input level increases above full-scale, the ADC saturates; the SNDR drops off relatively gradually for input levels below +2dBFS, although heavy clipping—as indicated by the sharp drop in SFDR and THD—occurs at levels beyond that. Throughout much of the valid input range, the

THD and SFDR is roughly 7 to 10dB below the ADC noise floor; at inputs beyond full-scale, the THD and SFDR suffer due the clipping effect mentioned above. At small inputs, THD and SFDR approach to within a few dB of the noise floor. This is most likely due to the fact that, at small inputs, very few ADC levels get toggled; as a result the quantization error no longer looks "white", but is instead heavy correlated to the input signal. This correlation between signal and quantization noise manifests itself as reduced THD, and it is a common effect in flash converters.

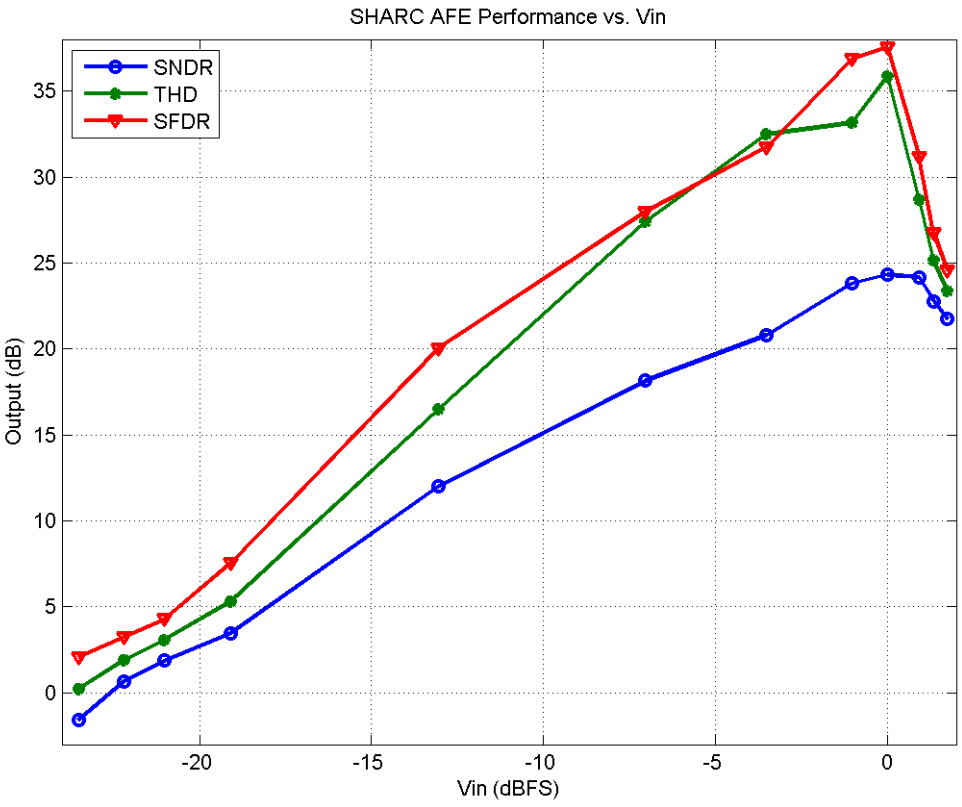

Figure 5.5: SNDR vs. input amplitude. (F_{SIG} = 1GHz. F_{CLK} = 2.048GHz)

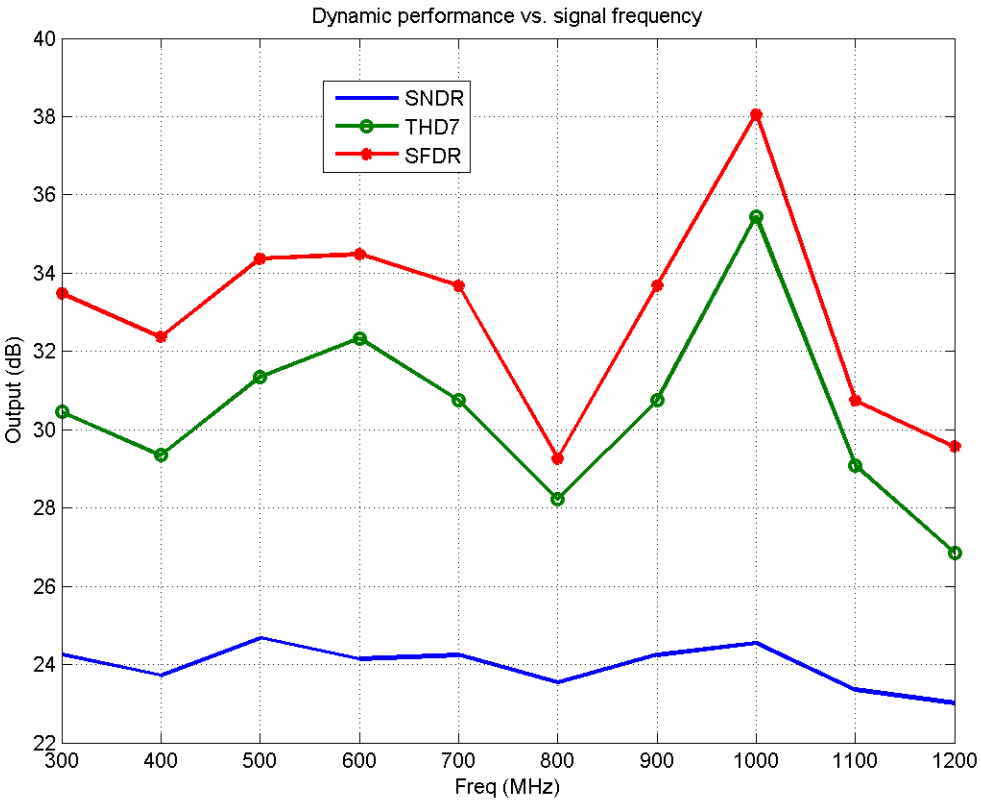

Figure 5.6: ADC dynamic performance vs. signal frequency. (F_{CLK} = 2.048GHz)

Figure 5.6 shows the dynamic performance of the ADC as the signal frequency is swept from 300MHz to 1200MHz. The clock frequency in this test is fixed at 2.048GHz. SNDR remains at roughly 24dB until about 1GHz, at which point it begins to slowly rolloff. THD and SFDR are typically above 30dB and 33dB respectively below the Nyquist frequency. Signal frequencies above 1200MHz were not tested, as the single-ended to differential converters used on the test PCB showed excessive attenuation at these frequencies.

Figure 5.7 shows the results of a two-tone test. The input frequencies were half-scale sinusoids at approximately 880MHz and 900MHz. All IM3 components were at least -44dBFS. An IM2 component at the equivalent frequency of $F_{sig,1}+ F_{sig,2}$ was roughly -40dBFS.

