Wideband Signal Acquisition via Frequency-Interleaved Sampling

Steven Callender

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2015-224
http://www.eecs.berkeley.edu/Pubs/TechRpts/2015/EECS-2015-224.html

December 1, 2015

Wideband Signal Acquisition via Frequency-Interleaved Sampling

by

Steven William Callender

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

in

Electrical Engineering and Computer Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Ali Niknejad, Chair
Professor Borivoje Nikolić
Professor Martin White

Spring 2015

Wideband Signal Acquisition via Frequency-Interleaved Sampling

Abstract

Wideband Signal Acquisition via
Frequency-Interleaved Sampling

by

Steven William Callender

Doctor of Philosophy in Electrical Engineering and Computer Science

University of California, Berkeley

Professor Ali Niknejad, Chair

High-speed analog-to-digital converters (ADCs) are key enabling blocks for emerging wideband applications in communication, high-end instrumentation, and medical imaging. Larger signaling bandwidths improves system performance, and necessitates high-speed ADCs for accurate digitization. As an example, current state-of-the-art oscilloscopes have an acquisition bandwidth exceeding 60GHz with effective sample rates greater than 100GS/s. This places significant difficulty in the design of sample-and-hold (S/H) and analog-to-digital conversion circuitry that can operate at such high speeds while providing moderate resolution. As a result, the front-end of these systems are often complex, multi-chip solutions that are fabricated in expensive processes such as indium-phosphide (InP). With the increased demand for battery-operable, low-power systems it is desirable to have these high-performance signal acquisition systems in a fully-integrated CMOS implementation in order to harness the power of scaling as dictated by Moore's law. To achieve this, several advancements on current data conversion techniques need to be made.

In this thesis, we explore the design and optimization of a frequency-interleaved ADC (FI-ADC) as an alternative to conventional high-speed ADC architectures, which are often heavily time-interleaved. Due to the large interleaving factor and timing sensitivity, the conventional architectures are often very power hungry and offer typical resolutions of 4 bits or less. FI-ADCs, in which the input signal is divided into various frequency bands which are independently digitized and digitally recombined, show less susceptibility to jitter, the primary bottleneck in high-speed ADCs. System simulations have shown a potential improvement in SNR performance for a frequency-interleaved ADC versus a direct sampling, time-interleaved architecture.

The focus of this thesis is to provide a fundamental understanding of the operation of the FI-ADC and investigate the similarities and differences to the conventional time-interleaved ADC in respect to design complexity, design challenges and overall performance.

To my loving family.
For always being in my corner, cheering me on.

Contents

List of Figures

List of Tables

Acknowledgments

This research would not have been possible without the guidance of my advisor, Ali Niknejad. Throughout my seven years at Berkeley, Ali has provided sound advice and encouragement that helped me develop both professionally and personally. There is no need to comment on his brilliance as a researcher and lecturer – this is already well-known. Instead, I would like to thank him for his patience as an advisor. Never once did I feel pressured to push out a publication or prematurely present my results. Ali insisted on digging deeper to make sure a fundamental understanding was obtained and to get to the root of the issue. One of my most testing moments in graduate school occured in 2013 when I missed the first tapeout of the FI-ADC chip after working tirelessly (14+ hour days) for 3 consecutive months. Feeling defeated and a bit embarassed, I had a meeting with Ali on the submission date to relay the bad news. Instead of responding with a tone of disappointment, Ali spoke words of encouragement that helped lift my spirits and get me back on my feet.

I would also like to thank my dissertation committee, Bora Nikolić, Elad Alon, and Martin White for overseeing my research and providing constructive feedback on points that needed strengthening. I would like to further thank Bora for the phone call he made on his way home from the office, back when I was a new admit, aiming to convince me to come to Berkeley. Berkeley was already my top choice, but it was good to hear a voice expressing Berkeley's interest in my talents. And Elad, for strengthening my circuits knowledge in the classroom and always having an open door when I needed to discuss research.

The FI-ADC chip presented in this thesis would not have been possible if it weren't for the help and support of Lucas Calderin and Greg LaCaille. As first-year graduate students with little-to-no tapeout experience, they joined the project and took on the resposibility of designing the LO generation and distribution for a tapeout that was five months away. It was quite impressive to see how swiftly they caught up to speed and how they quickly began contributing to design decisions in the overall FI-ADC. It was a pleasure working alongside them and making the late night trips to Safeway during tapeout.

I would like to thank SRC and the TxACE program for funding the FI-ADC project. I would also like to thank Robert Noyce Fellowship for funding part of my graduate studies, Texas Instruments for generously donating the ADCs that were used for testing the FI-ADC front-end, Keysight, Integrand and Ansoft for software donations. Thanks are also due to the BWRC industrial sponsors who have provided valuable feedback on this project at our semi-annual research retreats.

I would like to acknowledge the first group I ever worked with at Berkeley, the TUSI team: Amin Arbabian, Shinwon Kang, Jun-Chau Chien, Ali (Bagher) Afshar and Ehsan Adabi. It was an honor and privelege to work alongside such gifted and hard-working researchers. Working on the TUSI project, I learned how to effectively work in a group and collectively tackle a large problem. Amin, the headmaster of it all, unknowingly played the role of co-advisor in my formative years at Berkeley. He is truly a unique talent and one of the brightest individuals I have ever worked with.

I have also had the privelege to work alongside the talented individuals of Ali Niknejad's research group (past and present) : Siva Thyagarajan, Jiashu Chen, Maryam (Sahar) Tabesh, Ashkan Borna, Sriramkumar Venugopalan, Pramod Murali, Lu Ye, Cristian Marcu, Andrew Townley, Nai-Chung Kuo, Nima Baniasadi, and Paul Swihurn. These individuals were an invaluable resource throughout my graduate years, providing informative discussions and debates on various research topics and circuits questions.

Outisde of our research group, I have made many friends and colleagues at BWRC: Milos Jorgovanović, Katerina Papadopoulou, William Biederman, Dan Yeager, Matt Weiner, Matt Spencer, Rikky Muller, Simone Gambini, Yida Duan, Yue Lu, Charles Wu, Chintan Thakkar, Lingkai Kong, John Crossley, Angie Wang, Rachel Hochman, Amanda Pratt, and Sharon Xiao. I would like to thank them all for their friendsip and camaraderie and for making BWRC one of the best research centers to work in.

In addition, many thanks are due to the BWRC staff (past and present): Brian Richards, Fred Burghardt, Deirdre McAuliffe-Bauer, Gary Kelson, Leslie Nishiyama, Olivia Nolan, Sarah Jordan, Bira Coelho, Susan Mellers and Tom Boot. These individuals have helped keep the BWRC running smoothly over the years and tended to all research needs of the students.

I am fortunate to have also formed lasting friendships with EECS colleagues outside of BWRC: Michael Lorek, Henry Barrow, Bobby Schneider, Richard Przybyla and Mitchell Kline. We've shared many beers together in order to shed away the stresses of grad school which really helped me keep my sanity throughout the years.

I would also like to thank the EECS department and staff, namely, Shirely Salanio, Sheila Humphreys, Patrick Hernan, and Ruth Gjerde for making sure I properly progressed through my graduate studies.

Before Berkeley, there was Columbia. I would like to thank my friends who have been with me on this journey since undergrad: Brian Ramos, Carlene Liriano, Ed Liao, Josh Breslow, Jonathan Barbee, Keith Dronson, Frank Zovko, Jessica Lin, Dan Llamas, Emmanuel-Sean Peters, and David Diaz. I would also like to give a special thanks to my undergraduate research advisor at Columbia, Professor Yannis Tsividis, who took me under his wing as a bright-eyed freshman and nourished the beginnings of my career in research.

Last, but certainly not least, I would like to thank my entire family. I would not be where I am today if it were not for the love and support of my mother, Marcia Simms. She has been my number one supporter since day one. Words cannot express how indebted I am to her. I am also indebted to my father, Alfred (Billy) Callender for providing constant words of wisdom throughout the years. He has always emphasized the value of education and work ethic from when I was young. I would like to thank him for never settling for my 98% test scores and always questioning (half-jokingly) why I did not get a 100%. My #2 fans would have to be my beautiful sisters, Stephanie Callender and Raynelle Callender-McKissick. Thank you for always looking after your "little bro" and showing me unconditional love. And to my large, beautiful, and loving family – aunties, uncles, cousins, and in-laws – thank you for everything you have given me. You all are the village that raised me.

Chapter 1

Introduction

There is an ever-increasing demand for higher bandwidth systems. Modern day high-speed serial links, oscilloscopes, and pulsed mm-wave imagers have illustrated a need for analog-to-digital converters with sample rates exceeding 50GS/s and resolutions greater than 6 bits. These are extremely difficult specifications to meet using current approaches and the systems that do meet these specifications are often wall-powered and done in non-CMOS processes. In this thesis, we explore an alternative approach to high-speed analog-to-digital conversion that can provide improved resolutions at very high input frequencies with low power consumption in current-day CMOS technologies.

1.1 Motivation

High-speed analog-to-digital converters (ADCs) have traditionally been utilized in very niche applications (e.g. high-end oscilloscopes). With the increased speed offered by deeply scaled CMOS processes and the increased demand for communication data rates, high-speed ADCs have found more frequent usage in both traditional and emerging applications. The required specifications of these high-speed ADCs depends on application. Table 1.1 shows sample numbers for the required sample rate (SR) and effective number of bits (ENOB) of ADCs used in various applications.

For cellular LTE technology, the maximum sample rate is $200-250$MS/s which is considered high-speed for common consumer applications. The resolution of these ADCs is quite high, which is why $\Delta\Sigma$ ADCs are very popular for cellular receivers [1]. In contrast, mm-wave imaging for medical applicatons, an emerging application, requires sample rates two orders of magnitude higher than what is required for cellular applications but scales back on the necessary resolution. This application is of particular interest since it is the root project from which this work developed.

As discussed in [2], pulsed mm-wave imagers can be utilized in a myriad of applications such as medical diagnosis and gesture recognition. For these applications, transmission of narrow pulses is desirable since there is an inverse relationship between the pulse width and

the depth resolution of the imager. In the case of medical diagnosis,

	Cellular	**HS Links**	**HS OScopes**	**UWB Imaging (TUSI)**
F_s [GHz]	0.25	5-20	160	20-50
ENOB [bits]	9-10	3-6	8	6

Table 1.1: ADC specifications for various high-speed applications.

narrower pulses improve depth resolution and can lead to earlier detection of cancerous cells. In [3] and [4], 94GHz transceivers were designed for usage in a Time-domain Ultra-wideband Silicon Imager (TUSI) system. [3] demonstrated a 94GHz transmitter that was capable of transmitting record pulse widths down to 26ps. In [4], a similar transmitter was paired with a receiver to demonstrate the transmission and detection of pulses down to 30ps. Using a quadrature receiver, this translates to baseband bandwidths on the order of 20GHz. The receiver in that design was composed of a wieband LNA and downconverting micro-mixers, omitting the baseband digital processing. The reason for this was the inexistence of ADC techniques that could meet the required the specifications ($20 - 50$GS/s, $f_{in,max} > 25$GHz, and 6b ENOB). This is still the case as can be seen from Figure 1.1 which plots the ENOB of published ADCs as a function of the maximum input frequency.

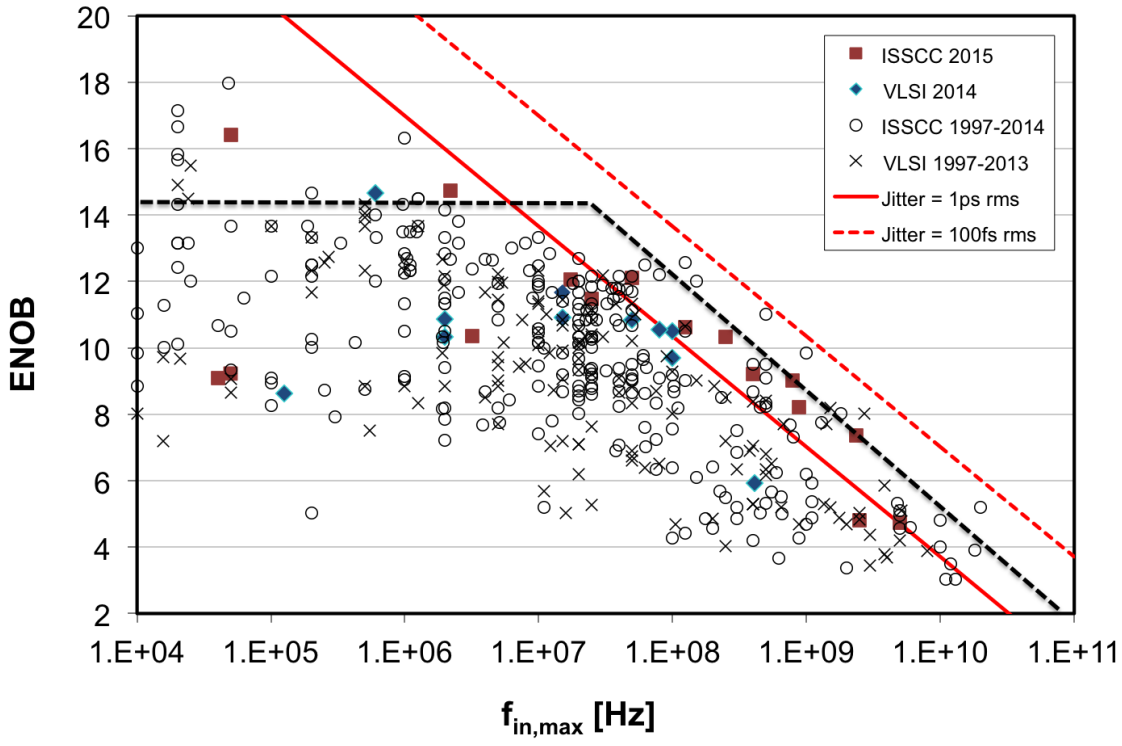

Figure 1.1: ENOB vs. $f_{in,max}$ for published ADCs.

From the above plot, it is seen that the ENOB begins to continuously decrease as the input frequency is increased above 100MS/s. This region is referred to here as the *jitter-limited regime*. In this region, the attainable ENOB of the ADC is no longer determined by the quantization noise, but is instead limited by the aperture jitter - timing errors on the sampling clock lead to voltage errors in the digital representation. Because of this issue, the majority of ADCs sampling above 20GS/s have ENOBs less than 5b. As can be seen in the plot, in order to have a chance at obtaining ENOBs greater than 6b for frequencies above 20GHz, the sampling clock must have an rms-jitter less than 100fs-rms which is a challenging task. As a result, it is worth exploring techniques that can potentially break the jitter-barrier faced by current ADC architectures.

Why Frequency-Interleaved?

The frequency-interleaved ADC (FI-ADC) is an alternative approach to high-speed data conversion [5]. It has the attractive quality that it is less sensitive to sampling jitter and can potentially improve the attainable resolution of high-speed ADCs given the same amount of rms-jitter. It accomplishes this by channelizing the input into various frequency sub-bands thereby limiting the maximum signal frequency presented to the sampler. Although there are other cited benefits of the FI-ADC [5], the decreased sensitivity to jitter is the most important feature when it comes to high-speed data conversion. As such, this thesis aims to provide a proper understanding of how the FI-ADC can (potentially) achieve higher resolution over the conventional architectures. This thesis also aims to help develop insight into how to properly design an FI-ADC system in order to reduce system complexity and power consumption.

1.2 Thesis Overview

This thesis is organized as follows:

Chapter 2 provides a background on the dominant high-speed ADC architecture, the time-interleaved ADC (TI-ADC). Operation of the TI-ADC is first discussed, followed by design challenges and performance limitations. The FI-ADC, as proposed by [5], is then introduced followed by a review of several FI-ADC designs.

Chapter 3 presents the proposed FI-ADC architecture that is used in this project. A discussion of the key design challenges such as wideband signal distribution, harmonic folding, LO generation and digital reconstruction is provided. Next, a detailed comparison between the FI-ADC and TI-ADC is presented. The heart of this section is found in the discussion of the impact of LO phase noise on FI-ADC performance. Until now, there as been little to no discussion on how LO phase noise impacts the overall performance of FI-ADCs. This chapter concludes with a summarization of system-level simulations which were performed to compare the two architectures.

Chapter 4 provides a guideline to the design of a 50GS/s 6-bit FI-ADC. System-level design considerations are discussed. The design of the FI-ADC analog front-end (AFE) which was taped out and measured is then presented. Chapter 5 reports the measurement results of the taped out chip.

Chapter 6 summarizes and concludes the thesis.

Chapter 2

Background

2.1 High-Speed ADCs - The Time-Interleaved ADC

The time-interleaved (TI) ADC [6] is the most commonly used architecture for high-speed ADCs, especially for designs aiming to achieve the highest sample rate in a given process. This can be seen in Figure 2.1 which plots the Walden figure-of-merit, FOM_W, versus the Nyquist sampling rate for all ADCs published in ISSCC and VLSI [7][8].

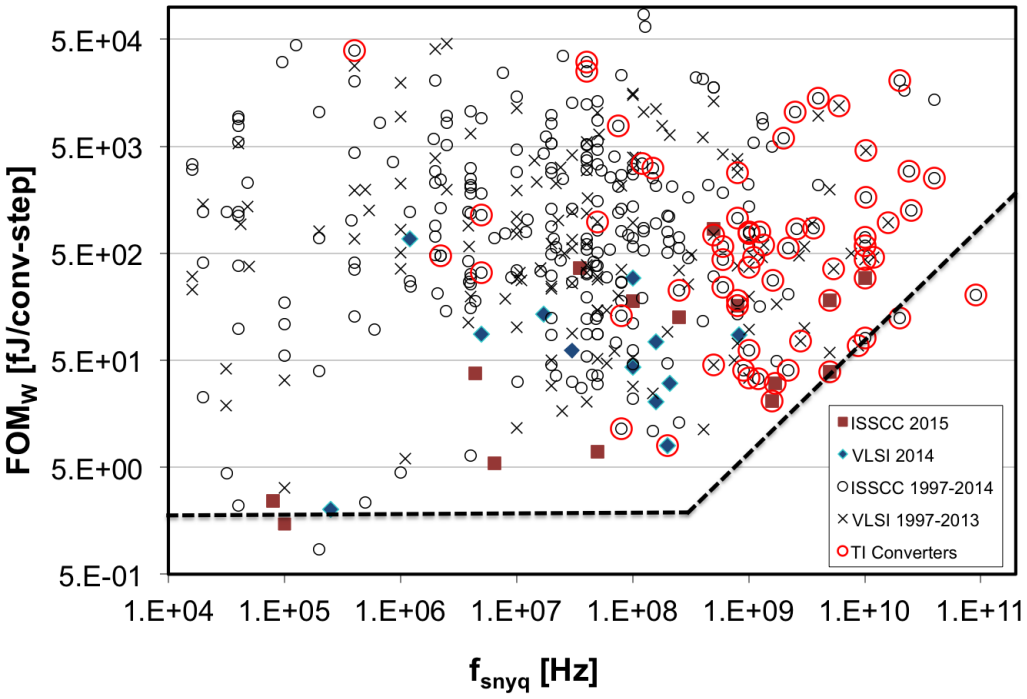

Figure 2.1: Walden FOM versus Nyquist sampling rate.

The TI-ADCs are circled in red and dominate the right-hand side of the plot ($f_{s,nyq} >$ 1GHz). The reason for its dominance at these high frequencies is that the TI-ADC enables more energy-efficient designs than the flash ADC, another popular architecture for high-speed ADCs. The ability to realize energy-efficient designs becomes apparent after looking at the architecture and operation of the TI-ADC.