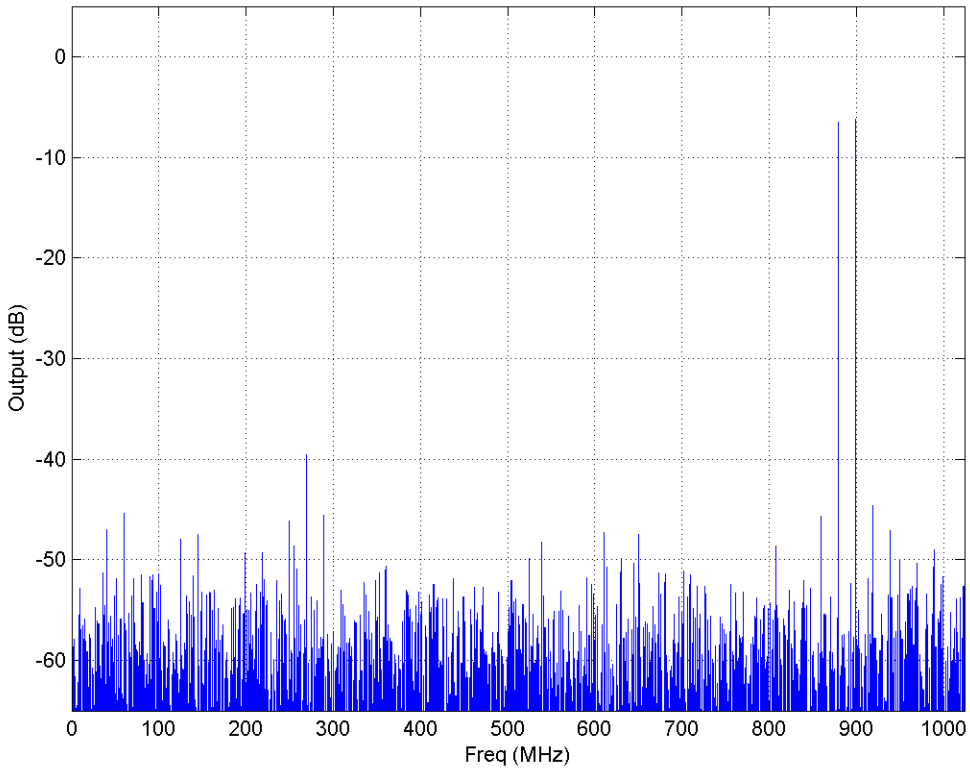

Figure 5.7: IMD performance of the ADC

5.2 Carrier Phase Rotator Performance

The role of the carrier phase rotator (CPR) block is to take an input quadrature signal and output a quadrature signal that has been phase-shifted or frequency-shifted by the desired amount. Therefore, the first metric to observe is the quadrature balance of the

CPR with no phase or frequency shift added. An input quadrature signal splitter at 1.014GHz was generated by using a Mini-Circuits 90° power-splitter [71]. The output of this block was applied to the I and Q chip inputs, and the ADC outputs were taken as the complex output, $ADC_I + j \cdot ADC_Q$. Figure 5.8 shows the complex FFT of the output; as expected, a tone at $+f_{SIG}$ is observed. Due to imperfect quadrature balance in the chip, the Mini-circuits quadrature generator and the other board components, an image tone at $-f_{SIG}$ is also present; however, this tone is 44dB below the desired tone, indicating quadrature matching better than 0.1dB and 0.7°. HD2 and HD3 tones are present at -33dB.

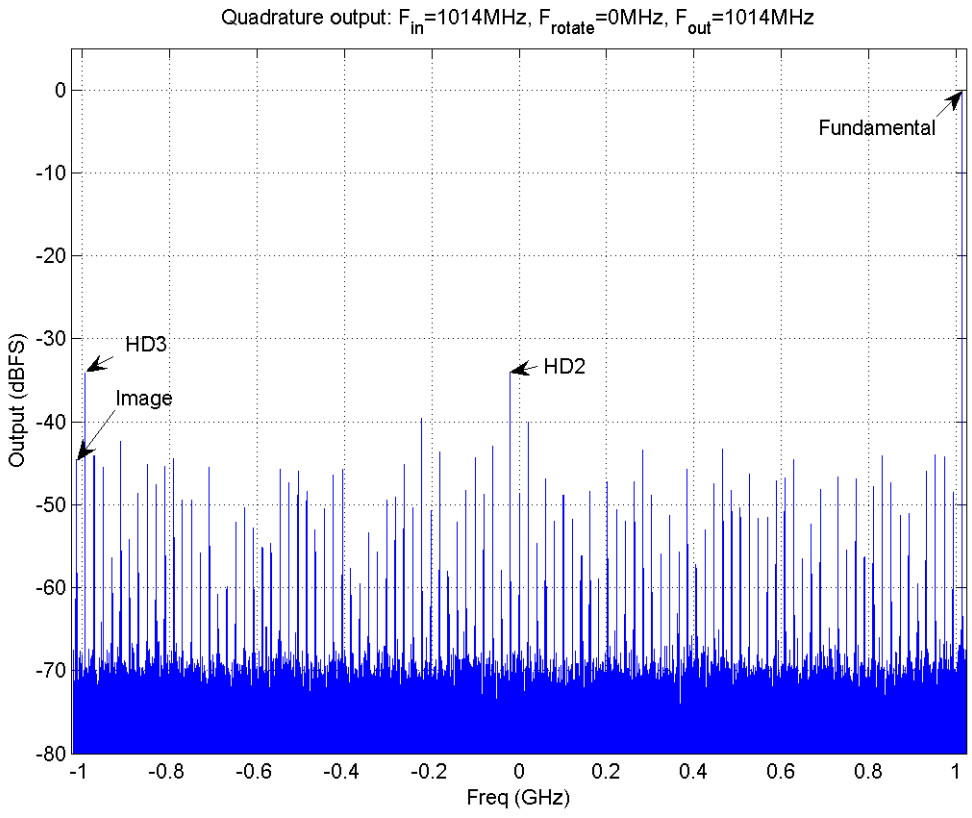

Figure 5.8: FFT of ADC I/Q output with no CPR frequency rotation

A second measurement was performed where the CPR was setup to rotate the same input signal by a -31MHz quadrature sinusoid. The results are shown in Figure 5.9. As expected, there is an output tone at 983MHz (1.014GHz + -31MHz). In this test, several other unwanted tones may be generated due to imperfect quadrature balance or other offsets in the CPR block. The tones, their descriptions, and the measured output levels are all summarized in Table 5.1.

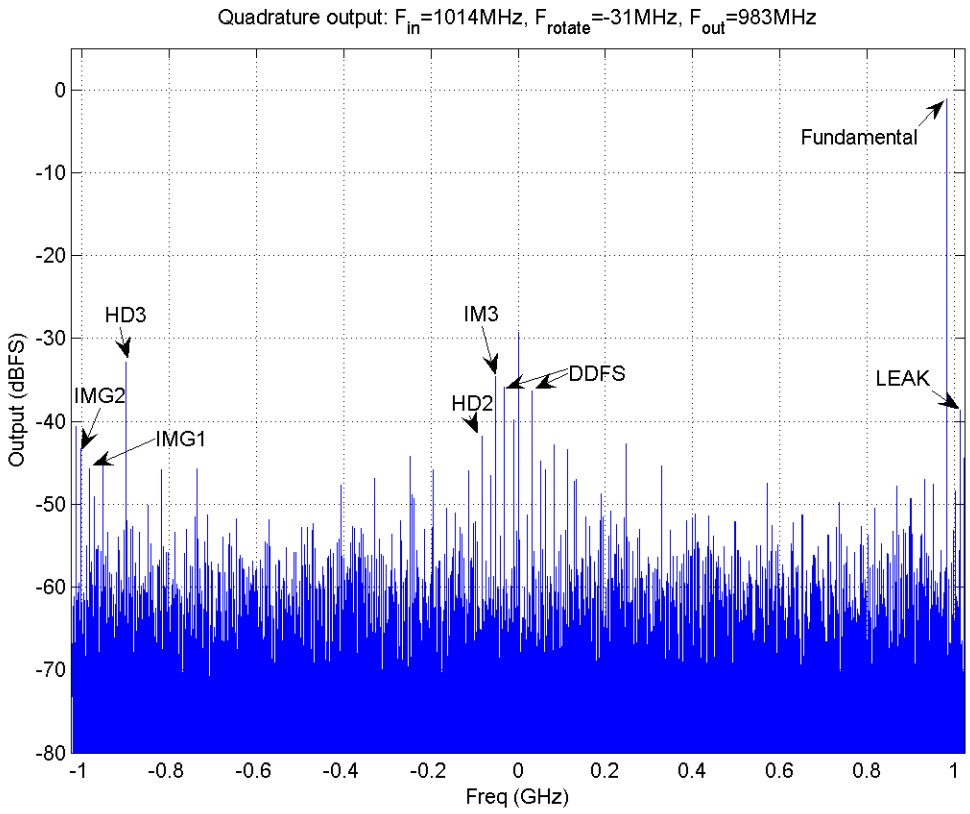

Figure 5.9: FFT of ADC I/Q output with CPR frequency rotation

Description	Notation	Frequency	Amplitude
Desired fundamental output tone	FUND	$F_{FUND}=F_{IN}+F_{CPR}$ = +983MHz	-1.1 dBFS
Image tone @ $-F_{FUND}$	IMG1	$F_{IMG1} = -F_{FUND}$ = -983MHz	-46.3 dBFS
Image tone @ $F_{IN} – F_{CPR}$	IMG2	$F_{IMG2} = F_{IN} – F_{CPR}$ = 1045MHz (Aliased to -1003MHz)	-43.3 dBFS
Leakage of input signal @ F_{IN}	LEAK	$F_{LEAK} = F_{IN}$ = 1014MHz	-38.6 dBFS
Tone at CPR rotation frequency	CPR	F_{CPR} = -31MHz	-36.4 dBFS
HD2 of output	HD2	$F_{HD3} = 3 \cdot F_{FUND}$ = 1966MHz (Aliased to -82MHz)	-40.6 dBFS
HD3 of output	HD3	$F_{HD2} = 2 \cdot F_{FUND}$ = 1996MHz (Aliased to +901MHz)	-31.7 dBFS
IMD3 of output	IMD3	$F_{IM3} = 2 \cdot F_{FUND} - F_{CPR}$ = 1997MHz (Aliased to -51MHz)	-34.5 dBFS

Table 5.1: Tones created by rotating 1.014GHz input signal by -31MHz in CPR block

5.3 Decision-Feedback Equalizer Performance

The accuracy of the DFE was measured by running various digital bit sequences into its input and observing its output. One challenge in this measurement is the fact that the DFE output is an analog quantity expressed in the current-domain and, as such, cannot be directly measured. Instead, it can only be indirectly observed by watching its effect on the ADC output; therefore, simple measurements of the DFE performance are limited by the 4-bit resolution of the ADC that follows it. Sub-LSB resolution can be obtained, however, by utilizing a scheme common to the serial link community known as “schmoo plots”. In generating schmoo plots, a typical serial link receiver with a very coarse quantizer (typically 1-bit or 2-bit) can measure to a much finer resolution by repeatedly changing the quantizer threshold levels while quantizing a repetitive sequence.

The SHARC chip utilized a similar technique to generate sub-LSB measurements of the DFE performance: the ADC had fixed threshold levels, but additional on-chip test circuitry allowed the injection of a high-resolution, programmable DC offset into the signal path; therefore, by quantizing the same repetitive DFE waveform with varying DC offsets, the DFE waveform can be observed with 7 – 8 bits of effective precision.

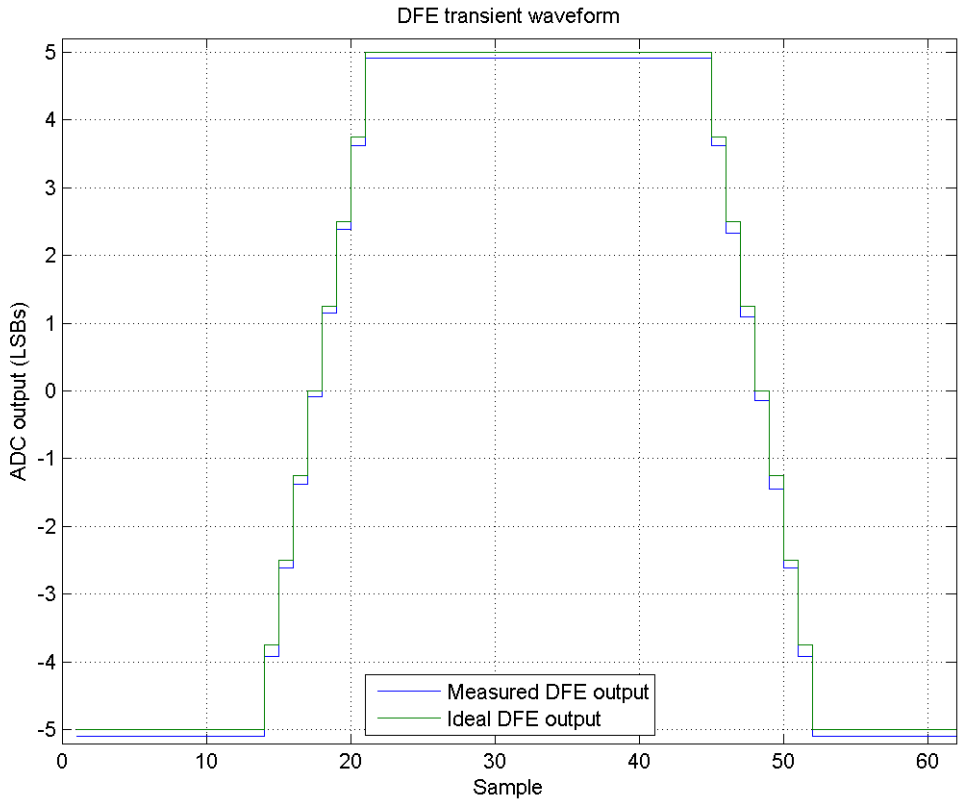

Figure 5.10: "Schmoo plot" of DFE waveform

A "schmoo plot" of an example DFE ramp waveform is shown in Figure 5.10. Both the ideal and measured values are shown. From this measurement, we can extract the effective INL and DNL of the DFE in order to get a sense of the current source matching present in this block. Figure 5.11 shows that the INL and DNL of the DFE current

sources are better than 0.1 LSBs; this is not surprising, as the channel length of the current sources used here are 2 μm.