A block diagram of the TI architecture is shown in Figure 2.2(a). It is comprised of M sub-ADCs working in parallel, each sampling the input signal at $\frac{1}{M^{th}}$ the full sample rate (i.e. $\phi_{1...M}$ are clocks of frequency $\frac{f_s}{M}$). The phases of the sampling clocks for each channel ($\phi_{1...M}$) are equi-spaced across the full sample period (Figure 2.2(b)). The outputs of each channel are time-multiplexed into a single stream at the output that represents the samples of the input signal at the full sample rate, f_s.

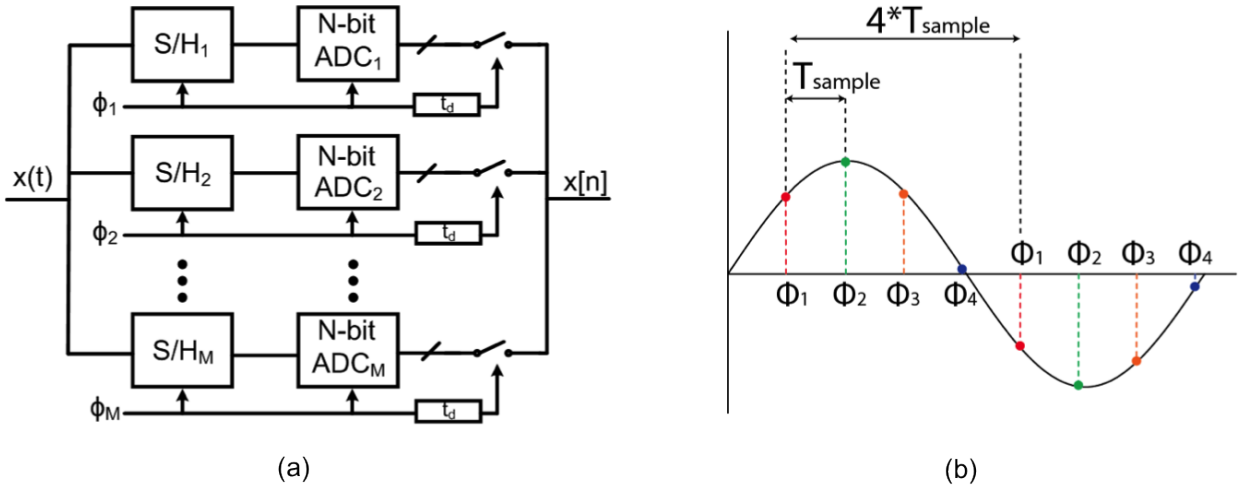

Figure 2.2: (a)Time-interleaved ADC architecture (b) Sampling diagram for 4-way TI-ADC.

Since each sub-ADC is operating at a fraction of the full sample rate, more energy-efficient designs can be used. This is because the sub-ADC has a longer period to perform the sample-and-hold and comparison operations, which can be very power hungry when very small conversion times are required. Take the CML comparator for example. It is comprised of a gain stage followed by a regenerative latch (Figure 2.3). The combined amplification time, t_{amp}, and regeneration time, t_{regen} must be less than the clock period. t_{regen} is typically the bottleneck of the two, and is given by[9]:

$$t_{regen} = ln(A_{latch}) \cdot \frac{C_L}{g_m} \tag{2.1}$$

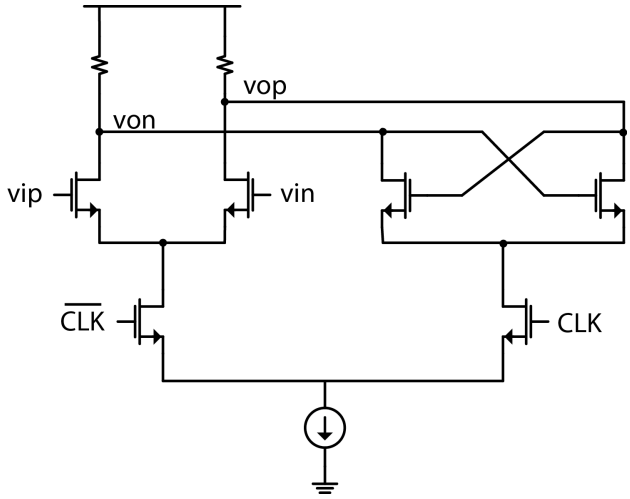

Figure 2.3: CML comparator.

where A_{latch} is the gain of the regenerative stage, C_L is the load capacitance on each output node, and g_m is the transconductance of each transistor. For high-speed comparators (i.e. when the allotted conversion time is small), external loads will be minimized and C_L will likely be dominated by the drain and gate capacitance of the transistors themselves. In this scenario, the transistors are considered to be *self-loaded*, meaning that they are mainly driving their own capacitance. When this is the case, the attainable speed is limited by the f_T of the process since:

$$t_{regen} = ln(A_{latch}) \cdot \frac{C_L}{g_m} \tag{2.2a}$$

$$\approx ln(A_{latch}) \cdot \frac{C_{gs} + C_d}{g_m} \tag{2.2b}$$

$$\approx ln(A_{latch}) \cdot \frac{1}{f_T} \tag{2.2c}$$

Thus, further speed improvements can only be obatined by proper biasing for peak f_T. Otherwise, attempting to improve speed by increasing the g_m, keeping V^* constant, would lead to an exponential increase in power consumption since the load capacitance will also grow almost linearly with the g_m, thus cancelling much of the gains afforded by the increase in g_m and breaking the inverse-linear relationship between g_m and t_{regen}.

When the comparator is given more time to make a decision, as is the case in the TI architecture, the comparator can be designed with much smaller transistors that are no longer *self-loaded*. The linear relationship between speed (t_{regen}) and power (g_m) is maintained and low power comparators can then be utilized, leading to a more energy-efficient design.

TI-ADC Designs and Challenges

Input Buffer

There has been significant progress over the past several years in improving the energy-efficiency and speed of TI-ADCs. Recently, a TI-ADC capable of sampling at rates up to 90GS/s was presented [10]. Designed in a 32nm SOI process, one of the key merits of this design was the sampling circuitry and clock generation which are critical for time-interleaved designs and can consume a significant fraction of the overall power. Although, the Nyqusit frequency is 45GHz at this sample rate, the acqusition bandwidth of the converter was limited to 20GHz. This is likely due to the limited bandwidth of the input buffer, the details of which was omitted by the authors. Nevertheless, the input buffer design is generally a huge bottleneck in TI-ADC designs with a large number of channels due to large input capacitance. As a result, the input buffers can ultimately limit the acquisition bandwidth of the converter and/or consume a lot of power[1].

Channel Mismatch and Calibration

Another design challenge for TI-ADCs is the effect of channel mismatch. Ideally, all channels should have an identical response to the input signal so that when their output streams are combined, it appears as if the input was sampled by a single-channel ADC. Unfortunately, there is always mismatch amongst the channel in the form of *gain*, *offset*, *clock skew*, and *bandwidth*. The presence of these types of mismatch leads to distortion tones to appear in the output spectrum, ultimately limiting the attainable SNDR [11].

There has been extensive work done on mitigating the effects of mismatch via calibration. [12] and [13] illustrate some of the modern calibration techniques, such as derivative estimation, and achieve 8b ENOB for speeds of 2.8 and 1.6GS/s, resepctively.

Clock Jitter

As mentioned in Chapter 1, the majority of high-speed ADCs operate in the *jitter-limited regime*. In this region of operation, the attainable ENOB is no longer limited by thermal noise or distortion, but is instead limited by the aperture error caused by jitter on the sampling clock. Timing errors in the sampling instant result in voltage errors in the sampled signal (Chapter 3) which degrade the overall SNR of the sampled signal. The theoretically attainable SNR in the presence of rms clock jitter, σ_j, is given by:

$$SNR_j = -20log(2\pi f_{in}\sigma_j) \tag{2.3}$$

According to Equation 2.3, in oder to achieve 25 GHz of acquisiton bandwidth and at least 6b of resolution, the front-end sampler needs to be driven by a (50GHz) clock with less than 80fs of rms-jitter, which is extremely difficult to obtain from an integrated source.

[1]The power consumption of the input buffer in [10] was not included in the reported power number.

Current state-of-the-art frequency synthesizers at similar frequencies achieve rms-jitter values on the order of 200-300fs [14][15]. This is why almost all high-speed ADC publications use an external clock source for testing.

For some high-speed applications (e.g. frequency-domain channel equalization), 4-5b resolution ADCs are sufficient to meet the system specifications and the TI architecture can be used. However, there are both current and emerging applications, such as real-time wideband signal capture (oscilloscopes) and mm-wave imaging, that require resolutions in the range of 6-8b which cannot be done without a very clean clock source and, in turn, a lot of power. Thus, it of interest to explore alternative topologies that can potentially relax the sensitivity of the ADC system to the clock jitter and improve the attainable SNR of high-speed ADCs. One potential candidate is the frequency-interleaved ADC.

2.2 Frequency-Interleaved ADCs

The frequency-interleaved ADC (FI-ADC) was first proposed in [5] as a means to circumvent the channel mismatch errors present in time-interleaved arrays. A block diagram of the *hybrid filter bank (HFB)*, as it is refered to in [5], is shown in Figure 2.4.

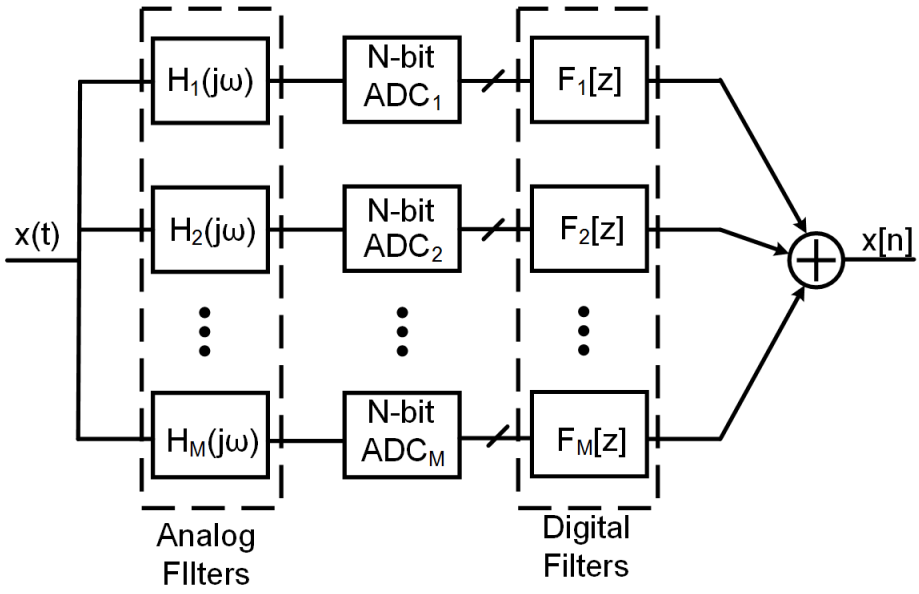

Figure 2.4: Block diagram of the hybrid filter bank.

The converter is composed of M sub-ADCs sandwiched between a bank of analog (analysis) filters and digital (synthesis) filters. Hence the name *hybid* filter bank. The analog filters assign a frequency band to each sub-ADC while the digital filters perform the reconstruction of the signal. It is shown that, on average, the filtering mitigates the non-ideal effects caused

by gain and phase mismatch between the channels. This is because a single-tone signal will ideally be processed by a *single* channel, whereas in the TI case, the output samples for a single-tone signal are taken from *all* channels, allowing the mismatch errors to compile together. The worst-case performance occurs in the transition band between channels since a single-tone in that band will be simultaneously processed by two channels.

Although the architecture and analysis presented in [5] focused on mitigating channel mismatch effects, the author briefly mentions the potential for reduced jitter sensitivity due to the reduced bandwidth presented to the sampler. Many works have since focused on applying this idea to various applications in order to take advantage of both the channel mismatch and jitter sensitivity benefits [16][17][18][19][20][21].

In [16], the authors used a three-channel frequency-intereaved (channelized) architecture for a 12.5GS/s serial link. Although the resolution is relatively low (3b), the authors were aiming to improve upon the performance of high-speed digitizers available at the time, citing jitter sensitivity and channel mismatch as the major bottlenecks for TI converters. Interestingly enough, the authors further interleaved the sub-ADCs using a 2-way TI-ADC archtiecture in order to further relax the sampling speed requirements of each ADC. In that design, the top-level frequency-interleaved architecture enables higher resolutions to be achieved via reduced jitter sensitivity, while the low-level time-interleaved architecture allows for energy-efficient sub-ADCs to be utilized. Lastly, I/Q downcoversion mixers are used in each passband channel in order to translate each frequency band down to baseband before sampling and reduce the required sampling speed of each sub-ADC (Figure 2.5). This downconversion was not needed in [5] due to the low overall sampling speed.

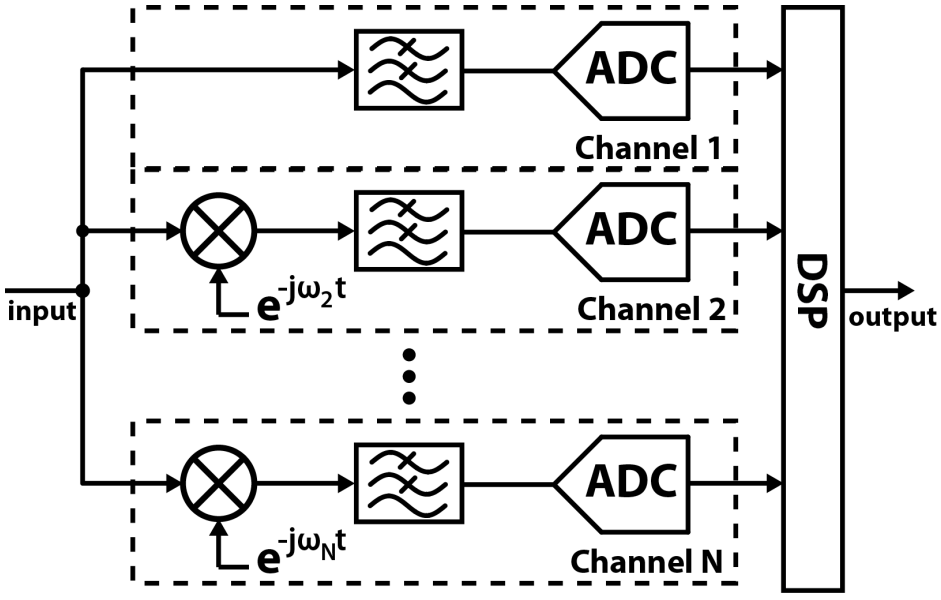

Figure 2.5: FI-ADC with downconversion mixers.

In [18], a two and three-channel FI architecture is used to design a 4 and 6GS/s ADC with 4b resolution in 90nm CMOS. The design integrates the full mixer-filter-ADC path. Reconstruction of the input signal is done offline and incorporates digital correction of gain and phase mismatches in the I/Q paths and gain/offset mismatches across the channels. A more detailed discussion of the digital compensation techniques are discussed in [22]. Measurements show a peak ENOB of 3.5b and futher demonstrate how a moderate ENOB is maintained over the entire Nyqust band (less than 1b degradation). This relatively flat ENOB response is a theoretical characteristic of the FI architecture and is further discussed in Chapter 3.

At this point, it should be noted that the designs in [16] and [18] restrict the number of channels in the FI architecture to two or three. The benefit of this decision is that these designs aren't susceptible to the harmonic folding problem that occurs once your number of channels is greater than three (Chapter 3). Unfortunately, as the speed of the converter is increased[2], a larger number of channels becomes necessary in order to allow for energy-efficient (and feasible to implement) sub-ADCs to be used and lead to an overall energy-efficient design. As such, harmonic rejection techniques need to be employed in the design of these FI architectures. Additionally, the larger channel count increases the input capacitive load and makes the distribution of the wideband signal to all channels more difficult and/or power hungry. Techniques to address these two key problems are presented in the later chapters.

In [19] and [20], the FI architecture is used to design wideband multi-channel receivers with more than three channels. [19] uses a five-channel FI architecture to simulate a 5GHz bandwidth OFDM receiver. The effective sample rate of the receiver is 10GS/s with a targeted SNR of 40dB (6.35b resolution). The authors propose the usage of "bandwidth-optimized" low-order filters in the sub-channels in order to reduce power consumption and design complexity. According to their simulations, reducing the bandwidth of first and second-order filters to below the pre-allocated channel bandwidth results in significant reduction in jitter sensitivity. This is due to the improved filtering of out-of-band signals which would normally reciprocally mix with the wideband phase noise of the LO and degrade the in-band noise floor. Since in-band signals are also partially attenuated in this scheme, they rely on their reconstruction algorithm to correct for any in-band losses. Furthermore, their simulations show that the SNR degradation due to jitter dominates the degradation due to noise added during mixing, which isn't the case for the design presented in this thesis. It is likely that the aggressive filtering in each channel helps mitigate the impact of the LO phase noise added during mixing (see Chapter 3).

While the techniques proposed in [19] are innovative and have the potential to significantly relax the design of the LOs and channel filters, they are not adopted in this design for two reasons: (1) Group delay variations across channels will be severe due to the purposeful in-band filtering. This may take more effort to correct in the digital domain since the bandwidth of the equalized channel would have to be extended further than the pre-allocated

[2]The target of design presented in this thesis is 50GS/s.

bandwidth[3]. For applications requiring the reconstruction of the full wideband signal, mismatch in the group delay between channels will result in imperfect reconstruction of the signal. (2) The corner frequency of the analog filters would need to be well controlled in order for the digital backend to properly equalize the in-band attenuation without causing further distortion. This neccessitates tuning circuitry for the filter's corner frequencies, and adds to the design complexity.

In [20], a four-channel FI architecture is used to design a 2GS/s sampling receiver in 65nm CMOS. The receiver has an acquisition bandwidth of 125MHz to 1GHz and maintains an impressive (mean) ENOB of 7.8b across the entire bandwidth and harmonic rejection greater than 59dB. At the time of this writing, this is one of the very few (and possibly only) receivers to address the haromonic folding problem that plagues the FI architecture when more than three channels are used. The harmonic rejection is performed in the digital backend via the shuffling of weighted coefficients to synthesize a sampled-and-held sinusoid with reduced harmonic content. The coefficient shuffling is also cleverly used to change the frequency of the effective LO, enabling programmable channel selection. The one drawback of this architecture is that it only allows the reception of a single channel at any given time. Although adequate for their targeted application, this is not sufficient for applications requiring real-time reception of an entire broadband spectrum.

It should also be noted that the FI-ADC technique has found usage in commercial applications, specifically high-speed oscilloscopes [23]. In this patent, a two-channel architecture is used to practically double the acquisition bandwidth of a real-time oscilliscope front-end. By simply highpass filtering and downconverting the upper-half of the spectrum, the same high-speed baseband ADC can be used. With only a two-channel architecture, harmonic folding is not a problem for their design. Additionally, since it isn't a power-constrained design (wall-powered), it isn't an issue to re-use the very power hungry 40GS/s sub-ADC for each channel. For low power designs, this architecture would of course have to be altered to incorporate more channels, as done in this thesis. However, it is still interesting to see that the FI architecture has been utilized to design modern oscilloscopes with effective sample rates of 240GS/s and resolutions of 6-8b.

The basics of digital reconstruction are also detailed in [23] and are used in the design of the offline digital reconstruction for the FI-ADC presented in Chapter 4.

[3]See Chapter 4 for further discussion.

2.3 Frequency Domain Sampling vs. FI-ADCs

Frequency-domain sampling is an alternate version of the frequency-interleaving topology which has also been researched [21][24]. The key difference is the replacement of the LPF with an integrator (Figure 2.6). The integrator and sampler effectively perform a DFT operation and, as a result, the reconstruction of the signal is done by performing an IFFT in the digital backend. As discussed in [21], because the DFT inherently assumes a periodic signal, the improved ADC resolution is only achieved for frequencies that are T_i-periodic, where T_i is the length of the integration period. In order to solve this problem, a windowing function must be applied to the signal before sampling[4]. Once this window function has been applied, the output signal becomes an *approximation* of the input signal, in the best case.

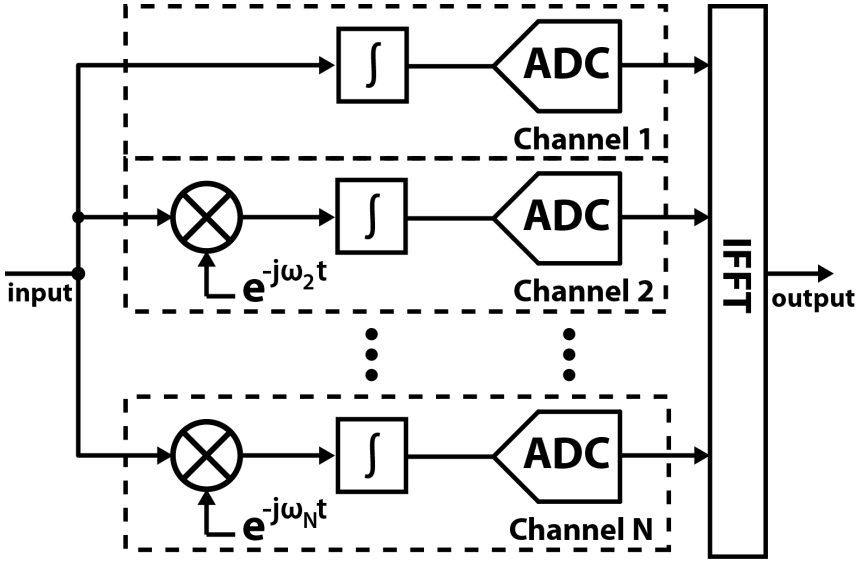

Figure 2.6: Frequency-domain sampling architecture.

Although the frequency-domain ADC may have merits for potential usage in multi-band receivers [24], it is not adopted here. For one, the *approximation* that is inherent in the architecture is not be suitale for applications requiring accurate reconstruction of the wide-band signal (e.g. oscilloscopes). Secondly, the windowing function effectively smooths out any sharp transitions/edges in the signal so that the signal could be represented by a finite number of DFT coefficients. This is undesirable for pulsed-radar imagers which rely on time-of-arrival (TOA) information. The TOA is encoded in the pulse edge and smoothing out this edge can alter the measured TOA and affect the accuracy of high-resolution imagers.

[4]This is commonly done in DSP for DFT/FFT calculations of arbitrary waveforms.

Chapter 3

The Frequency-Interleaved ADC

3.1 Proposed FI-ADC Architecture

A diagram of the proposed architecture for the FI-ADC analog front-end (AFE) is shown in Figure 3.1. The architecture is similar to that of [16][17][18] and [19] with the addition of a wideband distributed amplifier (DA) which is used to distribute the wideband signal to each channel of the ADC. Each passband channel is comprised of an I/Q mixer, lowpass filter (LPF), and sub-ADC. Using I/Q mixers halves the baseband bandwidth and allows for potential power savings in the design of the baseband filters and sub-ADC. The output of each sub-ADC is then sent to the digital backend for reconstruction of the original signal.

The digital backend for a single channel is shown in Figure 3.2 and is similar to that of [23]. The backend takes the samples from the sub-ADCs and immediately upsamples to the full sample rate. The lowpass filter following the upsampler acts as the interpolation filter. After upsampling, the signal is mixed with a digital LO signal whose frequency is the same as the frequency of the LO used for downconversion. This places the baseband signal back into its original location in the spectrum. I/Q paths are then re-combined. This processing takes place in each channel and the outputs of all channels are then summed in order to "re-stitch" the spectrum back together. Equations for the DSP processing can be found in [5] and [22].

The mixer-LPF combination performs the first level of channel filtering in lieu of bandpass filters (BPFs). Placement of BPFs in each channel is optional for added filtering and channel isolation but is not used here. The reason for omitting BPFs is that a modular design was desired in which all sub-blocks in each channel are identical so that a *design once and reuse* approach could be adopted. If BPFs were used, they would require independent designs that are tuned for each channel's center frequency. Unfortunately, the omission of BPFs gives rise to harmonic folding in low-frequency channels. Without a BPF preceeding the mixer, the entire input wideband sprectrum is present at the input of mixer and harmonics of the LO may downconvert out-of-band signals to baseband. The following section elaborates on this issue and presents a proposed method for solving the problem.

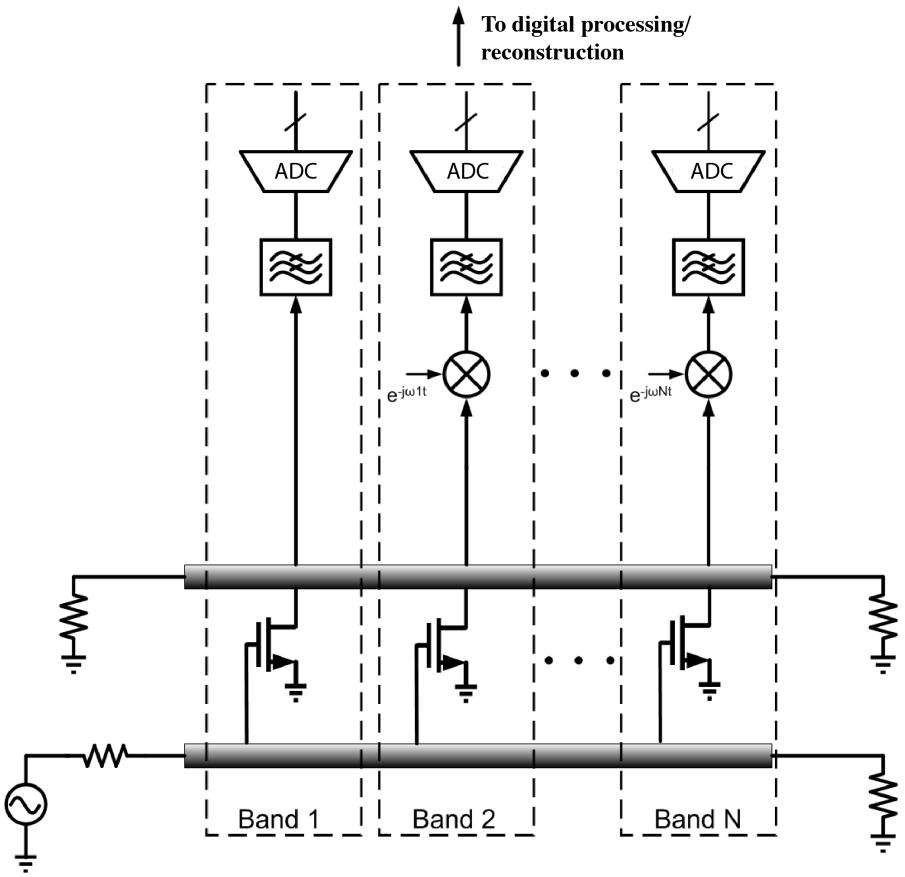

Figure 3.1: Proposed FI-ADC Architecture.

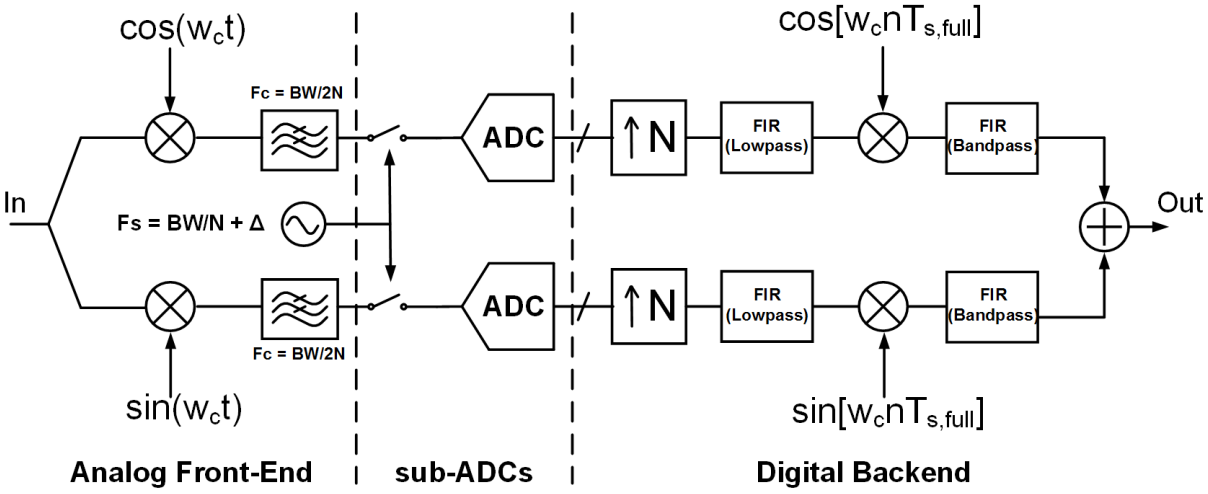

Figure 3.2: Simplified processing for digital reconstruction.

Key Challenges

Wideband Signal Distribution

Using a distibuted amplifier for distributing the signal to each channel is advantageous since the capacitive loading of each channel is absorbed into the drain transmission line of the DA, resulting in a wider bandwidth distribution network. In contrary, if a single buffer were used to drive all channels in parallel, the buffer would be extremely power hungry since it has to drive a very large capacitance. Furthermore, the achievable bandwidth would ultimately be limited by the f_T of the process, whereas the bandwidth of the DA is limited by the capacitive loading of each channel and the minimum inductance that can be accurately fabricated in the process (ignoring contraints on Z_o). The latter is usually much larger than the former[1]. Another advantage of using a DA is that it breaks the gain-bandwidth tradeoff that is faced by traditional amplifiers. This gives more freedom in the design, allowing for optimized system noise figure, by increasing DA gain for example, without sacrificing bandwidth.

An alternative to the above distribution approaches would be a passive tree distibution network comprised of Wilkinson dividers (Figure 3.3). This approach has two attractive qualities: (1) wideband and (2) zero power consumption. The major drawback, however, is that the inherent insertion loss (IL) from the input to each channel increases with the number of channels, N. This is the case even if the Wilkinson divider is comprised of lossless passives. The IL is further reduced once real, lossy passives are considered. Higher IL has a direct hit on the sensitivity of the ADC and results in a reduced dynamic range (resolution) for the ADC. As will be seen, the number of channels is a key design parameter in the optimization of the ADC. Thus, it is desirable to de-couple the sensitivity of the ADC from the number of channels in order to allow more freedom in choosing the number of channels. For this reason, the passive power distribution approach is not used.

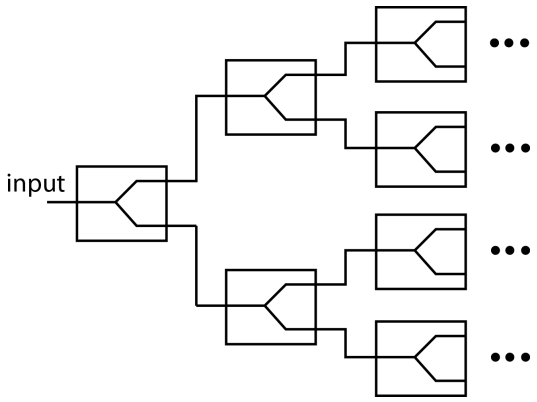

Figure 3.3: Wilkinson divider chain for broadband distribution.

[1]The extracted f_T of modern 65nm CMOS processes is $\sim$ 200GHz, while the theoretical cutoff frequency, $f_c = \frac{1}{\pi\sqrt{LC}}$, of an integrated transmission line would be $\sim$ 285GHz assuming L=50pH and C=25fF.

Harmonic Folding

As previously mentioned, the proposed FI-ADC architecture is susceptible to harmonic folding. Harmonic folding occurs in the proposed FI-ADC architecture for two reasons: (1) we are targeting larger than three channels and (2) there are no BPFs present before downconversion. These two design choices are critical for increasing the energy-efficiency of the architecture as well as simplifying the design. Due to the absence of bandpass filters, the full wideband signal is present at the RF input of the mixers. In an ideal mixer, the RF signal is mulitplied by a perfect sinusoid whose frequency domain representation is a delta function at $+w_{LO}$ and $-w_{LO}$. The output spectrum is then a sum of two frequency shifted versions of the RF input – a shift right by w_{LO} due to the delta function at $-w_{LO}$ and an identical shift left due to the delta function at $+w_{LO}$. After mixing, all energy around w_{LO} is now located in the baseband (Figure 3.4). Note that by using a complex mixer, we can effectively multiply the signal by a complex sinusoid (e^{jwt} or e^{-jwt}) and can extract either the upper sideband or lower sideband information.

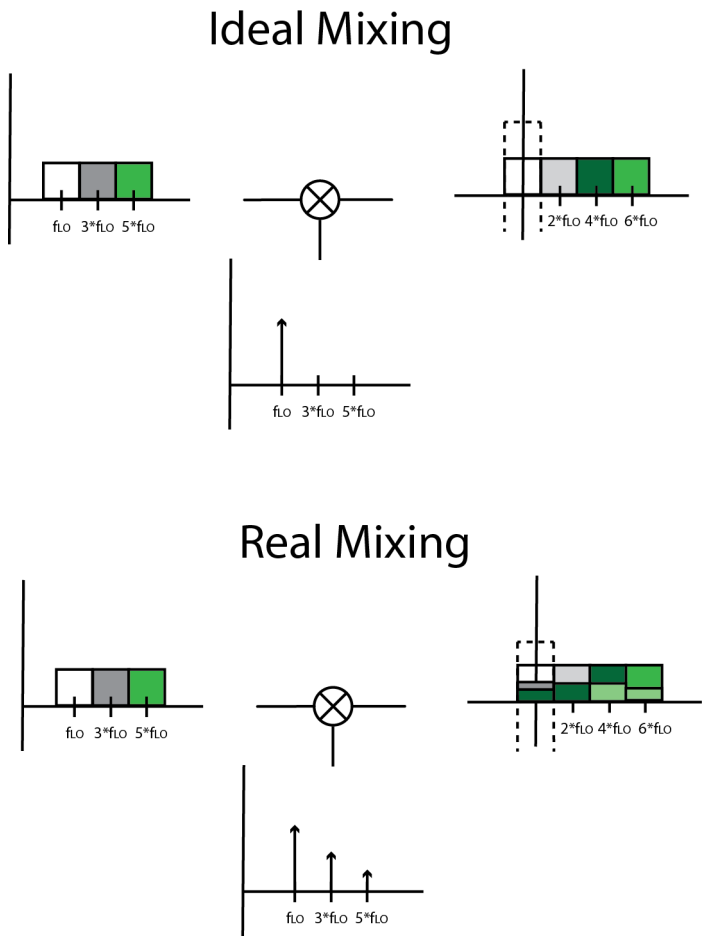

Figure 3.4: Harmonic folding corrupts the baseband signal.

In the case of real mixing, where a simple Gilbert switching quad is used, we are multiplying the RF signal with a square wave with a period of T_{LO}. The spectrum of the LO signal will thus have the fundamental tone along with its odd harmonics ($3w_{LO}$, $5w_{LO}$, $7w_{LO}$, etc.)[2]. As a result, energy around these harmonics will also be downconverted to the baseband and corrupt our desired signal. In the context of the FI-ADC architecture, this is a highly undesirable effect since our input signal is wideband and harmonics of the LO will likely fall within the bandwidth of the signal and downconvert out-of-band energy.

The lower frequency channels of the FI-ADC are most susceptible to harmonic folding since a larger number of their harmonics are likely to fall within the bandwidth of the input signal. For a fixed input signal bandwidth, increasing the number of channels leads to a larger number of channels that are susceptible to harmonic folding. Roughly speaking, the i^{th} channel will have an LO frequency of

$$\omega_{LO,i} = \frac{i \cdot \omega_{max}}{N} \tag{3.1}$$

where ω_{max} is the maximum input frequency of interest and N is the number of channels used in the FI-ADC. Considering the third harmonic of the LO, folding occurs if

$$3\omega_{LO,i} < \omega_{max} \tag{3.2}$$

or equivalently

$$\frac{i}{N} < \frac{1}{3} \tag{3.3}$$

Since $1 \leq i \leq N$, it can be seen that harmonic folding is a non-issue for FI-ADC architectures with three or less channels ($N \leq 3$). For $N > 3$, there are values of i for which the condition in Equation 3.3 is met. When N is increased, this condition is met for a larger number of values for i (i.e. more channels experience harmonic folding of the third harmonic). Similar conditions can be derived for the other harmonics of the LO. The key insight here is that increasing the number of channels exacerbates the harmonic folding problem. Since harmonic rejection techniques add complexity to the mixer design, there is a practical limit to how large N should be made.

Now, for an energy-efficient design, large N is desirable. Therefore, we will have to mitigate the impact of harmonic folding on our system performance. This can be done by utilizing harmonic rejection mixers (HRMs). Harmonic rejection techniques have been well researched in the context of wideband tranceivers for cellular and TV applications [25][26][27]. The majority of these techniques are based on the topology proposed by Weldon [28]. The core idea is to emulate a multiplcation by a sampled-and-held sinusoid (SHS). The SHS waveform can be produced by performing a weighted summation of phase shifted LO clock

[2]By using a differential topology in the mixer, the even-order components of the square wave are removed and only the odd-order components remain.

signals. By emulating a sinusoid, the harmonic content of the LO reduces and, in the limit, the spectrum of the LO coverges to a single delta function.

As discussed in [28], the level of harmonic rejection depends on the accuracy of the LO phase shift and magnitude weightings. For cellular and TV applications, the LO fundamental is typically below 1GHz. This allows for the usage of advanced digital techniques to generate accurate multi-phase LO signals without consuming considerable amounts of power. This is not the case for the wideband FI-ADC where the LO signals can easily be 3GHz or larger. For example, in our system the LO signals requiring harmonic rejection are 3 and 6GHz. At these frequencies, it is very difficult and power consuming to use digital techniques to generate the multiple phases of the LO required for Weldon's HRM architecture. Thus, we need to design a HRM architecture that is capable of operating at higher LO frequencies while providing adequate harmonic rejection.

An architecture to acheive harmonic rejection at high LO frequencies is shown in Figure 3.5. The high-frequency HRM (HF-HRM) architecture is comprised of a main mixer and an auxiliary mixer whose outputs are current-summed out-of-phase. The auxiliary mixer is driven by the harmonic of the LO that we wish to cancel (in this example we are cancelling the 3^{rd} harmonic).