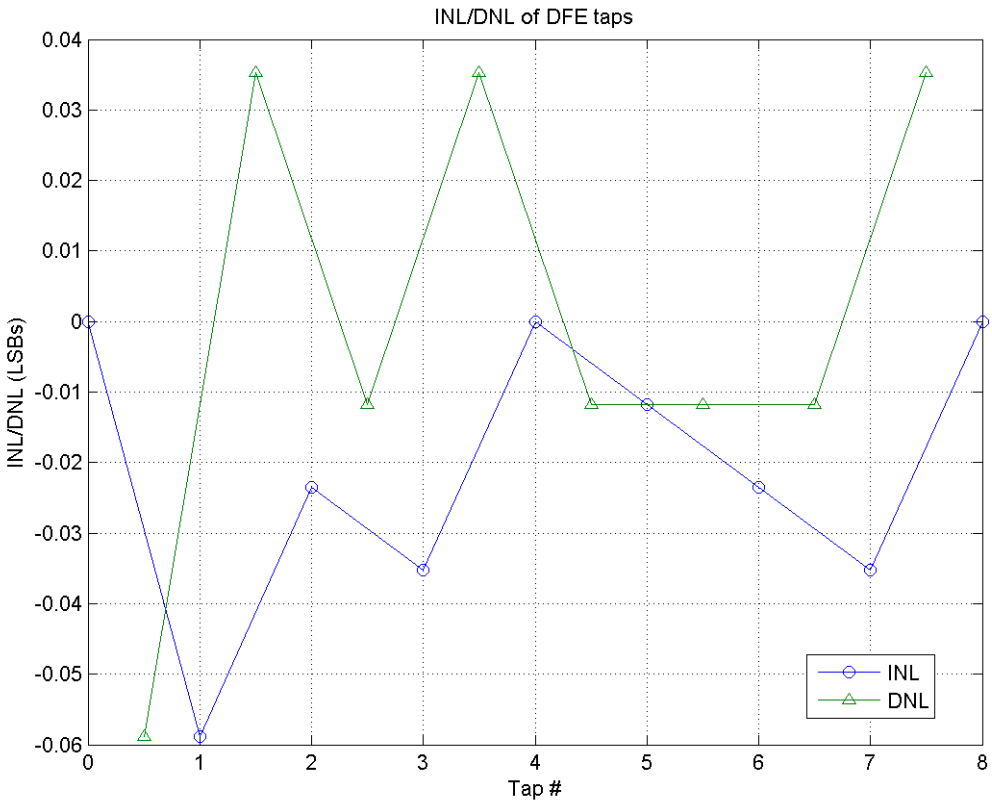

Figure 5.11: INL and DFE of DFE (Extracted from Figure 5.10)

For proper operation of the mixed-signal DFE utilized in the SHARC chip, it is important that the operation of the DFE not interfere with the performance of the ADCs. Two tests were performed to measure the impact of the DFE on ADC performance.

In the first test, the DFE taps were configured to be at their maximum value but with alternating polarity; therefore, the DFE would be actively switching during operation, but its net effect on the output signal should ideally be zero. Simultaneously, a -1dBFS sinusoid at 900MHz was injected into the primary chip inputs, and the dynamic

performance of the ADC was measured. This measurement was repeated with the DFE turned off. Figure 5.10 shows the FFTs from these measurements, with the DFE off and on, respectively. As can be observed from this figures, the dynamic performance of the main signal path is not significantly affected by the concurrent operation of the DFE. A table summarizing the measurements with and without the DFE running is presented in Table 5.2.

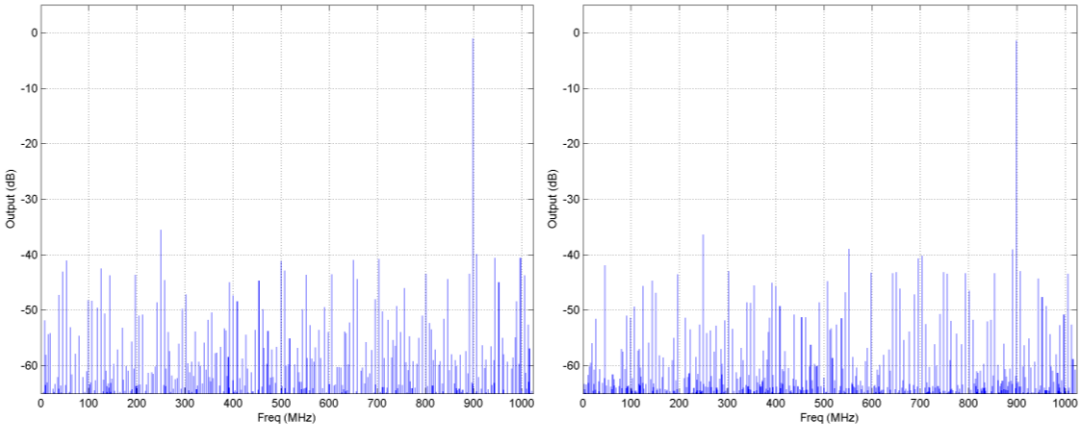

(a) **(b)**

Figure 5.12: FFT of ADC output with (a) DFE off and with (b) DFE on

Measurement	DFE off	DFE on
SNDR (dB)	23.8	23.6
THD (dB)	32.3	32.8
SFDR (dB)	34.5	35

Table 5.2: Dynamic performance metrics with DFE off and on

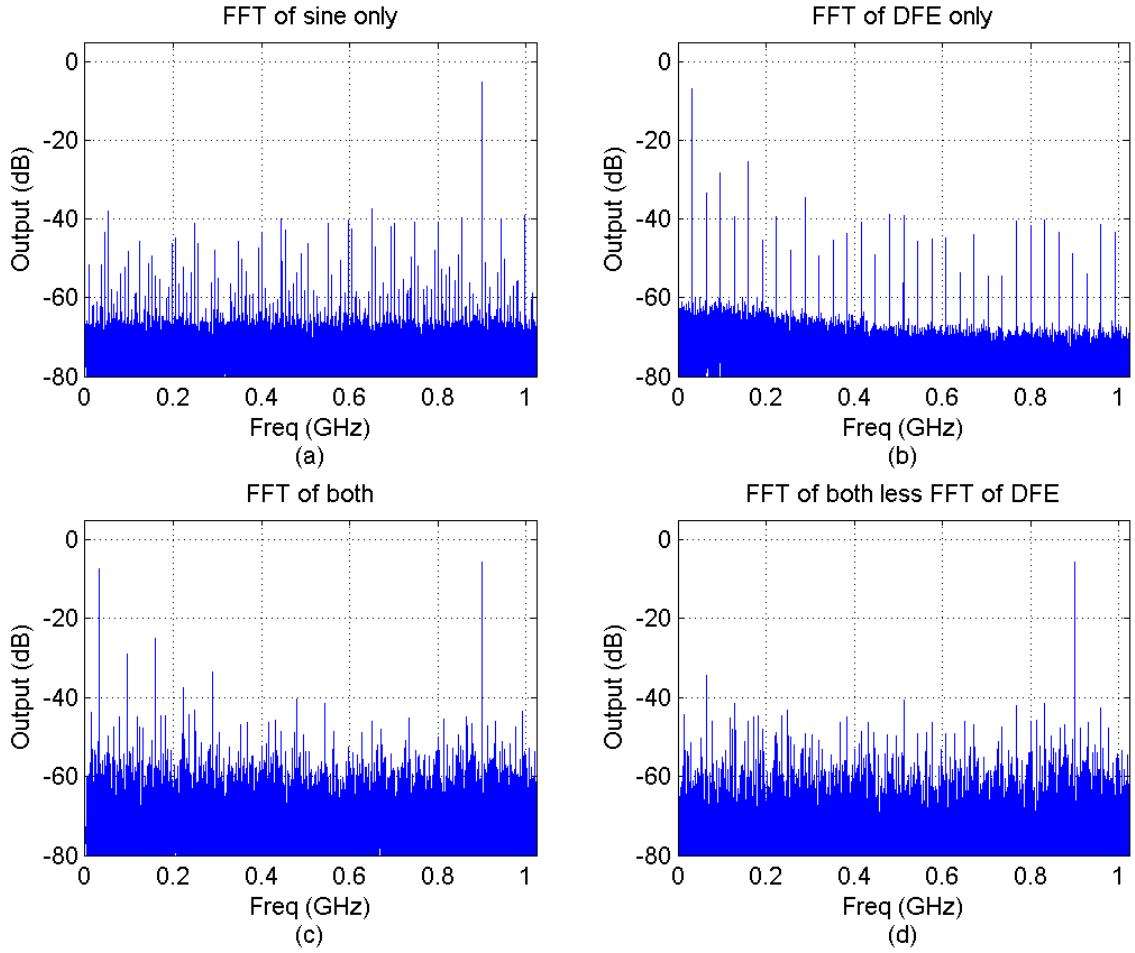

Figure 5.13: FFTs of various waveforms. (a) half-scale sine wave, (b) half-scale DFE pattern, (c) sine wave and DFE pattern, (c) sine wave and DFE pattern with DFE pattern digitally post-subtracted