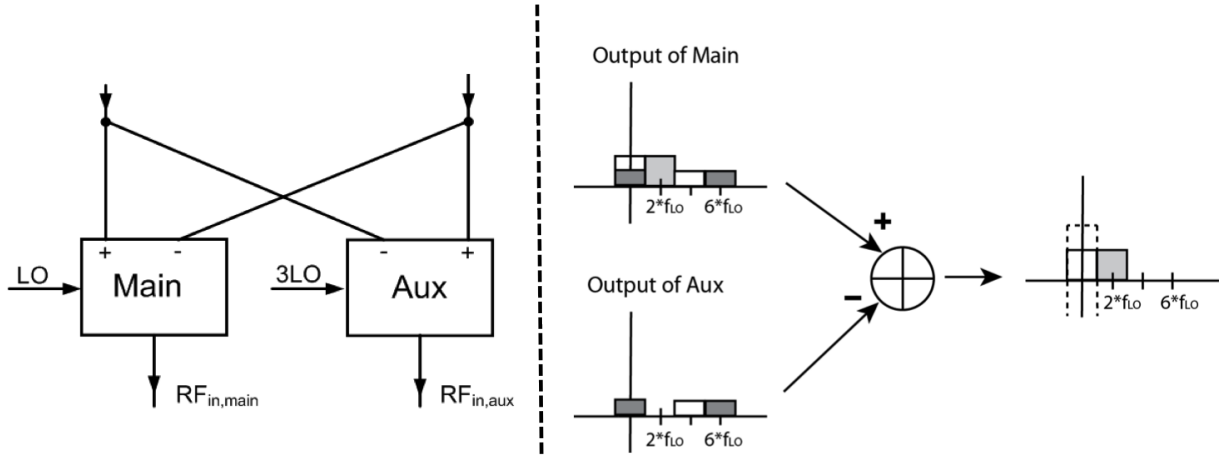

Figure 3.5: Block diagram of high frequency HRM.

From Figure 3.5 we see that the output spectrum of the main mixer consists of a corrupted baseband signal while the auxiliary mixer only produces the corrupting signal (energy around $3\omega_{LO}$ in the baseband). By summing the two paths out-of-phase we can generate a clean baseband signal that is absent of harmonic folding. This architecture can be extended to higher harmonic cancellation by adding additional auxiliary mixers that are driven by the harmonics to be cancelled. The appeal of this HRM design is that it fits naturally within the FI-ADC architecture since there is no overhead in the generation of the harmonic LO frequencies – they are already generated since they will be the fundamentals for higher frequency channels and are readily available from the LO generation block. In conventional

cellular/TV receivers, the harmonics of the LO are not readily available, making this approach unattractive for such systems.

As with any noise/interference technique that relies on the summation of two paths, gain and phase mismatch of the two paths limit the attainable level of cancellation. Further discussion of the HF-HRM and its design are provided in the next chapter.

LO Generation/Distribution

A major overhead in the FI-ADC design is the generation and distribution of the multiple LO frequencies needed for downconversion. For frequency generation, QVCO-based PLLs can be utilized for the highest required LO frequency. The frequency divider chain embedded inside the PLL can then be tapped to access the LO frequencies that are integer-related to the fundamental frequency. For example, the highest LO frequency required in this design is 24GHz and the 12GHz, 6GHz and 3GHz LOs required by other channels would be available from the dividers in the PLL. The challenge here is generating the frequencies that have a non-integer relation to the PLL fundamental frequency (e.g. generating 18GHz from 24GHz). In addition, all LOs must be phase locked so that the phase delay between channels doesn't drift over time and won't cause any errors during the reconstruction of the signal.

In this design, fundamental frequency generation wasn't performed since there have been works demonstrating the feasibility of designing low-power (< 50mW), very high-frequency (> 20GHz) PLLs in modern CMOS processes [14][15]. However, generation of all other LO frequencies (assuming two phase-locked fundamental frequencies are available) is performed on chip and techniques for generating the non-integer related frequencies are presented in Chapter 4.

Once the LO frequencies have been generated, the next challenge is in distributing the LO to the mixer inputs. As will be seen later in this chapter, LO phase noise plays a critical role in the performance of the FI-ADC and must be properly managed in order to achieve performance gains over the TI-ADC. As a result, any buffers used in distributing the LO must not add significant noise/jitter to the LO. This can result in high power consumption in the LO buffers. Thus, the number of LO buffers used should be minimized. Furthermore, care must be taken in the distribution of the highest LO frequencies in order to minimize signal loss. As such, transmission lines must be utilized and the floorplanning must be architected to minimize the trace lengths of these high-frequency signals.

Lastly, with numerous LOs being generated and distributed across the chip, crosstalk is another major concern. Any spurious tones that arise as a result of crosstalk between LO distribution lines will result in downconversion of spectrum from other channels, an undesirable effect. To meet system specifications, the crosstalk between LOs needs to be below ~ 45dB so that the undesired downcoverted spectrum falls below the noise floor. To meet these specs, proper layout techniques and floorplanning must be employed. Techniques for minimizing the LO crosstalk are presented in Chapter 4.

Digital Reconstruction

Proper reconstruction of the signal in the digital domain is another challenge in the FI-ADC. Although it is claimed that the FI architecture is less senstivite to channel mismatches than the TI architecture [5], equalization across channels is still necessary for applications requiring real-time reconstruction of a wideband signal. For a single-firequency excitation[3], the FI architecture will show better performance than the TI architecture because the output stream is given from a single channel as opposed to being a composite of output streams from various channels. In this case, the FI architecture does not require equalization across channels and only needs to equalize the I/Q paths within each channel for proper image rejection. On the other hand, for a wideband excitation (e.g. narrow-width pulse), multiple channels will contribute to the FI architecture's output and will therefore need to have identical performance so as to not introduce distortion when recreating the time-domain signal. As far as the AFE is concerned, the *gain* and *group delay* of all channels should be equal to allow the digital backend to perform proper reconstruction.

In addition, the AFE still needs to retain suffiicient *gain* and *phase* matching between the I and Q paths within in each channel in order to achieve adequate image rejection. Gain/phase mismatch results in finite image rejection. For a 6b system, the desired image rejection is $>$ 40dB. In order to obtain this level of rejection, the gain and phase of I/Q paths must match within 0.1dB and 1°, respectively.

Digital reconstruction and equalization techniques are not the focus of this work and any equalization required in the backend is performed offline in a manual fashion. Furthermore, [22] has proposed digital compensation techniques for channel mismatch in the downconverting FI architectures. As will be seen in the next chapter, many design decisions for the implemented FI-ADC were made with the intent to minimize the amount of digital compensation required.

[3]This is the standard excitation used for testing ADC performance. However, it should be noted that it is not sufficient to fully compare the FI and TI architectures.

3.2 FI-ADC vs. TI-ADC

One of the key merits of FI-ADCs is the potentital to break the jitter barrier faced by conventional sampling architectures. By limiting the bandwidth of the signal presented to the sample-and-hold circuitry, the jitter requirement on the sampling clock for achieving a specific ENOB is relaxed. In order to reduce the required sub-ADC sampling rate for passband channels, mixers are used to downconvert the passband channels to baseband and allow for a lower sampling rate to be used. Consequently, this "mix-then-sample" architecture has introduced a new process that isn't present in conventional ADC sampling architectures - mixing - and it is worth exploring how the two approaches compare.

The following sections present both a qualitative and quantitative comparison of the FI-ADC to the TI-ADC. First, conventional ADC sampling (direct-sampling) is compared to the mix-then-sample architecture in terms of noise performance and ENOB as a function of input frequency. Next, a detailed discussion on the impact of LO phase noise is presented followed by results from a system-level simulation. The simulation is used to accurately compare the FI-ADC to conventional direct-sampling architectures.

Direct Sampling vs. Mix-then-sample

It is of interest to compare the FI-ADC's mix-then-sample (MTS) front-end architecture to a direct-sampling (DS) architecture which feeds the wideband input signal directly into a Nyquist rate sampler (Figure 3.6).

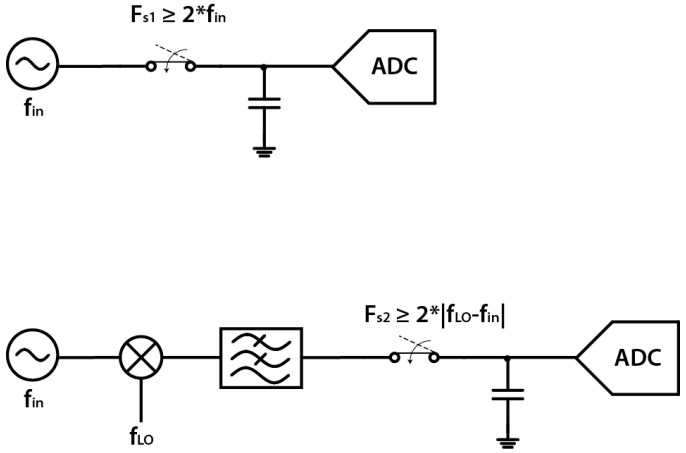

Figure 3.6: Sampling architectures: (top) Direct-sampling and (bot) Mix-then-sample.

In the DS approach (Figure 3.6(a)), the wideband signal is fed directly into a sampler that samples the signal at the Nyquist rate of $2 \cdot f_{in,max}$ and then passes the sampled signal to a quantizer. In the MTS approach, the input signal is downcoverted by an LO signal, low-pass filtered, sampled at the *baseband* Nyquist rate of $2 \cdot f_{IF,max} \approx 2 \cdot \frac{f_{in,max}}{2N}$ and then

quantized. The lowpass filter has a cutoff frequency dictated by the sub-channel bandwidth and limits the maximum frequency presented to the sampler to $\frac{f_{in,max}}{2N}$, where N is the number of channels in the FI-ADC[4]. We are now interested in comparing the two sampling architectures performance in regards to three metrics: (1) input noise, (2) ENOB v. input frequency, and (3) LO phase noise at the mixer input.

The impact of LO phase noise is of considerable interest. This is because the LO phase noise is transferred to the input signal during the mixing process and although we are less senstivite to jitter at the *sampling instant*, the jitter problem may have simply been transformed into a phase noise problem at the *mixing instant*. Almost all publications related to the FI architecture simply cite the decreased jitter sensitivitiy with no mention of the impact of phase noise. [19] does address this issue, but simply shows, via simulation results, that the sampling jitter was more a performance limiter than the "mixing jitter" and therefore they do not provide any detailed discussion on how the "mixing jitter" impacts the system. Similarly, [21] presents encouraging simulation results that incorporate LO phase noise but offers no supporting discussion. Here, we provide a quantitative discussion of the impact of LO phase noise.

Input Noise

One of the drawbacks of the MTS architecture is that it has a worse response to input noise than the direct-sampling architecture. The reason for the degraded response is noise folding that occurs during the mixing process. This can be better understood from Figure 3.7 which shows the spectrum of the input noise at the input of the mixer in the first passband channel (chan1) of a five channel architecture. Here, the input noise bandwidth is assumed to be limited to the maximum signal bandwidth, $f_{in,max}$. This is a legitimate assumption since all sampling systems have an anti-aliasing filter[5] placed before the sampler in order to prevent out-of-band signals from aliasing down to baseband and corrupting the desired signal. As a result, the input noise spectrum will be containted within this bandwidth and zero beyond that.

Due to lack of bandpass filters before the mixer, the entire input noise spectrum is present at the RF port of the mixer. The fundamental LO frequency, $f_{LO,1}$,will be centered in chan1 and downconvert the noise spectrum in its passband down to baseband. In addition, the third harmonic will also fall within chan3's passband and any noise in chan3's passband will also be downconverted to baseband. In the ideal case, we only want to process the noise located in the passband dedicated to that channel, but due to real mixing, the noise around the harmonics of the LO will also be folded down to baseband and degrade the SNR. Noise folding is identical to the harmonic folding problem since they both stem from the same non-ideality - harmonics of the LO downconverting out-of-band spectrum during the mixing

[4]The additional factor of 2 is due to the fact that we are using I/Q mixers which further cuts the baseband bandwidth in half

[5]An ideal anti-aliasing filter is a brickwall lowpass filter with a bandwidth equal to the maximum frequency of interest.

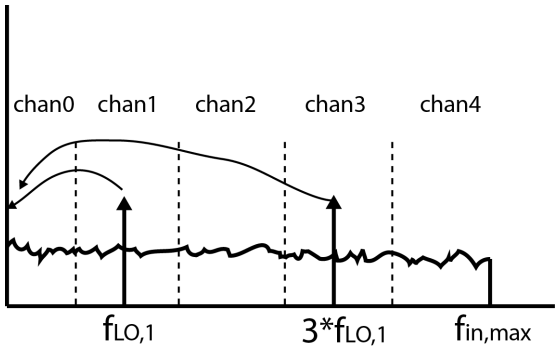

Figure 3.7: Input noise spectrum for mix-then-sample.

process. This means that the input-noise folding problem can also be mitigated by using harmonic-rejection mixers.

The DS architecture does not suffer from this issue of input noise folding. This can be easily seen by observing typical noise input spectrum for the DS case (Figure 3.8).

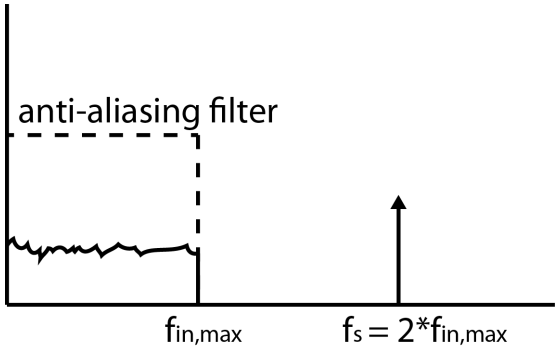

Figure 3.8: Input noise spectrum for direct-sampling.

An anti-aliasing filter will limit the noise bandwidth to $f_{in,max}$ and the sampling clock's LO frequency is located at the Nyquist sample frequency of $2 \cdot f_{in,max}$. As can be seen in Figure 3.8, the fundamental tone of the sampling clock falls out of band of the input noise and thus all the subsequent harmonics that comprise the spectrum of the sampling process will also fall out of band and there will be no noise folding in the direct sampling system.

ENOB vs. f_{in}

As illustrated in [18], the FI-ADC can maintain close to peak ENOB across the entire input frequency range as opposed to the TI-ADC which exhibits a degradation in ENOB as the input frequency approaches the Nyquist frequency. This superior ENOB performance is attributed to the bounds placed on the signal frequency at the input of the sampler.

As shown by Equation 2.3, SNR in the presence in the presence of aperture jitter decreases 20dB per decade increase in frequency. The SNR due to quantization noise is given by the well known formula:

$$SNR_q = 6.02 \cdot Q + 1.76 \quad [\text{dB}] \tag{3.4}$$

where Q is the number of bits used in the digital representation. There is a frequency, f_{corner}, for which $SNR_j = SNR_q$. Below f_{corner}, $SNR_j > SNR_q$ and the ENOB is set by the quantization noise. Therefore, ENOB is independent of input frequency in this region (Figure 3.9). Above f_{corner}, $SNR_j < SNR_q$ and aperture jitter becomes the performance limiter. This is the aforementioned *jitter-limited region*. The ENOB in this region is dictated by SNR_j. Using Equations (2.3) and (3.4), we see that the ENOB degrades 3b per decade increase in frequency in this region.

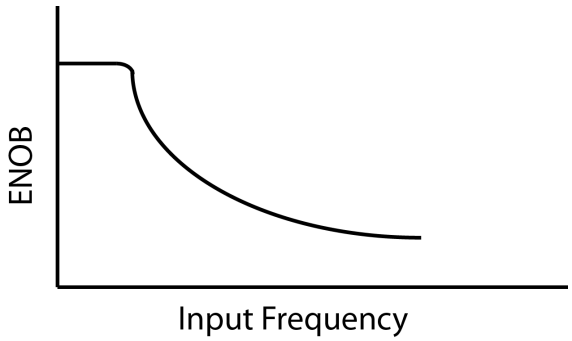

Figure 3.9: ENOB vs. input frequency (linear scale) for direct-sampling.

For the FI-ADC the theoretical[6] ENOB plot is shown in Figure 3.10.

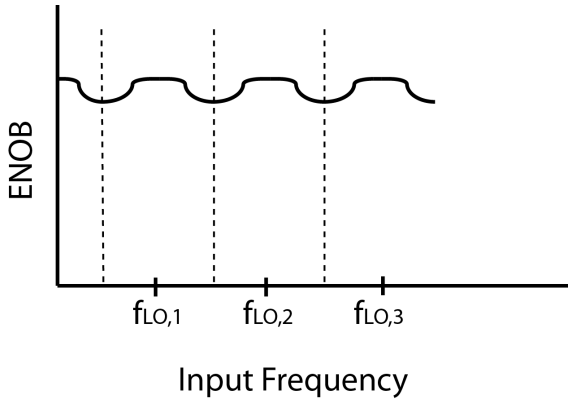

Figure 3.10: ENOB vs. input frequency (linear scale) for mix-then-sample.

[6]In this example, the only nonidealities are quantization noise and jitter on the sampling clock. Phase noise on the mixing LO is not considered.

To best understand how this plot is generated, consider an input frequency equal to the LO frequency of the first channel, $f_{LO,1}$. After mixing, the signal is translated to a DC baseband frequency (0 Hz). Because this is a very low frequency, SNR will be set by SNR_q. As the frequency is increased (or decreased) slightly, the ENOB will remain flat since the baseband frequency, $|f_{LO,1} - f_{in}|$, will still be a relatively low frequency. Increasing (or decreasing) the input frequency further, the baseband frequency becomes larger and we will enter the jitter-limited region and the ENOB will begin to degrade just like Figure 3.9[7]. However, unlike Figure 3.9 where the ENOB degrades indefinitely, the ENOB degradation will hit a minimum at the channel's frequency edge. Above that frequency, the input signal is now processed by the adjacent channel and will begin approaching that channel's LO frequency. For example, increasing the input frequency will push the input signal into channel 2 and the new baseband frequency is $|f_{LO,2} - f_{in}|$ which decreases in magnitude until $f_{in} = f_{LO,2}$. A decrease in baseband frequency leads to an increase in ENOB and an eventual re-entry into the *quantization-limited* region. After this point, the behavior repeats itself.

So to first-order, it appears that the FI-ADC can provide a relatively constant ENOB over the entire Nyquist frequency band. When LO phase noise is considered, this is not the case, and slight degradation in ENOB at higher frequencies may be observed. Interestingly, enough, the FI-ADC can still outperform the TI-ADC even when this high-frequency degradation is taken into account. There are many factors that determine the true relative ENOB performance in the presence of LO phase noise and sampling jitter. The following discussions elaborate on this topic.

Impact of LO Phase Noise

The last, and most important, point of comparison is the impact of LO phase noise. The key argument for the FI architecture is the potential for breaking the jitter barrier which plagues current high-speed ADCs. In order to do so, we have downconverted and lowpass filtered the signal before sampling. This limits the bandwidth of the signal going into the sampler and hence, relaxes the sensitivity to the jitter on the sampling clock. The question here is: what impact does the LO phase noise have on the noise performance down the chain. Afterall, phase noise and jitter are closely related. Also, it is likely that the LO used for downconversion and the sampling clock are derived from a common source and will have similar phase noise/jitter performance. So it isn't clear if there is a net gain since we potentially have a phase noise constraint on the mixing LO and this may be equivalent to the jitter constraint on the sampling clock in the DS achitecture.

Before we assess how LO phase noise impacts the noise performance of the MTS archtiecture, we first derive how jitter at the sampling instant limits the SNR performance of the DS architecture. This analysis will be valuable when we want to assess the SNR performance of

[7]Depending on the amount of sampling jitter, σ_j, and the channel bandwidths, there is a possiblity that the jitter-limited region is never observed

the MTS archtiecture. First, we observe that any timing error (i.e. jitter) on the sampling clock will result in an error in the sampled voltage of the input signal (Figure 3.11).