A second test of the impact of DFE operation on ADC performance was performed. In this setup, a half-scale sinusoid was injected into the primary inputs while a half-scale DFE pattern was simultaneously excited. These two signals were added together at the receiver summing node and the resultant signal was quantized by the ADC. The signal amplitudes were intentionally chosen at half-scale so that they would not cause the ADC to clip when summed together. As shown in Figure 5.13 (a) and (b), the input sinusoid

had a frequency of approximately 900MHz, while the DFE pattern had a fundamental tone at roughly 70MHz. The FFT of the summed signal is shown in Figure 5.13 (c). In order to determine whether the presence of the DFE pattern degrades the ADC performance, the DFE pattern was digitally post-subtracted from the waveform of the quantized sum in Figure 5.13 (c) to create the FFT shown in Figure 5.13 (d). As can be seen in Figure 5.13 (d), the presence of the DFE waveform does not significantly degrade ADC performance. The SNDR of the original sine wave in Figure 5.13 (a) was 20dB; after digital post-subtraction of the DFE pattern, the SNDR is 19.7dB, a drop of only 0.3dB .

A final test of the DFE was performed to determine the DFE's capability to remove unwanted signal energy from the incoming signal. In typical operation, the role of the DFE is to remove the channel multipath energy from the incoming signal, leaving only the desired signal for subsequent demodulation and detection. Since the test setup could not create a 1GHz-bandwidth modulated multipath waveform for signal input, a related test had to be performed. The input signal consisted of two components: the desired signal, consisting of a -1.5dBFS sine wave at approximately 252MHz; and the unwanted signal, a 65% full-scale square-wave at approximately 256MHz. An FFT of just the desired signal is shown in Figure 5.14 (a) and an FFT of just the unwanted signal is shown in Figure 5.14 (b). With the DFE turned off, the composite incoming signal is larger than the fullscale range of the ADC, so heavy clipping results. Figure 5.15 (a) shows the ADC output in the time-domain, and clipping is evident. This in turn results in

significant intermodulation distortion; Figure 5.14(c) shows the FFT of the ADC output from Figure 5.15(a), and IMD of -19dBc is evident.

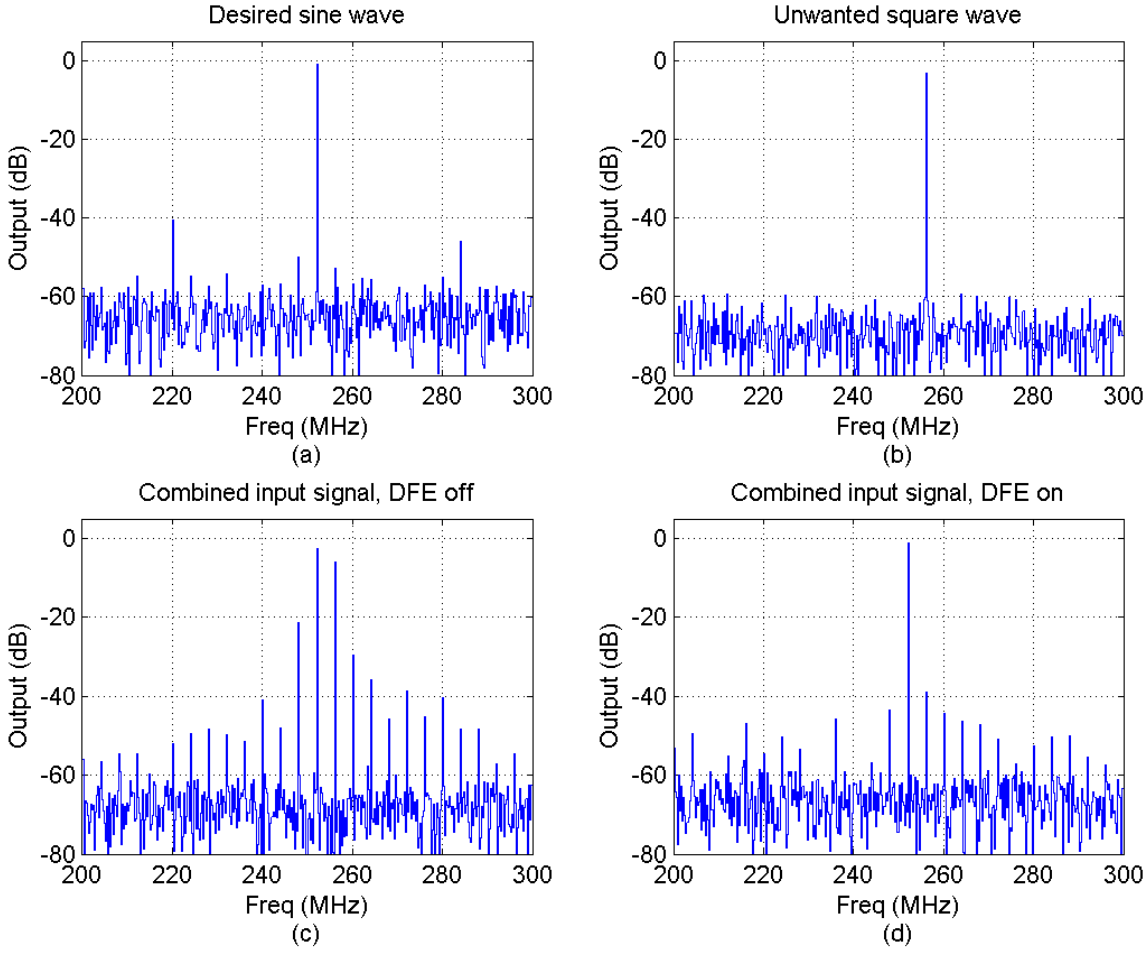

Figure 5.14: Demonstration of DFE's capability to subtract out unwanted signal. (a) Desired 250MHz sine wave, (b) Unwanted 250MHz square wave, (c) Combined input signal, DFE off, (d) Combined input signal, DFE on

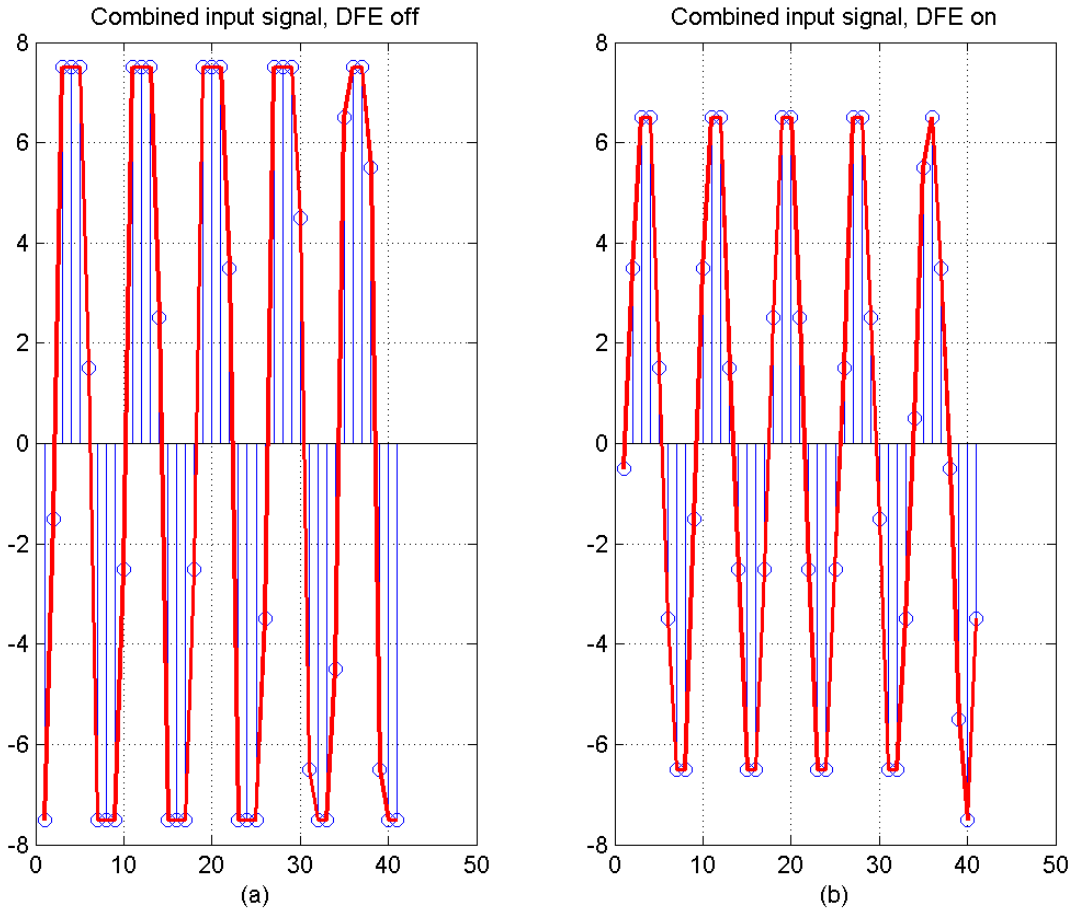

Figure 5.15: Transient waveforms of ADC output with combined input signal. (a) DFE off, (b) DFE on

When the DFE is properly configured and turned on, its output can precisely subtract out the incoming unwanted signal. Figure 5.15(b) shows a time-domain plot of the ADC output after the DFE is activated; clipping no longer occurs. An examination of the FFT of this output, shown in Figure 5.14 (d) shows that the DFE was able to suppress the unwanted signal by an additional 33dB; additionally, since the ADC no longer clips, IMD drop below the -42dBc level.

5.4 Conclusions

The measurement results have confirmed the performance capabilities of the SHARC receiver chip. The 4-bit 2Gs/s ADCs deliver an SNDR of 24.65dB at the Nyquist rate and THD and SFDR performance near the 6-bit level. The CPR block has demonstrated quadrature matching on the order of 0.1dB and 0.7° and is capable of frequency shifting a quadrature input signal, while generating image and leakage tones below the -37dBc level. The DFE is capable of producing high-accuracy current output waveforms to cancel incoming multipath signal. It has been shown that the operation of the DFE does not degrade ADC performance, and in fact the DFE effectively extends the dynamic range of the entire receiver architecture by as much as 7dB. The chip consumes roughly 55mW. A performance summary is shown in Table 5.3

Technology		90nm 9M1P Digital CMOS	
Package		Chip on board	
ADC	Sample Rate	2 Gs/s	
	Peak SNDR	24.8dB	
	SFDR	38dB	
	IMD3	-40dB	
CPR	DDFS clock speed	500MHz	
	DDFS resolution	8 bits (1.4°)	
	Image tones	< -37dB	
	I/Q matching	< 0.1dB, < 0.7°	
DFE	Clock speed	1GHz	
	INL/DNL	< 0.06 LSB	
	Linearity	>40dB	
Power	CPR	11mW	Total Power: 55mW
	DFE	14mW	
	THTIA	2 @ 5mW	
	ADC	2 @ 10mW	

Table 5.3: Chip measurement performance summary

6 Conclusion and Future Work

This research work has presents a mixed-signal baseband receiver architecture for 60GHz 1Gb/s wireless communications. The research conducted proposed both system and circuit solutions to minimize the power consumption and complexity of the receiver. At very high speeds, it was found that a combination of system optimization and proper partitioning between analog and digital signal processing can lead to significant power savings. In addition, several analog circuit techniques were introduced to enable high-speed operating at very low levels of current consumption.

This work represents just one step towards the ultimate goal of realizing a fully integrated CMOS 60GHz transceiver. Much of the microwave circuits work has already been completed by other researchers at the BWRC [9], [16], [72], [73], [74]. However, a major step that has yet to be completed is the integration of the microwave front-end with the baseband section of the receiver. Integrating the sensitive 60GHz circuitry on the same die as the "noisy" mixed-signal baseband while maintaining overall system performance will present a significant challenge. New techniques to maintain signal integrity and isolate the sensitive RF section from the baseband circuitry will need to be investigated.

Also, another obvious step towards an integrated transceiver would be the design of a fully integrated transmitter. This work focused on the design of the baseband receiver. To date at the BWRC, there has been work on some components within the transmitter, but little study of the entire transmitter portion of the system.

Another area of study that remains ripe for new research is the design and control of the beam forming antenna array. Prior work at the BWRC [9] has demonstrated a 4-element array, but high beamforming gain can be achieved by increasing the number of elements, perhaps to as many as 16. Two major challenges seem likely in this area of study: first, significant progress in packaging and interconnect design will be required to enable 16 antennae (and 16 LNA's and 16 PA's) to coexist on a single substrate without causing unwanted cross-coupling or interference. Second, further work is needed in the digital control algorithms in order to solve the beamforming alignment problem that arises with very narrow beamwidths.

Another natural extension of this work would be a deeper examination of the digital algorithms and architectures used in the transceiver. This future research could take one of a number of different directions. In this research, CMOS microwave performance limitations motivated the use of a constant envelope single-carrier system. However, there are digital techniques that could enable a multicarrier approach while potentially mitigating the circuit performance constraints of such a system. (Constant envelope OFDM, as proposed in [75] is an example of such a technique.) Further study of this and other approaches may make a multicarrier 60GHZ CMOS solution viable. Also, while the research presented in this work presented some digital algorithms for synchronization, the other algorithms required for the PHY and MAC layers were left unexamined. Implementing all these high-rate digital algorithms at power levels conducive to mobile applications is another potential area of research.

Bibliography

[1]. RNCOS. *Emerging Mobile Wireless Technologies -- Road to a Wireless World.* 2005.

[2]. WiMedia Alliance Home Page. [Online] http://www.wimedia.org/.

[3]. *Rules to permit use of radio Frequencies Above 40GHz for New Radio Applications.* FCC ET Docket No. 94-124. December 1995.