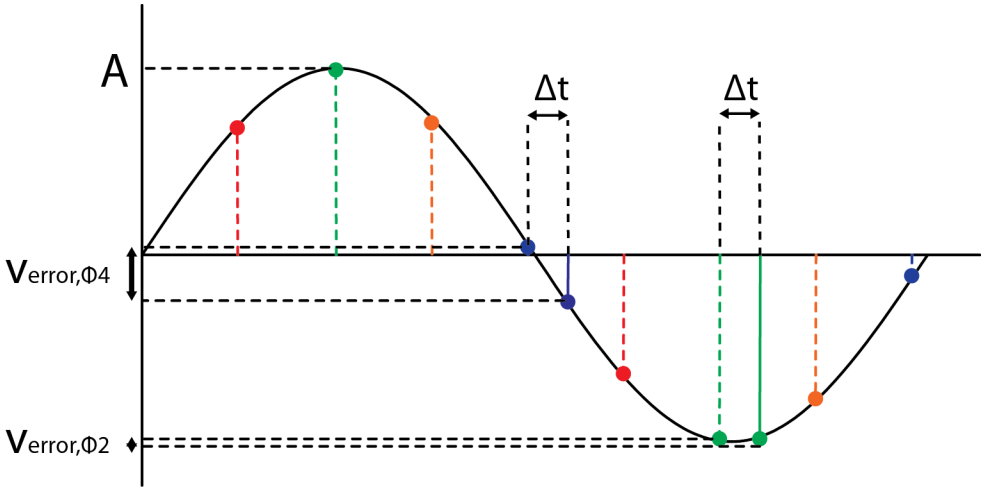

Figure 3.11: Timing error on sampling clock translates to a voltage error proportional to the slope of the signal.

The voltage error is proportional to the slope of the signal and is given by [29]:

$$v_{err}(t) = \frac{ds(t)}{dt} \cdot \sigma_{jitter} \tag{3.5}$$

where s(t) is the input signal being sampled and σ_{jitter} is the rms-jitter on the the sampling clock. The signal output of the sampler will consist of the desired signal, s(t), and the error signal, $v_{err}(t)$. Assuming a single sinusoid at the input:

$$\begin{aligned} v_{out,DS}(t) &= s(t) + v_{err}(t) && \text{(3.6a)} \\ &= A\cos(2\pi f_{in}t) + A2\pi f_{in}\sigma_{jitter}\sin(2\pi f_{in}t) && \text{(3.6b)} \end{aligned}$$

To obtain the SNR, we take the ratio of the power in the desired signal to power of the error signal. Note that in the time domain, the error signal is a sinusoid of amplitdue $A2\pi f_{in}\sigma_{jitter}$ and is phase shifted by 90° w.r.t. the input signal. Intuitively, this makes sense since the error signal should be largest at the zero crossings of the input sinusoid. Taking the ratio of signal powers from Equation 3.6b, we get:

$$\frac{S}{N} = \frac{\left(\frac{A}{\sqrt{2}}\right)^2}{\left(\frac{A2\pi f_{in}\sigma_{jitter}}{\sqrt{2}}\right)^2} = \frac{1}{(2\pi f_{in}\sigma_{jitter})^2} \tag{3.7}$$

Or equivalently in dB:

$$SNR = -20\log(2\pi f_{in}\sigma_{jitter}) \quad [\text{dB}] \tag{3.8}$$

This is the well-known classical result that says that for a given jitter on the sampling clock, σ_{jitter}, the SNR of a jitter-limited system degrades 20dB/decade w.r.t. input frequency. In order to avoid operating in the jitter-limited regime, the SNR due to jitter (Equation 3.8) needs to be greater than the SNR due to quantization noise (Equation 3.4). As an example, to obtain 6b of resolution for a 25GHz signal the rms-jitter would have to be less than 123fs-rms.

The above analysis provides the mathematical evaluation of the SNR at the *sampling instant*. In order to evaluate the SNR of the entire MTS chain, we can take a two-step approach. First, we determine the expression for the analog signal at the ouput of the low-pass filter, $v_{LPF}(t)$. Next, we substitute s(t) with $v_{LPF}(t)$ in Equation 3.6b and follow the analysis used above for evaluating SNR since $v_{LPF}(t)$ is now the input of our baseband sampler.

At the output of the mixer, we have a multiplication between our input sinusoid and the LO with phase noise. For simplicity, we choose the amplitudes of the input signal and LO in such a way that the baseband amplitude becomes unity. Thus, we have at the output of the mixer:

$$\begin{aligned} v_{mix}(t) &= v_{in}(t)\cdot v_{LO}(t) && \text{(3.9a)}\\ &= \cos(2\pi f_{in}t)\cdot 2\cos(2\pi f_{LO}t+\phi(t)) && \text{(3.9b)}\\ &= \cos(2\pi(f_{in}-f_{LO})t-\phi(t))+\cos(2\pi(f_{in}+f_{LO})t+\phi(t)) && \text{(3.9c)} \end{aligned}$$

The lowpass filter will filter the components at the sum frequency, $f_{in}+f_{LO}$, and thus at the output of the LPF we have:

$$\begin{aligned} v_{LPF}(t) &= \cos\left(2\pi(f_{in}-f_{LO})t-\phi(t)\right) && \text{(3.10a)}\\ &= \cos\left(\phi(t)\right)\cos(2\pi f_{IF}t)+\sin\left(\phi(t)\right)\sin(2\pi f_{IF}t) && \text{(3.10b)}\\ &= \cos(2\pi f_{IF}t)+\phi(t)\sin(2\pi f_{IF}t) && \text{(3.10c)} \end{aligned}$$

where $f_{IF}=|f_{in}-f_{LO}|$ is the baseband intermediate frequency (IF) and we have used the small angle approximations $\cos(\phi(t))\approx 1$ and $\sin(\phi(t))\approx\phi(t)$.

Observing Equation 3.10c, we see that after lowpass filtering we have the desired baseband signal along with the LO phase noise modulated up to the IF frequency. In essence, the phase noise profile of the LO signal has been transferred to the baseband signal (Figure 3.12).

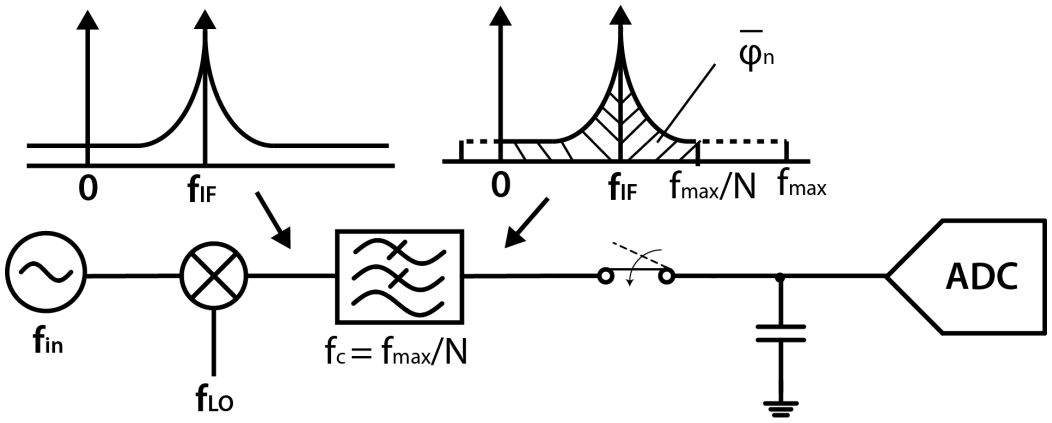

Figure 3.12: Mixing of input signal and LO with phase noise.

Equation 3.10c is now the expression for the input signal to the sampler. The sampler in the MTS system samples at the baseband Nyquist rate of $2 \cdot f_{IF,max} \approx 2 \cdot \frac{f_{in,max}}{2N}$. Using Equation 3.5 and Equation 3.6b, and substituting s(t) in these equations with $v_{LPF}(t)$ from Equation 3.10a, we can determine the signal at the output of the sampler:

$$v_{out,MTS}(t) = v_{LPF}(t) + \frac{\delta(v_{LPF}(t))}{\delta t} \cdot \sigma_{jitter} \tag{3.11a}$$

$$= \cos(2\pi f_{IF}t) + \phi(t)\sin(2\pi f_{IF}t) + \sigma_{jitter}\left[2\pi f_{IF} - \frac{\delta\phi(t)}{\delta t}\right]\sin(2\pi f_{IF}t - \phi(t)) \tag{3.11b}$$

The first two terms on the right hand side of Equation 3.11b are directly from the v_{LPF} expression and represent the baseband signal (first term) along with the additive noise introduced during the mixing process (second term). The last term is the voltage error caused by sampling a non-zero slope signal using a clock with jitter.

One thing to note at this point is that the presence of the second term is new as compared to Equation 3.6b. There, we only had two terms – the desired signal and the error due to jitter. Here, even if we sample the signal with a perfect clock ($\sigma_{jitter} = 0$), we will still have the downconverted phase noise that will degrade the SNR at the output of the sampler. For the DS case (Equation 3.6b), $\sigma_{jitter} = 0$ would result in an infinite SNR (i.e. no SNR degradation due to sampling) and the system performance would then be dictated by the quantization noise. Thus, it is apparent that the presence of this second term plays a crucial role in determining the net benefit of the MTS architecture over the conventional DS architecture and we will return to this discussion. But first, we can further simplify Equation 3.11b by analyzing the various parts of the third term.

The third term represents the voltage error due to jitter on the sampling clock and, as discussed earlier, is proportional to the radian frequency of the signal being sampled. The

two expressions in brackets are measures of the radian frequency; the first term is the absolute frequency of $2\pi f_{IF}$ while the second term represents the uncertainty in instantaneous frequency due to the phase noise. As such, the second term can be thought of as the effective widening of the spectrum that occurs when phase noise is present on a tone. For general oscillators, this widening is generally contained to a very small fraction of the absolute frequency because it would otherwise render it useless in any communication system. This is because the spectrum of the oscillator should not spill over into adjacent channels since it would lead to downconversion of out-of-band signals, thereby corrupting the baseband signal. Therefore, this second term is typically much less than the first term and can be ignored. Equation 3.11b then becomes:

$$v_{out,MTS}(t) = \cos(2\pi f_{IF}t) + \phi(t)\sin(2\pi f_{IF}t) + \sigma_{jitter}2\pi f_{IF}\sin(2\pi f_{IF}t - \phi(t)) \tag{3.12}$$

Using the identity $\sin(u-v) = \sin(u)\cos(v) - \cos(u)\sin(v)$ the third expression in Equation 3.12 can be expanded:

$$= \sigma_{jitter}2\pi f_{IF}\left[\sin(2\pi f_{IF}t)\cos(\phi(t)) - \cos(2\pi f_{IF}t)\sin(\phi(t))\right] \tag{3.13a}$$
$$= \sigma_{jitter}2\pi f_{IF}\left[\sin(2\pi f_{IF}t) - \phi(t)\cos(2\pi f_{IF}t)\right] \tag{3.13b}$$

The second term in (3.13b) will produce a term similar to the second term in (3.12). Both terms represent a modulated version of $\phi(t)$ up to the IF frequency. However, the additional factor of $2\pi f_{IF}\sigma_{jitter}$ in (3.13b) will make its term less significant and we can therefore ignore it for the purpose of simplifying the analysis[8]. Plugging the first part of (3.13b) into (3.12), we arrive at our final expression for the signal at the output of the sampler:

$$v_{out,MTS}(t) = \cos(2\pi f_{IF}t) + \phi(t)\sin(2\pi f_{IF}t) + \sigma_{jitter}2\pi f_{IF}\sin(2\pi f_{IF}t) \tag{3.14}$$

In order to calculate the SNR of this signal we need to evaluate the power of the desired signal (first term) and the noise/error signals (last two terms). The powers of the first and third terms are straightforward to calculate (same calculation as in (3.7)). As for the second term, $\phi(t)$ is a stochastic process that can be fully described by it's frequency domain representation, $\mathscr{L}(f)$, which is the SSB phase noise spectrum commonly used to characterize oscillator phase noise. Using Parseval's theorem, we can calculate the power of this signal by integrating in the frequency domain [30][31]:

$$\overline{\phi_n^2} = 2\cdot\int_0^\infty 10^{\mathscr{L}(f)/10}\,\mathrm{d}f \quad [\text{radians}^2] \tag{3.15}$$

The rms value of the phase noise is then given by:

$$\overline{\phi_n} = \sqrt{2\cdot\int_0^\infty 10^{\mathscr{L}(f)/10}\,\mathrm{d}f} \quad [\text{radians}] \tag{3.16}$$

[8] As an example, for $f_{IF} = 1.5GHz$ and $\sigma_{jitter} = 1ps\ rms$, $2\pi f_{IF}\sigma_{jitter} \approx 10^{-2}$ which is enough to make this term insiginicant when compared to the second term in (3.12)

In both equations, the SSB phase noise, $\mathscr{L}(f)$, is first converted to linear units before being integrated. The factor of two in both equations is needed to capture integration across both sidebands since $\mathscr{L}(f)$ is a SSB spectrum. Now that we have the power of each term in (3.14), we can solve for the SNR of the MTS architecture:

$$\frac{S}{N} \approx \frac{1}{(2\pi f_{IF}\sigma_{jitter})^2 + \left(2 \cdot \overline{\phi_n^2}\right)} \tag{3.17}$$

Comparing (3.17) to (3.7), there are apparent similarities and differences. Ignoring, the $\overline{\phi_n^2}$ term, the two equations are almost identical. This, of course, comes as no surprise since they are modeling the same nonideality - voltage error due to non-zero slope. The only difference here is that the f_{in} in (3.7) has been replaced by f_{IF}. Since f_{in} can be as high the full Nyqusit frequency, f_{max}, while f_{IF} is bounded to roughly $\frac{1}{N}$ of this maximum frequency[9], we see that there is an immediate gain in SNR of N^2. This is the desired result we were targeting. We have bounded the frequency of the signal that is subject to the jitter on the sampling clock. As a result, the resultant rms voltage error due to jitter at the sampling instant has been significantly reduced and we can now obtain larger SNR after sampling in the presence of a given σ_{jitter}.

Unfortunately, this is not the end of the story due to the presence of the $\overline{\phi_n^2}$ which takes away from our immediate N^2 SNR gain. Due to the presence of the $\overline{\phi_n^2}$ term, we take a penalty in SNR for low input frequencies and the DS architecture actually outperforms the MTS architecture. This is because for values of $f_{in} \leq f_{IF}$, the $2\pi f_{IF(in)}\sigma_{jitter}$ terms are equal in both SNR equations and thus the $\overline{\phi_n^2}$ term knocks the SNR of the MTS architecture below that of the DS architecture.

As f_{in} is increased to higher frequencies, (3.7) continues to decrease without bound, while the worst-case value for (3.17) is:

$$\frac{S}{N} \approx \frac{1}{(2\pi f_{IF,max}\sigma_{jitter})^2 + \left(2 \cdot \overline{\phi_n^2}\right)} \tag{3.18}$$

where $f_{IF,max}$ is the maximum IF frequency ($\approx \frac{f_{max}}{2N}$). Thus, there is a crossover frequency where the MTS SNR becomes larger than the DS SNR and the SNR gain continues to increase with increasing f_{in}. Since $\overline{\phi_n^2}$ acts as a "one-time hit", the higher the value of this integrated phase noise, the higher the crossover frequency past which it makes sense to use the MTS architecture over the DS (Figure 3.13). This means that the phase noise of the LO cannot be ignored in the design of the MTS architecture and care should be taken to ensure that the crossover frequency is low enough such that the MTS architecutre make sense for the given application. For example, if the application requires digitization of signals up to 10GHz yet the crossover frequency is 9GHz, then it may not make much sense to implement the MTS architecture since it's providing very marginal gains for frequences between 9 and

[9]Assuming that the wideband spectrum is dissected to N sub-bands of equal bandwidths.

10GHz, while providing worse SNR for frequencies below 9GHz. On the other hand, if the crossover frequency was somehow designed to be 1GHz, then the MTS architecture is more attractive since it is providing significant SNR gains for a larger fraction of the frequencies ($\sim 2-10$GHz) while sacrificing the SNR of a smaller fraction of frequencies ($DC-1$GHz).

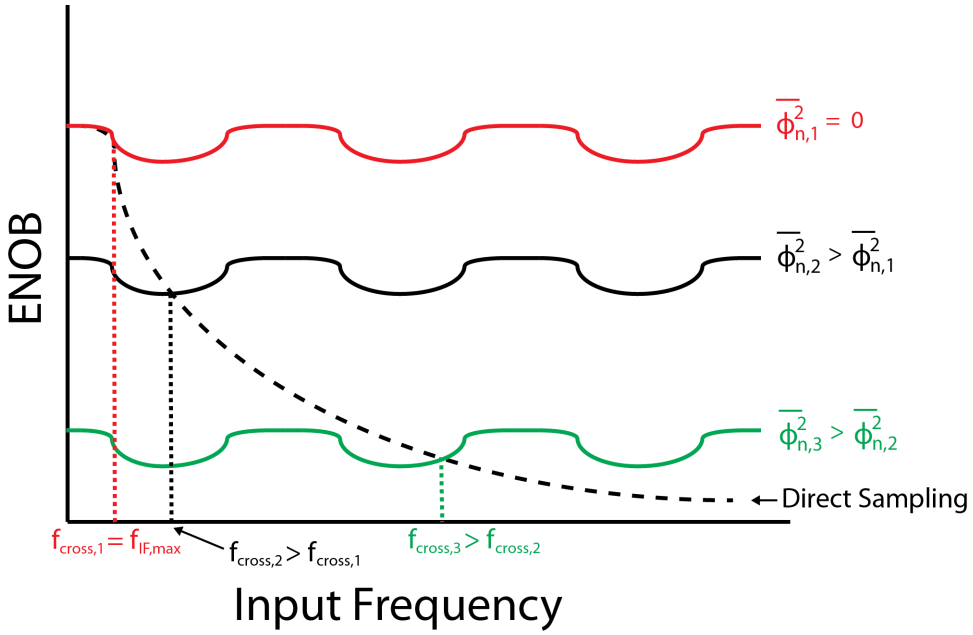

Figure 3.13: Impact of $\overline{\phi_n^2}$ on crossover frequency.

Discussion on $\overline{\phi_n^2}$

Since a lower $\overline{\phi_n^2}$ makes the MTS more attractive over the DS architecture, it is of interest to know how to lower the value of this term. At this point, it is important to note that the $\overline{\phi_n^2}$ in (3.17) is not directly calculated according to (3.15). The difference is that for $\overline{\phi_n^2}$ in (3.17) we should only integrate over the bandwidth of the baseband which is determined by the lowpass filter. Phase noise above the LPF corner frequency will be filtered out and will have negligible contribution to the integral. This is a key point, since it enables the MTS architecture to have any chance against the DS architecture.

To understand this better, consider the relationship between σ_{jitter} and $\overline{\phi_n^2}$. The jitter on the sampling clock is related to the phase noise on the sampling clock by [31]:

$$\sigma_{jitter,s} = \frac{\overline{\phi}_{n,sample}}{2\pi f_s} \tag{3.19}$$

where $\overline{\phi}_{n,sample}$ is the rms phase noise on the sampling clock calculated according to (3.16). The limit of integration in this case will be determined by the bandwidth of the sampler,

which for the DS sampling case must be equal to the full bandwidth of the input signal, f_{max}.

$$\bar{\phi}_{n,sample} = \sqrt{2 \cdot \int_0^{f_{max}} 10^{\mathscr{L}_s(f)/10}\,\mathrm{d}f} \quad \text{[radians]} \tag{3.20}$$

where $\mathscr{L}_s(f)$ is the SSB phase noise spectrum of the sampling clock given in dBc/Hz. We can similarly write an equation for the jitter on the LO used for downconversion:

$$\sigma_{jitter,LO} = \frac{\bar{\phi}_{n,LO}}{2\pi f_{LO}} \tag{3.21}$$

where $\bar{\phi}_{n,LO}$ is given by:

$$\bar{\phi}_{n,LO} = \sqrt{\int_0^{\frac{f_{max}}{2N}} 10^{\mathscr{L}_{LO}(f-f_{IF})/10}\,\mathrm{d}f} \quad \text{[radians]} \tag{3.22}$$

where $\mathscr{L}_{LO}(f - f_{IF})$ is the phase noise spectrum of the LO (in dBc/Hz) centered around the IF frequency, f_{IF}[10]. The factor of 2 before the integral is omitted since we are now integrating across both sidebands of the phase noise spectrum – the part of the spectrum that falls within the baseband bandwidth [29]. In Equations 3.16 and 3.20, the phase noise spectrum centered around 0Hz, $\mathscr{L}_{LO}(f)$, could be used since we wanted to calculate the *total* phase noise power from $-\infty$ to ∞ and the frequency shift was immaterial since we were still integrating across the entire spectrum. Here, the frequency shift matters since we only want to capture the phase noise power that is contained within the baseband bandwidth. For practical calculations, the integral in Equation 3.22 should be broken into two integrals:

$$\bar{\phi}''_{n,LO} = \sqrt{\int_0^{f_{IF}-f_\Delta} 10^{\mathscr{L}_{LO}(f-f_{IF})/10}\,\mathrm{d}f + \int_{f_{IF}+f_\Delta}^{\frac{f_{max}}{2N}} 10^{\mathscr{L}_{LO}(f-f_{IF})/10}\,\mathrm{d}f} \quad \text{[radians]} \tag{3.23}$$

Where $2 \cdot f_\Delta$ is the width of the phase noise spectrum around the downconverted signal that is dominated by *phase* perturbations and not *amplitude* noise. This is sometimes referred to as the close-in phase noise, often exhibiting a $\frac{1}{f^3}$ or greater slope. The *phase* perturbations affect the spectral resolution of the system, but do not contribute to the SNR degradation and should not be included in the calculation [31]. This is illustrated in Figure 3.14.

Now, if we assume that the sampling clock and the LO are derived from a common source with jitter, $\sigma_{jitter,source}$, via the use of frequency multipliers and/or dividers, and further assume that these multipliers/dividers are noiseless, then the *total* jitter on each clock is equal to the source jitter and therefore to each other.

$$\sigma'_{jitter,source} = \sigma'_{jitter,s} = \sigma'_{jitter,LO} \tag{3.24}$$

[10]Recall that f_{IF} varies with input frequency. f_{LO} is fixed.

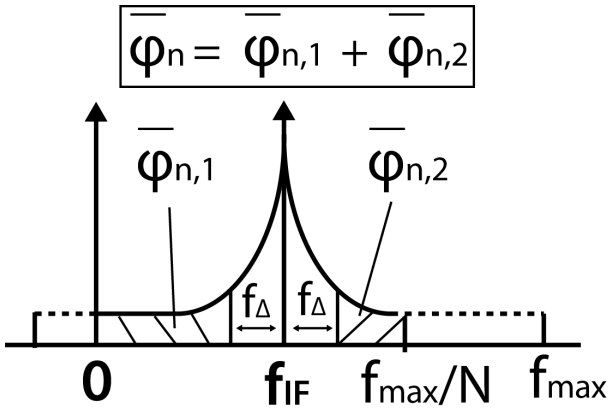

Figure 3.14: LO phase noise power w/o close-in phase noise.

where primed notation is used to indicate the *total* integrated jitter from 0 to ∞:

$$\sigma'_j = \frac{\sqrt{2 \cdot \int_0^\infty 10^{\mathscr{L}(f)/10}\,\mathrm{d}f}}{2\pi f} \tag{3.25}$$

Using (3.24) and the *primed* versions of (3.19) and (3.21) we can write an expression relating the phase noise of the LO to the jitter on the sampling clock:

$$\bar{\phi}'_{n,LO} = 2\pi f_{LO} \cdot \sigma'_{jitter,s} \tag{3.26}$$

Before we continue, let's take a step back to remind ourselves that these primed variables, which denote integration from 0 to ∞, are representative of the MTS sytem *without* a LPF in the baseband. Without the LPF, the limit of integration is no longer set by the IF bandwidth and is instead integrated up to the highest frequency with non-neglible phase noise. Generally, this limit would be f_{max} since that's the highest bandwidth any node in either architecture has to achieve, but we lose nothing by integrating to ∞. The key takeway here is that without the LPF, all of the phase noise that was integrated in the DS architecture to determine jitter, is also integrated when determining the noise added during the mixing process. Immediately, it becomes apparent that, with baseband LPFs omitted, it is implausible that the MTS architecture can outperform the DS architecture since we are subjecting ourselves to the phase noise *twice* – once at the mixing node and again at the sampler. This is seen by plugging (3.26) into (3.17) and simplifying:

$$\frac{S}{N}_{,noLPF} \approx \frac{1}{(2\pi\sigma'_{jitter})^2 \cdot (f_{IF}^2 + 2f_{LO}^2)} \tag{3.27}$$

Comparing (3.27) to (3.7) and plugging in the best-chance scenario of $f_{in} = f_{LO} + \frac{f_{in,max}}{2N}$[11] it can be shown that (3.27) is always less than (3.7), except for a subset of frequencies in the

[11]For a given channel in the MTS architecture, the DS SNR is minimized at the edge frequency.

first channel. Another takeaway from (3.27) is that the SNR degrades as we go to higher frequency channels (larger f_{LO}). We will later see that this general behavior holds with the LPF present[12]. Further analysis of this "LPF-omitted" scenario is not done here since it is not representative of the true MTS architecture.

In the real MTS architecture, the LPF is a integral part of the operation and is what allows the escape from the grim prospect of re-introducing all of the phase noise at the mixer. With the LPF, (3.24) no longer holds since the limits of integration are given according to (3.20) and (3.22) and we can now write:

$$\sigma_{jitter,LO} < \sigma_{jitter,s} \tag{3.28}$$

The key variable in determining the magnitude of this inequality, is the number of channels, N. With a larger number of channels, the limit of integration in (3.22) becomes smaller, resulting in a smaller $\overline{\phi}_{n,LO}$ and hence a smaller "one-time" hit in SNR. As a result, we can reduce the crossover frequency and make the MTS architecture beneficial for a larger range of frequencies. Of course other problems arise with increasing N, and there is a practical limit to how large N should be made.

Aside: Using direct-sampling as a basis for comparison

The direct-sampling architecture is a good basis for comparison as it captures the constraints and fundamental performance trends that we are interested in assesing for all Nyquist-rate ADC architectures – bandwidth requirements, ENOB vs. frequency, and noise performance. This is easy to see for single-channel architectures since their front-end will closely resemble that of Figure 3.6(a). Furthermore, one can argue that single-channel sampling sets the upper-bound on the performance of multi-channel architectures (i.e. time-interleaved) with the same effective sample rate, if we neglect power consumption. First, the multi-channel architectures degrade the noise performance as compared to single-channel since there are more components and thus, more noise sources needed to process the signal. The bandwidth requirements of multi-channel is the same as single-channel since the full signal bandwidth needs to be passed to all channels even though the sub-channel sample rate is lower. Thus, the same bandwidth requirement is placed on the front-end T/H amplifier. Lastly, if we assume a perfect quantizer, then the ENOB performance is determined by the jitter of the sampling clock and the maximum input frequency, both of which are the same for multi-channel and single-channel architectures. Therefore, the DS architecture is a proper basis for comparison.

[12] Assuming certain conditions on the LO phase noise.

System-Level Simulations

Simulation Setup

In order to assess the performance FI-ADC as compared to the TI-ADC and other DS architecures, a system level simulation was performed using Keysight's SystemVue software. A block diagram of the models used for comparison is shown in Figure 3.15.

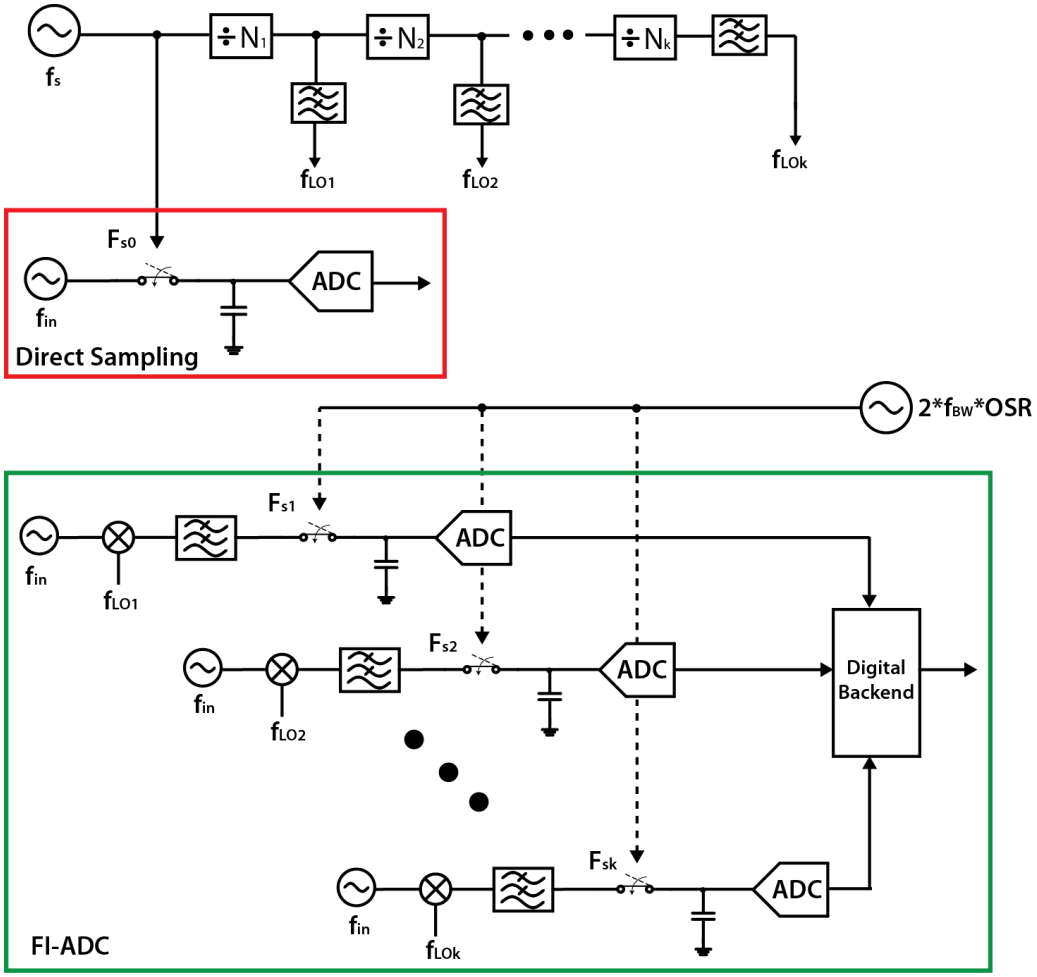

Figure 3.15: Block diagram of models used for system simulation.

For the DS architecture, a model similar to that of Figure 3.6 is used. The sampling frequency is chosen to be 50GS/s for an input Nyquist frequency of 25GHz. An 8-bit ADC is used to sample and quantize the input signal. For the FI-ADC system, an 8-channel architecture is chosen with passband bandwidths of 3GHz. The IF frequency is cut down to 1.5GHz by using complex (I/Q) downconversion (not shown in Figure 3.15). The LO frequencies are integer multiples of 3GHz up to 24GHz (i.e. 3,6,9...21,24). The baseband filters are 6^{th} order Butterworth filters with a corner frequency of 2GHz. The output of

each LPF is oversampled (2.35x) and quantized by an 8-bit ADC[13]. The I/Q outputs of each channel is then sent to the digital backend to perform signal reconstruction and output a single stream of data. An FFT is then performed on the output streams of each system and the SNRs at each output are compared.

The clocks used for sampling and downconversion are all related in a manner similar to what is depicted in Figure 3.15. The only difference being that all sources have uncorrelated phase noise whereas Figure 3.15 suggests that the phase noises of the clocks are correlated. The LO chain begins with the master frequency of 50GHz used to drive the ADC in the DS architecture. The phase noise of this sampling clock is defined such that it integrates to a desired σ_{jitter}. Next, we want to assess if the FI-ADC architecture can perform better in the presence of the same σ_{jitter}. Given this, we derive the phase noise of each subsequent LO frequency such that it's phase noise integrates to the same σ_{jitter}[14]. The assumption here is noiseless frequency division. The jitter on the sampling clocks for the FI-ADC are also equal to the same σ_{jitter}. Lastly, low Q bandpass filters (2^{nd} order, Q=4) are placed on the LO paths before driving the mixers. These filters are representative of the resonant tanks that would typically be used in the LO chain to drive the mixer LO ports with large swing sinusoids.

Simulation Results

Given this simulation setup, the input frequency to each system was swept and the output SNR analyzed for a single-tone excitation. This sweep was performed with $\sigma_{jitter} = 200$fs and $\sigma_{jitter} = 400$fs. The results from this simulation are shown in Figure 3.16. In this plot, the jitter-limited behavior of the DS architecture (dotted lines) is immediately apparent. The SNR degrades at a rate of 20dB/decade for frequencies well past the corner frequency.

For the MTS architecture (solid lines), the SNR does degrade slightly as the input frequency is increased, but the degradation is not as severe as the DS case. At the Nyqusit rate, we see a 15dB improvement in SNR over the DS architecture which is equivalent to 2.5b of added ENOB. Note that the SNR plot exhibits the periodic dipping as predicted by the theory and shown in Figure 3.10[15,16]. However, unlike Figure 3.10, the SNR response is not flat across the entire Nyqusit band. This is because LO phase noise during mixing was not considered for that earlier figure. Equation 3.27 captured this dependence of MTS architecture's SNR on the LO frequency for the case where no baseband LPF was included and this dependence holds true even when the LPF is included even though the new governing equation may be slightly different than Equation 3.27.

[13] See Section 4 to better understand how these design parameters were chosen.

[14] In order to keep σ_{jitter} constant as frequency is scaled by $\frac{1}{M}$, the integral needs to be scaled by $\frac{1}{M^2}$. This is accomplished by decreasing $\mathscr{L}(f)$ by $20\log(M)$.

[15] A coarse logarithmic frequency sweep was used in the simulation and is why the exact shape of Figure 3.10 is not captured in the plot.

[16] For very low frequencies, the two architectures should have the same SNR. The higher MTS SNR at low-frequencies is due to the 2.35x oversampling ratio. This oversampling gain only affects the quantization-noise-limited region ($f <\sim 2.5$GHz).

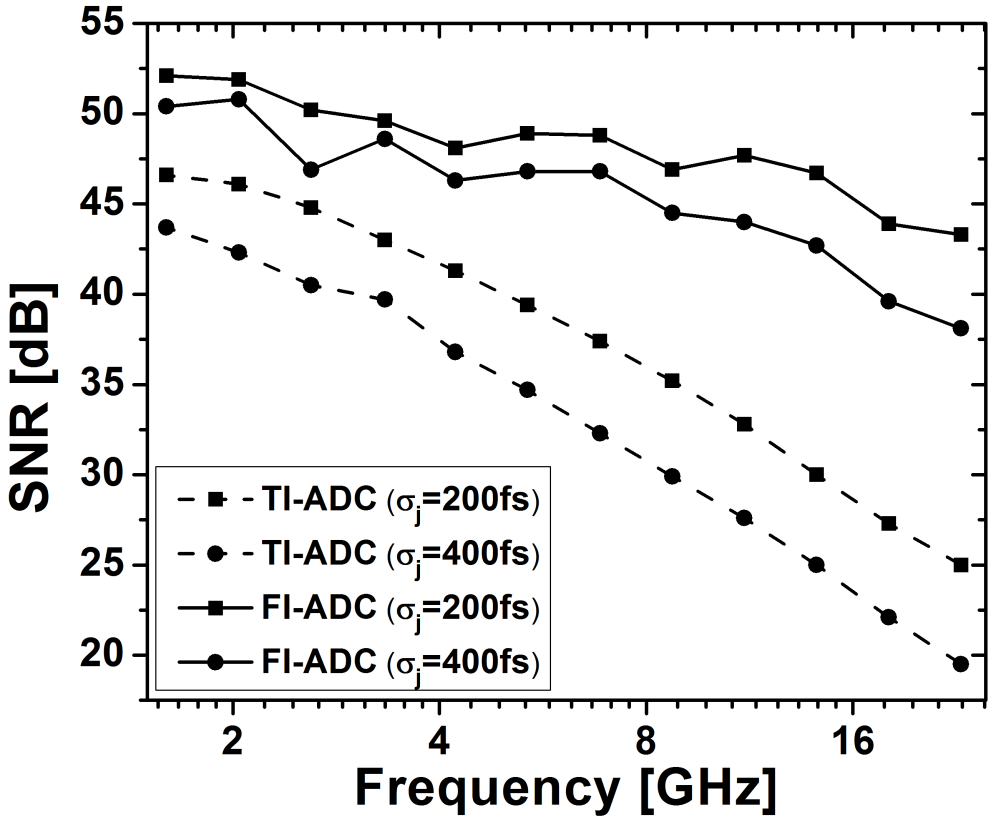

Figure 3.16: Comparision of SNR vs. f_{in} for different values of σ_{jitter}.

The dependence of SNR on f_{LO} is due to the increased phase noise levels for higher LO frequencies. From Equation 3.25, it is easy to see that if two clocks have the same rms-jitter, $\sigma_{j,1} = \sigma_{j,2}$, and one clock frequency is half of the other, $f_{clk,1} = 0.5 \cdot f_{clk,2}$, then the SSB phase noises must have the relation: $\mathscr{L}_1(f) = \mathscr{L}_2(f) - 20 \cdot \log(2) = \mathscr{L}_2(f) - 6\text{dB}$. In general, if $f_{clk,1} = \frac{f_{clk,2}}{M}$, then the SSB phase noise relationship is given by [32]:

$$\mathscr{L}_1(f) = \mathscr{L}_2(f) - 20 \cdot \log(M) \tag{3.29}$$

As a result of Equation 3.29, the higher frequency channels will have a higher phase noise level and hence a higher $\overline{\phi_n^2}$ value in Equation 3.17 thereby lowering the SNR. Figure 3.17 shows the simulation results for the MTS architecture with and without LO phase noise. Without phase noise, we see that the ENOB is relatively flat across the entire Nyqusit band (focusing on the jitter-limite regime, $f > 2.5$GHz) and more closely matches Figure 3.10.