[4]. *Amendment of Part 2 of the Commission's Rules to Allocate Additional Spectrum to the Inter-Satellite, Fixed, and Mobile Services and to Permit Unlicensed Devices to Use Certain Segments in the 50.2-50.4GHz and 51.4-71.0 GHz Bands.* FCC ET Docket No.99-261. December 2000.

[5]. IEEE 802.15.3c WPAN Task Group. [Online] http://www.ieee802.org/15/pub/TG3c.html.

[6]. C. H. Doan, S. Emami, A. M. Niknejad, and R. W. Brodersen. "Design of CMOS for 60GHz applications". *ISSCC Digest of Technical Papers.* 2004, pp. 440-441.

[7]. C. H. Doan, S. Emami, D. A. Sobel, A. M. Niknejad, and R. W. Brodersen. "Design considerations for 60GHz CMOS radios". *IEEE Communications Magazine.* Dec. 2004, Vol. 42, 12, pp. 132-140.

[8]. Balanis, C. A. *Antenna Theory: Analysis and Design.* 2nd Edition. New York : John Wiley and Sons, Inc., 1997.

[9]. Alalusi, S. H. *"A 60GHz Adaptive Antenna Array in CMOS".* Ph.D. Dissertation, University of California at Berkeley. Berkeley, CA : s.n., 2005.

[10]. M. R. Williamson, G. E. Athanasiadou, and A. R. Nix. "Investigating the effects of antenna directivity on wireless indoor communication at 60GHz". *Proceedings of the PIMRC Conference.* 1997, pp. 635-639.

[11]. Williamson, M. *60GHz Measurement Program.* s.l. : Hewlettt Packard Laboratories Internal Report, Dec 1997.

[12]. H. Xu, V. Kukshya, and T. S. Rappaport. "Spatial and Temporal Characteristics of 60-GHz Indoor Channels". *IEEE Journal on Selected Areas in Communications.* Apr. 2002, Vol. 20, 3, pp. 620-630.

[13]. P. F. M. Smulders, G. Vervuurt. "Influence of Antenna Radiation Patterns on Mm-wave Indoor Radio Channels". *Second International Conference on Universal Personal Communications.* 1993, pp. 631-635.

[14]. S. Reynolds, B. Floyd, U. Pfeiffer, and T. Zwick. "60GHz Transceiver Circuits in SiGe Bipolar Technology". *ISSCC Digest of Technical Papers.* 2002, pp. 442-443.

[15]. K. Ohata, et al. "Wireless 1.25GB/s Transceiver Module at 60GHz-Band". *ISSCC Digest of Technical Papers.* 2002, pp. 236-237.

[16]. C. H. Doan, et al. "A Highly Integrated 60GHz CMOS Front-End Receiver". *ISSCC Digest of Technical Papers.* 2007.

[17]. A Babakhani, et al. "A 77GHz 4-Element Phased Array Receiver with On-Chip Dipole Antennas in Silicon". *ISSCC Digest of Technical Papers.* 2007.

[18]. L Nathawad, et al. "A Dual-Band CMOS MIMO Radio SoC for 802.11n Wireless LAN". *ISSCC Digest of Technical Papers.* 2008. To be published.

[19]. C. Lee, S-I. Liu. "A 58-60.4GHz Frequency Synthesizer in 90nm CMOS". *ISSCC Digest of Technical Papers.* 2007.

[20]. B. Floyd, et al. "A Silicon 60GHz Receiver and Transmitter Chipset for Broadband Communications". *ISSCC Digest of Technical Papers.* 2006.

[21]. A. Dec, K. Suyama, and T. Kitamura. "A 4.5GHz LC-VCO with Self-Regulating Technique". *ISSCC Digest of Technical Papers.* 2007.

[22]. D. Chowdhury, P. Reynaert, A. Niknejad. "A 60GHz 1V +12.3dBm Transformer-Coupled Wideband PA in 90nm CMOS". *ISSCC Digest of Technical Papers.* 2008. To be published.

[23]. A. Komijani, A. Hajimiri. "A Wideband 77GHz, 17.5dBm Power Amplifier in Silicon". *Proceedings of the Custom Integrated Circuits Conference.* 2005, pp. 571-574.

[24]. J. Kang, A. Hajimiri, B. Kim. "A Single-Chip Linear CMOS Power Amplifier for 2.4GHz WLAN". *ISSCC Digest of Technical Papers.* 2006.

[25]. Gary, Bruce. [Online] http://Brucegary.net/MTP_tutorial/MTP_ch5.html.

[26]. L.M. Correia, J. J. Reis, P. O. Frances. "Analysis of the average power to distance decay rate at the 60 GHz band". *Proceedings of the 47th Vehicular Technology Conference.* 1997, Vol. 2, pp. 994-998.

[27]. C. R. Anderson, T. S. Rappaport. "In-Building Wideband Partition Loss Measurements at 2.5 and 60GHz". *IEEE Transactions on Wireless Communications.* May 2004, Vol. 3, 3, pp. 922-928.

[28]. R. Davies, et al. "Wireless Propagation Measurements in Indoor Multipath Environments at 1.7GHz and 60GHz for Small Cell Systems". *Proceedings of the Vehicular Technology Conference.* 1991, pp. 589-593.

[29]. G. Durgin, T. S. Rappaport, and H. Xu. "MEasurements and Models for Radio Path Loss and Penetration Loss In and Around Homes and Trees at 5.85GHz". *IEEE Transactions on Communications.* Nov. 1998, Vol. 46, 11, pp. 1484-1496.

[30]. L. M. Correia, P. O. Frances. "Transmission and Isolation of Signals in Buildings at 60 GHz". *Proceedings of PIMRC.* 1995, pp. 1031-1034.

[31]. R. Davies, et al. "Wireless propagation measurements in indoor multipath environments at 1.7GHz and 60GHz for small cell systems". *Proceedings of the IEEE 41st Vehicular Technology Conference.* 1991, pp. 589-593.

[32]. Smulders, P. "Exploiting the 60 GHz Band for Local Wireless Multimedia Access: Prospects and Future Directions". *IEEE Communications Magazine.* Jan. 2002, pp. 140-147.

[33]. N. Guo, R. C. Qiu, et al. 60-GHz Millimeter-Wave: Radio: Principle, Technoloy, and New Results. *EURASIP Jounal on Wireless Communications and Network.* 2007, Vol. 2007. Article ID 68253, 8 pages.

[34]. H. Yang, et al. Channel Characteristics and Transmission Performance for Various Channel Configurations at 60 GHz. *EURASIP Journal on Wireless Communications and Networking.* 2007, Vol. 2007. Article ID 19613, 15pages.

[35]. A. S. Y. Poon, D. N. C. Tse, and R. W. Brodersen. Multiple-antenna channels from a combined physical and networking perspective. *Conference Record of the Thirty-Sixth Asilomar Conference on Signals, Systems and Computers.* Nov. 2002, Vol. 2, pp. 1528-1532.

[36]. Proakis, J. G. *Digital Communications.* 3rd Edition. Boston : McGraw-Hill, 1995.

[37]. B. Langen, G. Lober, W. Herzig. "Reflection and Transmission Behaviour of Building Materials at 60 GHz". *Proceedings of PIMRC.* 1994.

[38]. P. F. M. Smulders, L. M. Correia. Characterisation of propagation in 60 GHz radio channels. *Electronics and Communication Engineering Journal.* April 1997, pp. 73-80.