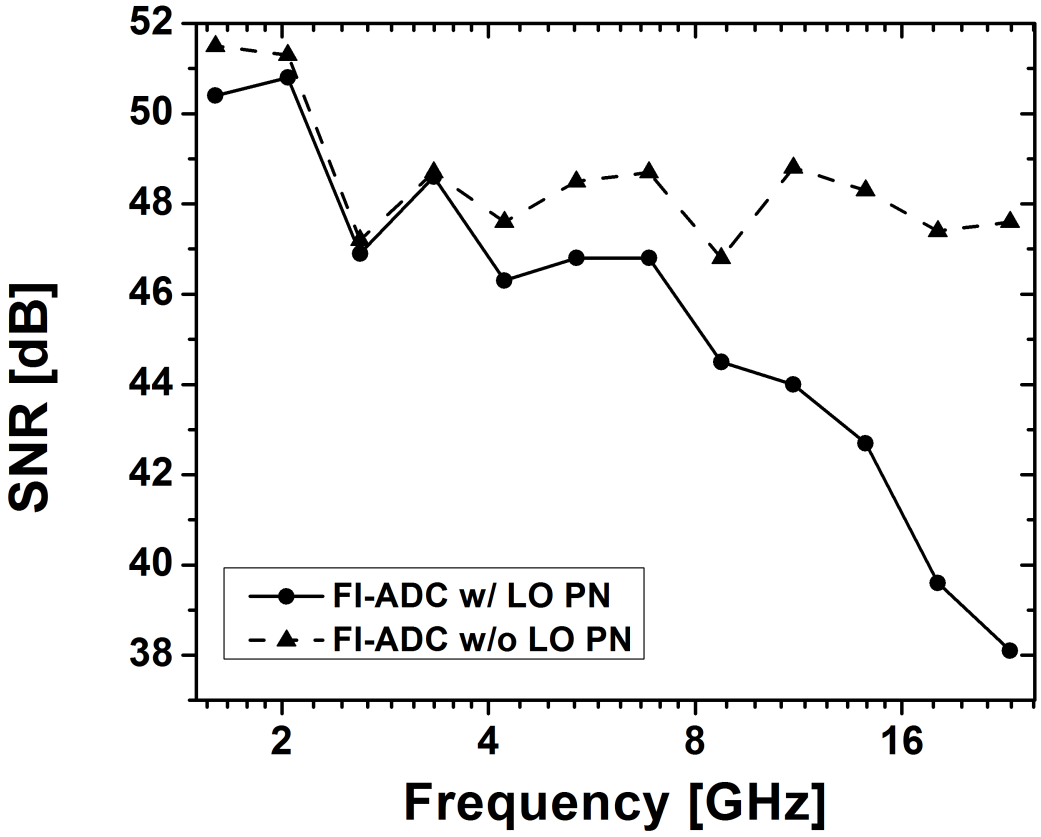

Figure 3.17: MTS SNR vs. f_{in} with and without LO phase noise.

The above simulation results were for a system that was representative of the designed and measured FI-ADC implementation (Chapter 4). In particular, the bandpass filters (which were included on the LO distribution chain to model the tuned buffers used for LO distribution) partially filter out the wideband phase noise of the LO and help decrease the $\overline{\phi_n^2}$ term for each channel (Equation 3.17). It is of interest to see how the MTS system performs when each channel is subjected to the full LO phase noise spectrum without any second-order bandpass filtering. This, afterall, is a more fair comparison to the DS system which did not have any LO filtering. Figure 3.18 plots the new comparison plot for the two systems.

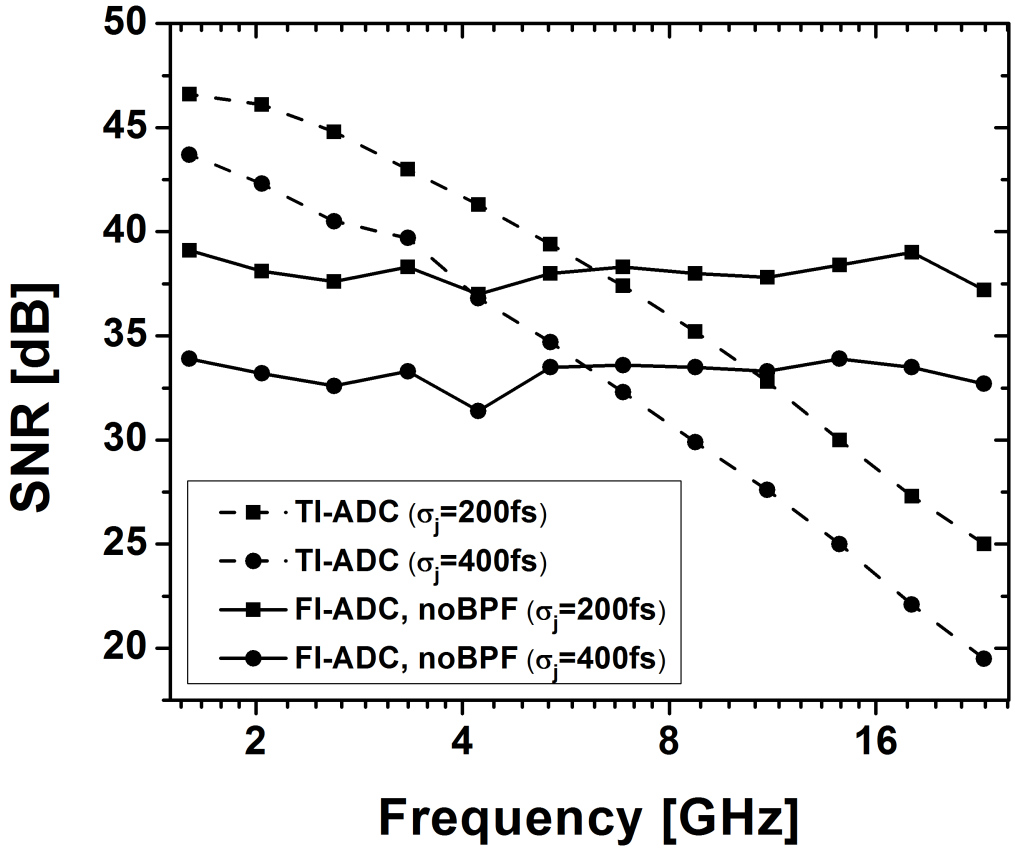

Figure 3.18: Comparision of SNR vs. f_{in} without LO BPF in the MTS system.

The plots of Figure 3.18 now resemble that of Figure 3.13. Looking at the TI-ADC $\sigma_j = 200$fs plot and the two plots for the FI-ADC, we see the predicted trend of higher crossover frequency as $\overline{\phi_n^2}$ is increased (higher σ_j is equivalent to higher integrated phase noise). It is also interesting to note that the crossover frequencies for the pair of plots ($\sigma_j = 200$fs and $\sigma_j = 400$fs) are identical.

Now, following the previous discussion of how SNR is dependent on f_{LO}, the plot of Figure 3.18 begs the question of why isn't there any frequency dependence exhibited? The answer is reciprocal mixing. The wideband phase noise of the LOs in each channel will mix with the input signal regardless of which channel is responsible for processing that signal. This wideband phase noise is downconverted to baseband and, if large enough, will be quantized by the ADC and create a noise spectrum at the output of each of channel. Note that these channels which don't process the input signal, should normally be off. Thus, the output spectrum will include the input signal tone, the noise of the excited channel *and* the noise of all other channels. Because all channels are always contributing to the noise floor of

the output spectrum, the output noise floor is constant and output SNR is independent of input frequency. When BPFs were used to model the LO distribution, the far-out wideband phase noise was sufficiently filtered such that only the excited channel and its adjacent channels contributed to the output noise spectrum. Channels that were further away didn't have large enough phase noise to "wake-up" their sleeping ADCs since the reciprocal mixing would occur with the wideband phase noise that was already sufficiently filtered by the BPF. As such, only a subset of channels contribute to the output noise floor and this subset is dependent on the input frequency, leading to the SNR dependence on f_{LO} that was seen in Figure 3.16.

These simulation results show that the answer to the question of which architecture is better (TI or FI) isn't straightforward. In general, the FI architecture has great potential to outperform the TI architecture at very high input frequencies. Depending on the LO generation/distribution (e.g. whether or not BPFs are used), the FI architecture may only provide gains for frequencies which are above the crossover frequency. When this is the case, the integrated phase noise power introduced during mixing, $\overline{\phi_n^2}$, needs to be minimized as much as possible in order to make the FI architecture more appealing. Fortunately for the design presented in Chapter 4, the usage of tuned LO buffers was appropriate (and necessary). This helps reduce the impact of the wideband phase noise and in turn results in the FI architecture outperforming its TI counterpart. In the next chapter, an overview of practical design considerations for FI-ADCs will be presented along with the design of a 25GHz channelized analog front-end (AFE).

Chapter 4

Design of a 50 GS/s 6-bit FI-ADC

The target application for this project was pulsed-radar imaging applications [4]. In these systems, a pulse-modulated carrier signal is transmitted to a target and the reflected signal is detected. Based on the reflection properties (amplitude and time-of-arrival) measured at various carrier frequencies, the dielectric properties of the material can be determined [2]. In order to improve the depth resolution of these imagers, smaller pulse-widths must be used. In [3], pulse-widths below 30ps were demonstrated in 0.13μm SiGe technology, resulting in I/Q baseband bandwidths up to 20 GHz. In order sufficiently digitize these bandwidths for proper detection, an ADC with an effective sample rate of 50GS/s and 6b resolution is needed.

This chapter discusses the design of a 50GS/s 6b frequency-interleaved ADC. The focus of the design was on the analog front-end (AFE) which is responsible for the wideband distribution and channelization of the input signal. However, since the AFE, sub-ADC, and digital back-end specifications are tightly coupled with one another, a system-level design approach was initially performed in order to determine the necessary specifications of the AFE.

This chapter begins with a discussion of system-level design considerations. The design of the AFE is then presented, focusing on the circuit design of key building blocks.

4.1 System-Level Design

Number of Channels

The first key design parameter in the FI-ADC is the number of channels, as it has a direct impact on every design block in the system. For example, a larger number of channels results in a smaller processing bandwidth per channel which means the sub-ADCs can sample at much lower rates and be designed to be very energy-efficient. In addition, the low-pass filter cutoff frequencies are reduced, which can greatly reduce design complexity. On the other hand, a larger number of channels increases the amount of channels requiring harmoinc

	Pros	Cons
Wideband Distr.		- Lower BW due to higher capacitive loading
Mixer		- More channels require HRMs - Larger # of harmonics need to be cancelled
LPF	- Lower corner frequency - Feasibility of higher-order filters	
Sub-ADC	- Lower sampling rate. - More energy-efficient - Easier to oversample	
LO Gen.		- Larger # of LO frequencies to generate/distribute
DSP	- Lower data rate on each path - Reduced FIR filter complexity	- Larger # of input datapaths - Larger # of channels to calibrate

Table 4.1: Pros and Cons for larger number of channels.

rejection mixers, as discussed in Chapter 3, adding to design complexity. Furthermore, the wideband distribution becomes more difficult due to the increase in capacitive loading that needs to be driven. Table 4.1 summarizes some of the major design parameters impacted.

Assuming power consumption is to be minimized, determining the true optimal number of channels requires accurate modeling of all the blocks in the design as well as a modeling of how the design parameters of each block are coupled. In order to avoid this arduous task, a more heuristic approach was taken. There are a few observations considered in this huerisitic decision process:

1. The replicated sub-blocks (sub-ADC and LPF) should operate at frequencies far enough from the f_T of the process such that the power vs. frequency trade-off is linear rather than exponential. Once in this linear tradeoff region, further increase in the number of channels does not result in as much power savings since the component count may also increase linearly, resulting in zero power savings, yet added system complexity.

2. Op-Amp RC filters are limited to practical corner frequencies of 100s of MHz. Other LPF topologies (g_m-C, current-mode) are limited to frequencies below 10GHz in the process of interest.

3. There is an upper-limit to the number of stages in the DA due to losses on the transmission line. Larger number of stages leads to longer T-lines and hence, increased losses. This added loss can overtake the gain being added by the additional stages, resulting in no increase in DA gain and a waste of power [33].

4. Distributing LOs is a significant challenge in regards to the area needed for routing and the coupling due to proximity of the traces. As such, it is desirable to minimize the number of LOs needed to generated/routed.

Taking these points into consideration, it is best to minimize the number of channels used to the point where the sub-ADCs are energy-efficient ($<$ 75fJ/conv-step) and the LPF corner frequencies are no more than a few GHz. As such, a 9-channel (1 baseband, 8 passband) architecture was chosen. Each of the passband channels process a bandwidth of 3GHz which results in baseband I/Q bandwidths of 1.5GHz. The baseband channel processes frequencies from DC-1.5GHz. The LO frequencies that need to be generated are 3, 6, 9, 12, 15, 18, 21, and 24GHz. Figure 4.1 depicts the processing bandwidths of each channel.

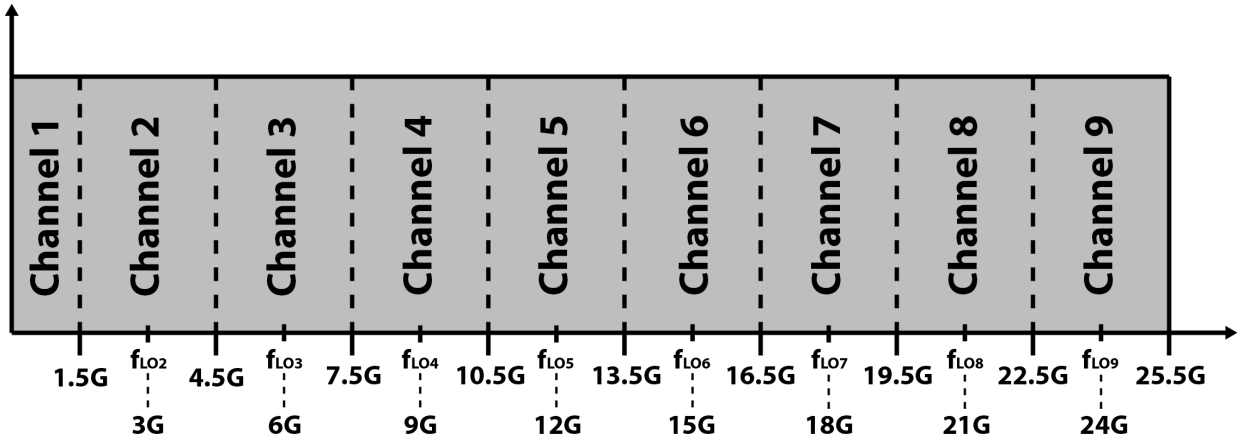

Figure 4.1: Processing bands for each channel.

Channel Filtering - Analog vs. Digital

Adequate filtering is an integral part of the FI-ADC. In [5][16][18][19], the analog filters are responsible for channelizing the input signal. These filters are designed to precisely determine the edge frequencies of each channel. Due to process variations, the edge frequencies of the analog filters will deviate from their ideal values. This causes complications in the design of the digital synthesis filters used for reconstruction since they require apriori knowledge of the response of the analog filters [5][34].

In order to relax the system's sensitivity to the analog filters in the front-end, the job of setting channel bandwidth is transferred to digital filters. Figure 4.2 shows the proposed digital backend architecture which is a slightly modified version of Figure 3.2. In this architecture, the LPFs in the AFE perform the first level of filtering, however their corner frequencies are set to be slightly larger than the processing bandwidth of each channel. For example, the baseband processing bandwidth is 1.5GHz while the LPF corner frequency would be set to $\sim$ 3GHz. Thus, in the AFE, each channel will handle signals from the lower

edge of the succeeding channel and the upper edge of the preceeding channel. Lowpass FIR filters are then used as the first stage in the DSP for setting the proper corner frequency and filtering out the adjacent channels' signals[1]. The advantage of this approach is that channel edge frequencies are no longer set by the analog filters which are subject to process variations, but are instead determined by digital filters with well-defined frequency responses.

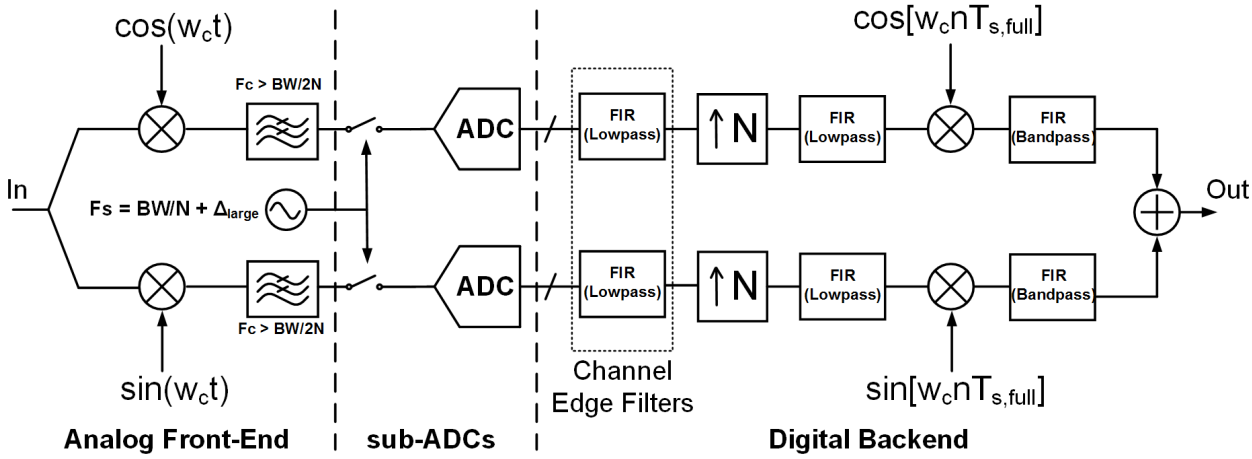

Figure 4.2: Proposed architecture for digital backend.

Using this approach, the sub-ADCs now have to sample slightly faster due to the increased analog bandwidth at its input. The SNDR performance of these sub-ADCs only needs to be maintained over the channel processing bandwidth and can degrade in the excess bandwidth region that overlaps into the adjacent channel, alleviating the design of this higher-speed ADC. In addition, this increase in sample rate is already accommodated by the need for oversampling in order to reduce the magnitude of aliased signals, as will be discussed in the following section.

Analog Filtering

In addition to first-level channel filtering, the analog LPFs also function as anti-aliasing filters for the sub-ADCs. Like conventional ADC designs, there is a tradeoff between the required order of the filter and the ADC oversample ratio (OSR). For a targeted ENOB, the aliased signals need to be sufficiently below the noise floor as to not impact the SNDR. There are two ways to mitigate aliasing: 1) oversampling and 2) high-order lowpass filters (anti-aliasing filters). Oversampling increases the frequency that is first susceptible to aliasing, pushing it further into the attenuation band of the anti-aliasing filter. Using a higher order

[1]Although the channel filtering could be absorbed into the interpolation filter that is used for upsampling, it is better to seperate the two so that the channel FIR filter can be done at a lower sample rate with lower power consumption.

filter provides larger attenuation at aliasing frequencies due to the sharper rolloff of the filter. Figure 4.3 illustrates the tradeoff.