[39]. G. E. Athanasiadou, A. R. Nix, and J. P. McGeehan. Indoor 3D ray tracing predictions and their comparison with high resolution wideband measurements. *Proceedings IEEE Vehicular Technology Conference.* 1996.

[40]. D. Dadari, L. Minelli, V. Tralli and O. Andrisano. "Wideband Indoor Communications Channels at 60GHz". *Proceedings of PIMRC.* 1996, Vol. 3, pp. 791-794.

[41]. H. Droste, G. Kadel. "Measurement and Analysis of Wideband Indoor Propagation Characteristics at 17 GHz and 60 GHz". *Ninth International Conference on Antennas and Propagation.* 1995, pp. 288-291.

[42]. Kajiwara, A. "Effects of Polarization, Antenna Directivity, and Room Size on Delay Spread in LOS Indoor Radio Channel". *IEEE Transactions on Vehicular Technology.* Feb 1997, Vol. 46, 1, pp. 169-175.

[43]. T. Manabe, Y. Miura, and T. Ihara. "Effects of Antenna Directivity and Polarization on Indoor Multipath Propagation Characteristics at 60 GHz". *IEEE Journal on Selected Areas in Communications.* Apr. 1996, Vol. 14, 3, pp. 441-448.

[44]. P. F. M. Smulders, A. G. Wagemans. "Frequnecy-Domain Measurement of the Millimeter Wave Indoor Radio Channel". *IEEE Transactions on Instrumentation and Measurement.* Dec. 1995, Vol. 44, 6, pp. 1017-1022.

[45]. P. F. Smulders, A. G. Wagemans. "Wide-band Measurements of mm-Wave Indoor Radio Channels". *Proceedings of PIMRC Conference.* 1992, pp. 329-333.

[46]. D. Dardari, et al. "Wideband Indoor Communication Channels at 60GHz". *Proceedings of the PIMRC Conference.* 1996, pp. 791-794.

[47]. Saleh, A. A. M. "Frequency-Independent and Frequency-Dependent Nonlinear Models of TWT Amplifiers". *IEEE Trans. on Comm.* Nov. 1981, Vols. COM-29, 11, pp. 1715-1720.

[48]. Rapp, C. "Effects of HPA-nonlinearity on 4-DPSK-OFDM-signal for a digital sound broadcasting system". *Proceedings of the 2nd European Conference on Satellite Communications.* Oct. 1991, pp. 179-184.

[49]. Crawford, J. A. *Frequency Synthesizer Design Handbook.* Norwood, MA : Artech House, 1994.

[50]. D'Andrea, U. Mengali and A. *Synchronization Techniques for Digital Receivers.* New York : Plenum Press, 1997.

[51]. D. Sun, A. Xotta, A. A. Abidi. "A 1 GZ CMOS analog front-end for a generalized PRML read channel". *IEEE Journal of Solid-State Circuits.* Nov. 2005, Vol. 40, 11, pp. 2275-2285.

[52]. J. L. Zerbe, et al. "Equalization and clock recovery for a 2.5-10-Gb/s 2-PAM/4-PAM backplane transceiver cell". *IEEE Journal of Solid-State Circuits.* Dec. 2003, Vol. 38, 12, pp. 2121-2130.

[53]. P. M. Figueiredo, et al. "A 90nm CMOS 1.2V 6b 1Gs/s Two-Step Subranging ADC". *ISSCC Digest of Technical Papers.* 2006, pp. 568-569.

[54]. Z. Cao, S. Yan, Y. Li. "A 32mW 1.25GS/s 6b 2b/step SAR ADC in 0.13um CMOS". *ISSCC Digest of Technical Papers.* 2007.

[55]. K. Deguchi, et al. "A 6-bit 3.5-Gs/s 0.9V 98-mW Flash ADC in 90nm CMOS". *2007 IEEE Symposium on VLSI Circuits.* 2007, pp. 64-65.

[56]. Ammer, M. J. *Low Power Synchronization for Wireless Communication.* Berkeley : UC-Berkeley Ph.D. dissertation, 2004.

[57]. H. Meyr, M. Moeneclaey, S. A. Fechtel. *Digital Communications Receivers.* New York : Wiley, 1998.

[58]. Mathworks Inc. *Matlab SIMULINK.*

[59]. Yee, D. *A Design Methodology for Highly Integrated Low-Power Receivers for Wireless Communications.* Berkeley : UC-Berkeley Ph.D dissertation, 2001.

[60]. P. J. Sullivan, et al. "Low voltage performance of a microwave CMOS Gilbert cell mixer". *IEEE Journal of Solid-State Circuits.* July 1997, Vol. 32, 7, pp. 1151-1155.

[61]. Gilbert, B. "A High-Performance Monolithc Multiplier Using Active Feedback". *IEEE Journal of Solid-State Circuits.* Dec. 1974, Vol. 9, 6, pp. 364-373.

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

[63]. T. Kwan, K. Martin. "An Adaptive Analog Continuous-time CMOS Biquadratic filter". *IEEE Journal of Solid-State Circuits.* 1991, Vol. 26, 6, pp. 859-867.

[64]. Samueli, H.T. Nicholas and H. "A 150-MHz direct digital frequency synthesizer in 1.25um CMOS with -90dBc spurious performance". *IEEE Journal of Solid-State Circuits.* Dec. 1991, Vol. 26, 12, pp. 1959-1969.

[65]. B. D. Yang, et al. "A direct digital frequency synthesizer using a new ROM compression method". *Proc of 27th European Solid-State Circuits Conference.* Sept. 2001, pp. 273-276.

[66]. Lee, S. Mortezapour and E. K. F. "Design of low-power ROM-less direct digital frequency synthesizer using nonlinear digital-to-analog converter". *IEEE Journal of Solid State Circuits.* Oct. 1999, Vol. 34, 10, pp. 1350-1359.

[67]. Razavi, B. *Principles of Data Conversion System Design.* Piscataway : IEEE Press, 1995.

[68]. Plassche, R. van de. *CMOS Integrated Analog-to-Digital and Digital-to-Analog Converters, 2nd ed.* Boston : Kluwer Academic Publishers, 2003.

[69]. Barrow, K. Kattmann and J. "A technique for reducing differential nonlinearity errors in flash A/D converters". *ISSCC Digest of Technical Papers.* Feb. 1991, pp. 170-171.

[70]. Mangelsdorf, C. W. "A 400-MHz INput Flash Converter with Error Correction". *IEEE Journal of Solid-State Circuits.* Feb. 1990, Vol. 25, 1, pp. 184-191.

[71]. [Online] Datasheets available at. http://www.minicircuits.com.

[72]. B. Heydari, M. Bohsali, E. Adabi, A. M. Niknejad. "Low-Power mm-Wave components up to 104GHz in 90nm CMOS". *ISSCC Digest of Technical Papers.* 2007, pp. 200-201.

[73]. B. Afshar, Y. Wang, A. M. Niknejad. "A Robust 24mW 60GHz Receiver in 90nm Standard CMOS". *ISSCC Digest of Technical Papers.* 2008.

[74]. E. Adabi, A. M. Niknejad. "Broadband Variable Passive Delay Elements Based on an Inductance Multiplication Technique". *RFIC Digest of Papers.* 2008.

[75]. M. Kirivanta, et al. "Constant envelope multicarrier modulation: Performance Evaluation in AWGN and fading channels". *Proceedings of IEEE Military Communications Conference.* 2005, pp. 807-813.