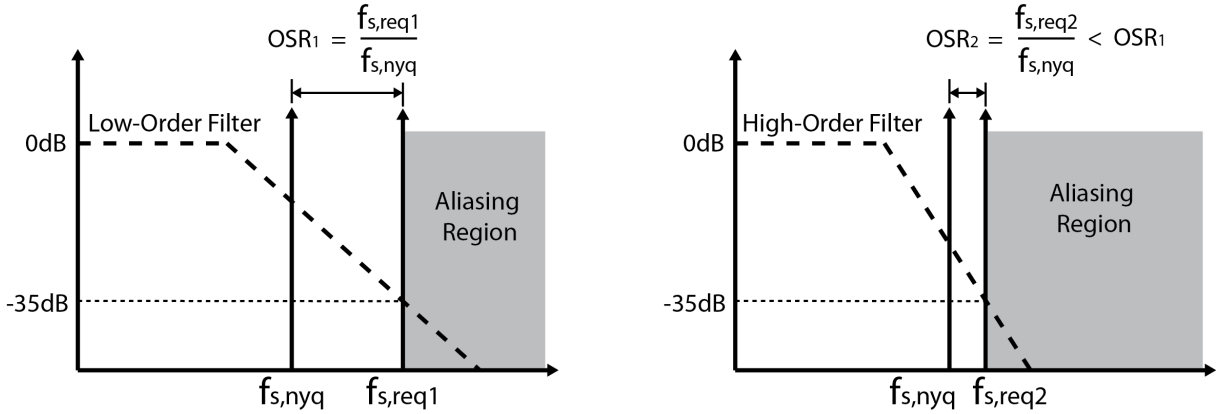

Figure 4.3: Tradeoff between filter rolloff and oversampling ratio.

Generally, there is a practical limit to the order of baseband analog filters, with 5^{th} or 6^{th} order being an upper-limit. Given this upper-limit, the attenuation provided by the filter is fixed, and the only way to further reduce the magnitude of aliased signals is through oversampling. Table 4.2 shows the required sample rates of the sub-ADCs for various targeted ENOBs and filter types with a filter bandwidth of 2GHz. The specification is that all aliased signals are at least 3dB below the quantization noise floor (e.g. 5b ENOB is equivalent to an SNDR of $5 * 6.02 + 1.76 = 32$dB, so the aliased signals are all 35dB below the signal level).

The three filter types shown in Table 4.2 are chosen for their superior in-band phase and group delay response. In-band group-delay variation needs to be minimized in order to reduce dispersive effects. This is important in pulsed-radar applications where the pulse-width determines the depth resolution [2]. Dispersion in the medium and/or the receiver causes a widening of the received pulse, thereby reducing the effective depth-resolution of the imager. In addition, the group delay and amplitude of all channels should also match for the case of a wideband signal that is processed by multiple channels. Thus, it is desirable to have a filter that requires a minimal amount of calibration for equalizing the group delay and amplitude response across all channels.

The Bessel filter is known for its maximally flat group delay, yielding the best response of all filter types. Unfortunately, this comes at the price of very poor rolloff characteristics and thus, poor out-of-band attenuation. As a result, the uasge of Bessel filters requires very high oversampling ratios to meet the resolution specifications.

The Butterworth filter provides a maximally flat amplitude response, which is good for minimizing amplitude calibrations. It also provides better rolloff when compared to the Bessel filter, but is still gradual when compared to its Chebyshev (I/II) counterparts. As seen from Table 4.2, a 6^{th} order Butterworth filter would still need to be paired with a 8+GS/s

	Bessel			Butterworth			Inverse-Chebyshev		
	4th	**5th**	**6th**	**4th**	**5th**	**6th**	**4th**	**5th**	**6th**
5 bits	14	12	11.4	9.4	8.5	7.6	7.6	6	5.4
6 bits	16.6	14	13	11.4	9.4	8.4	8.8	6.6	6
7 bits	19.5	16	14.6	13	10.6	9.4	10.6	7.4	6.4
8 bits	22.8	19	16.6	16	12.2	10.4	12.2	8.4	7.1

Table 4.2: Sub-ADC sample rate (in GS/s) vs. Filter order/type.

ADC for a 6b system. ADCs operating at this sample rate have reported efficiencies larger than 100fJ/conv-step [8], which is above the targeted efficiency.

The Inverse-Chebyshev (Chebyshev-II) filter provides an adequate solution to both problems – it provides the sharp rolloff known to the Chebyshev family of filters while maintaining good in-band performance. The group delay of these filters exhibit no ripples in the passband (good for equalization), and can yield a sufficiently flat delay response if the corner frequency is set to be appropriately higher than the max frequency of interest. In addition, the magnitude response of the inverse-Chebyshev filter is flat within the passband. As seen in Table 4.2, using a 6^{th} order inverse-Chebyshev filter drops the required ADC sample rate to 6GS/s for a 6b system. ADCs with FOM less than 100fJ/conv-step have been reported at this sample rate [8].

Analog Front-End Specifications

The first step towards the design of a full 50GS/s, 6b FI-ADC system is the design of the analog front-end. The AFE is responsible for the wideband distribution and channelization of the input signal. Given the previous discussions, a 9-channel front-end architecture was chosen. The targeted input bandwidth is 25GHz. The *processing bandwidth* of each passband channel is 3GHz resulting in I/Q baseband *processing bandwidths* of 1.5GHz. The actual bandwidth of the I/Q baseband is designed to be 2GHz in order to reduce sensitivity to analog filter variations, while digital FIR filters are used to enforce the 1.5GHz channel corner frequency.

In order to meet the SNDR requirements for a 6b system, the noise and distortion power introduced by the AFE must fall 38 dB below the maximum input signal level. The main distortion products of interest are caused by 2^{nd} and 3^{rd} order intermodulation products ($IM_{2/3}$) and harmonic folding. As a result, we require[2] $IM_{2/3} < -40$dB and $HRR > 40$dB. To meet the noise specifications, the required noise figure of the AFE can be calculated using link budget equations similar to wireless receivers.

For the 6b FI-ADC, the maximum input power, $P_{i,max}$, is the input power that yields $IM_{2/3}$ levels equal to the threshold specification (40dB in this example). Assuming that IM_3

[2]Additional margin should be added to these specifications so that the combined power of these undesired signals is 38 dB below the input power.

is the dominant distortion, this means that $P_{i,max}$ can be written in terms of the IIP_3 of the receiver:

$$P_{i,max} = IIP_3 - 20\text{dB} \tag{4.1}$$

Equation 4.1 is derived from the 2 dB/dB relation between IM_3 and input power. Once $P_{i,max}$ is determined, the noise figure of the AFE must be sufficient to yield a minimum detectable signal (MDS) that is lower than $P_{i,max}$.

$$P_{MDS} < P_{i,max} \tag{4.2}$$

P_{MDS} is the minimum input signal power that yields the required SNR at the sub-ADC input. This constraint can be written using the link budget equation:

$$SNR_{ADC} = SNR_{IN} - NF_{AFE} \tag{4.3a}$$
$$SNR_{ADC} = P_{MDS} + 174\text{dBm} - 10 \cdot \log(BW) - NF_{AFE} \tag{4.3b}$$

where SNR_{ADC} is the required SNR at the output of the AFE, BW is the channel *processing bandwidth*, and NF_{AFE} is the noise figure of the AFE. Solving for P_{MDS} and plugging in $SNR_{ADC} = 38dB$ and $BW = 1.5GHz$, we get:

$$P_{MDS} = SNR_{ADC} - 174\text{dBm} + 10 \cdot \log(1.5e9) + NF_{AFE} \tag{4.4a}$$
$$P_{MDS} = 38\text{dB} - 174\text{dBm} + 92\text{dB} + NF_{AFE} \tag{4.4b}$$
$$P_{MDS} = -44\text{dBm} + NF_{AFE} \tag{4.4c}$$

Using Equations (4.1), (4.2), and (4.4c), we arrive at a relation between IIP_3 and NF_{AFE}:

$$NF_{AFE} < IIP_3 - (-24\text{dBm}) \tag{4.5}$$

Equation 4.5 sets an upper bound on the noise figure of the AFE, given its IIP_3 performance. The first observation is that it immediately sets a lower bound on the IIP_3 of the AFE since NF of any receiver must be greater than 0dB. Assuming a practical lower bound on NF of 4dB, the lower bound on IIP_3 becomes:

$$IIP_3 > -20\text{dBm} \tag{4.6}$$

This lower bound can be easily met, but it should be noted that it can become more constraining for higher resolution systems (larger SNR_{ADC}) and/or implementations with a fewer number of channels (larger channel processing bandwidth, BW). Another observation from Equation 4.5 is that a more linear front-end relaxes the noise performance requirement. This provides another handle for optimizing the AFE design - by maximizing the linearity of the linearity-limiting block (the mixer in this design), the required gain of the front-end

DA can be reduced. Since the DA runs from a higher supply voltage, power can be saved. In addition, the lower gain required from the DA can yield increased DA bandwidth.

At this point, it becomes apparent that Equations (4.5) and (4.6) drive the design methodology for the AFE, as will be discussed in the following section.

4.2 AFE Circuit Design

The first phase of this design focused on the analog front-end. The AFE consists of the distributed amplifier, harmonic rejection mixers, lowpass filter, and 8-frequency LO generation. Figure 4.4 shows the block diagram of the AFE architecture that was designed. The input signal is applied to a distributed amplifier that provides front-end gain and wideband distribution to the various channels. Each stage of the DA is tapped by a single channel which contains a transconductor, passive mixer and lowpass filter. Channels 2 and 3 utilize harmonic rejection mixers since the harmonics of their LO frequencies fall within the input bandwidth.

The following sections discuss the designs of each of the building blocks shown in Figure 4.4. The first block that was designed was the high-frequency harmonic rejection mixer. The mixer dominates the linearity performance of the AFE. Once the linearity and NF of the mixer was optimized, the required gain of the distributed amplifier was calculated to set an appropriate NF_{AFE} in accordance with Equation 4.5.

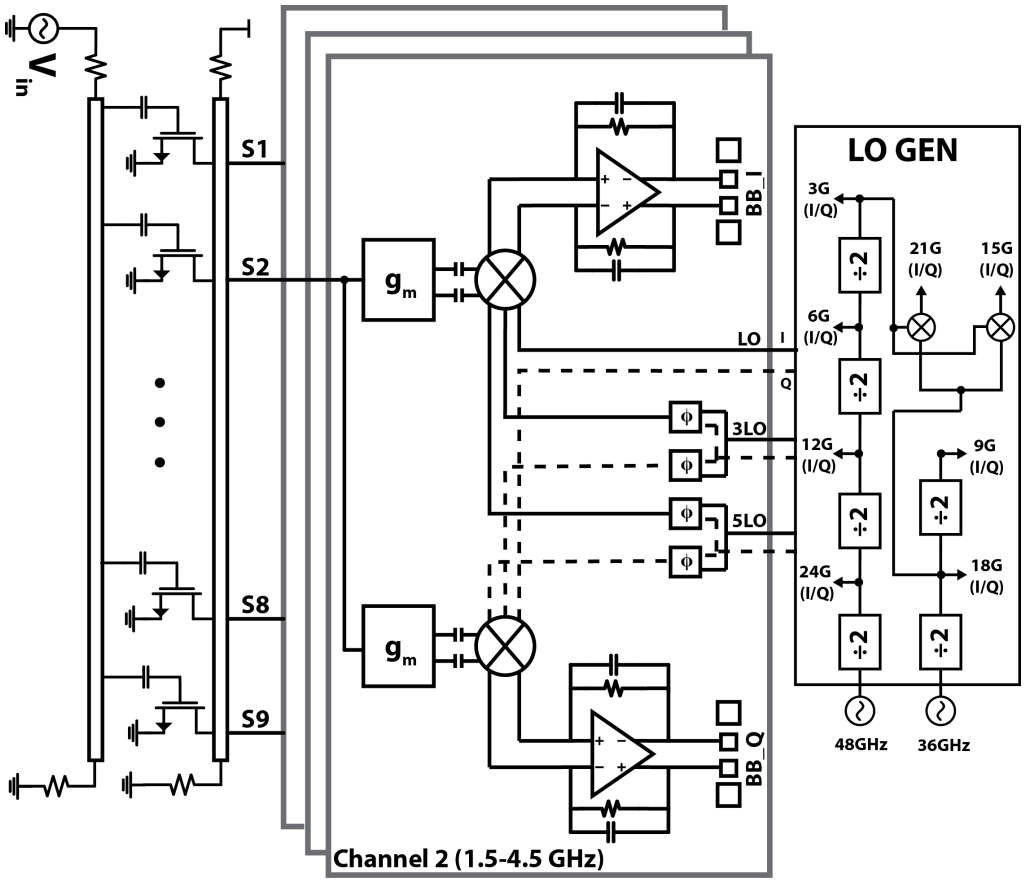

Figure 4.4: FI-ADC AFE Block Diagram.

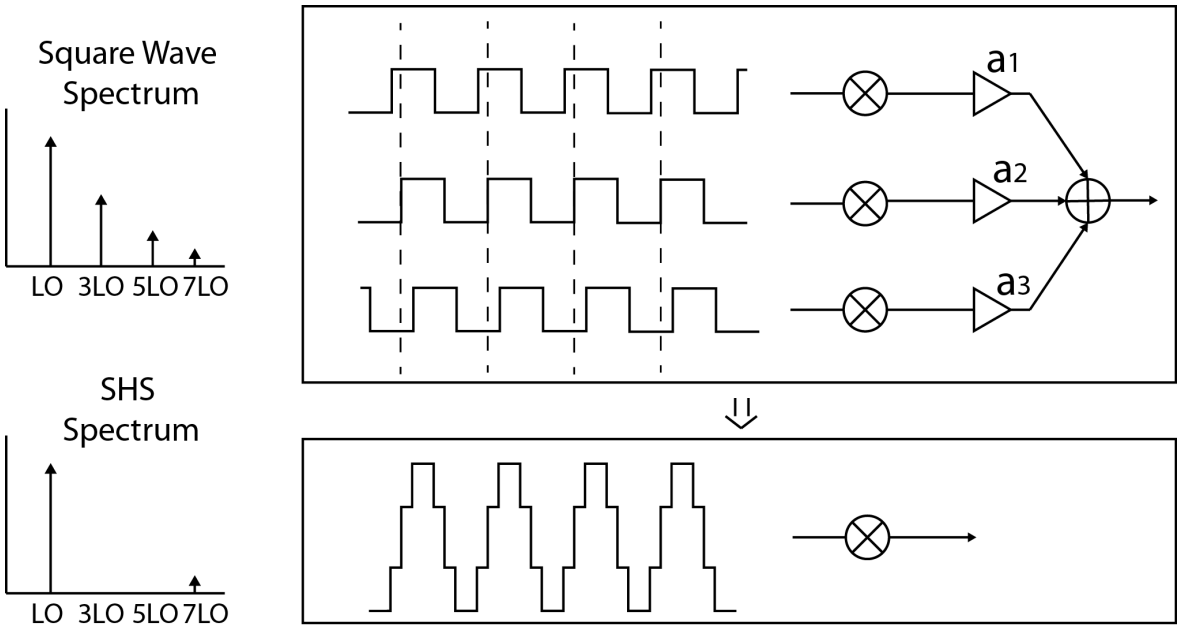

Figure 4.5: Conventional method for harmonic rejection.

Wideband Harmonic Rejection Mixer

Chapter 3 discussed the harmonic folding problem. In this design, channels 2 and 3 are susceptible to harmonic folding. For channel 2, the 3^{rd}, 5^{th}, and 7^{th} harmonics of its 3GHz LO all fall within the 25GHz input bandwidth and can downconvert signals that are out of its processing bandwidth. For channel 3, only the 3^{rd} harmonic of its 6GHz LO falls within the input bandwidth. Due to the high LO frequencies (3GHz and 6GHz), conventional harmonic rejection techniques cannot be easily adopted. In the conventional HRM architecture, complex digital logic is used to generate various non-overlapping clock phases. These clock signals then feed the gain-weighted mixing paths whose outputs are summed. This process emulates multiplication by a sampled-and-held sinusoid (SHS) which has reduced harmonic content (Figure 4.5). Designing the digital logic to run at these frequencies is challenging and can consume considerable amounts of power. As an example, in [26], a clock of $N \cdot LO$ was needed in order to achieve harmonic rejection for harmonics up to (N-1). In our system, for 3^{rd} and 5^{th} harmonic rejection in Channel 2, we would need to run the digital logic at 18GHz, which is not practical at this technology node. Thus, there is a need for a high-frequency HRM architecture that is capable of obtaining greater than 40dB of HRR.

Figure 4.6 shows the schematic of the proposed high-frequency HRM that was introduced in Chapter 3. The main and auxiliary mixers each contain identically-sized passive switching quads. Independent gm-stages are used to generate the AC current inputs to each switching

quad[3]. Identically-sized switching quads are used for improved matching of the gain ratio between the two paths. The gain ratio is set by the strengths of the gm-stages.

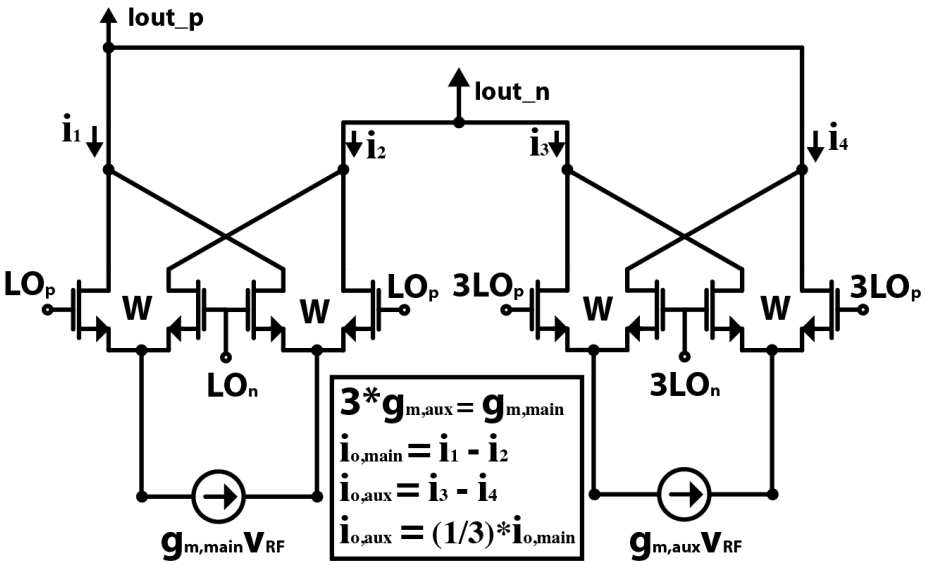

Figure 4.6: Schematic of proposed HRM.

An alternative implementation is shown in Figure 4.7. In this implementation, the gain ratio is achieved through relative sizing of the switching quads. Although this architecture offers superior linearity performance (elimination of transconductor stage), it has a few disadvantages when it comes to setting relative gains. First, it is difficult to set the gain ratio since the effective gain of each path isn't solely determined by the ratio of device sizes – parasitic source/drain capacitance and rise/fall time of the LO each change the effective impedances of the switching paths. For instance, with matched rise/fall times at the input of each quad, the rise/fall time observed on the output current waveforms would be different due to the different time constants of each path. By using identical paths, we can eliminate this second-order impact on the gain.

Figure 4.8(a) shows how the attainable HRR is limited by these second-order effects as compared to the identical switching quad topology. The figure shows a plot of the attainable HRR of each topology as a function of the rise/fall time. The rise/fall times at the input of each switching quad are matched for these simulations. As illustrated in the plot, the identical path topology can achieve 20dB larger HRR, on average, than the alternative design. Furthermore, the identical path topology shows very low sensitivity to the input rise/fall times, unlike the alternative design which has a HRR that sharply drops as rise/fall times are increased.

[3]The gm-stages are AC coupled to the mixing quad, as illustrated in Figure 4.4. The current sources in Figure 4.6 only represent an AC signal.

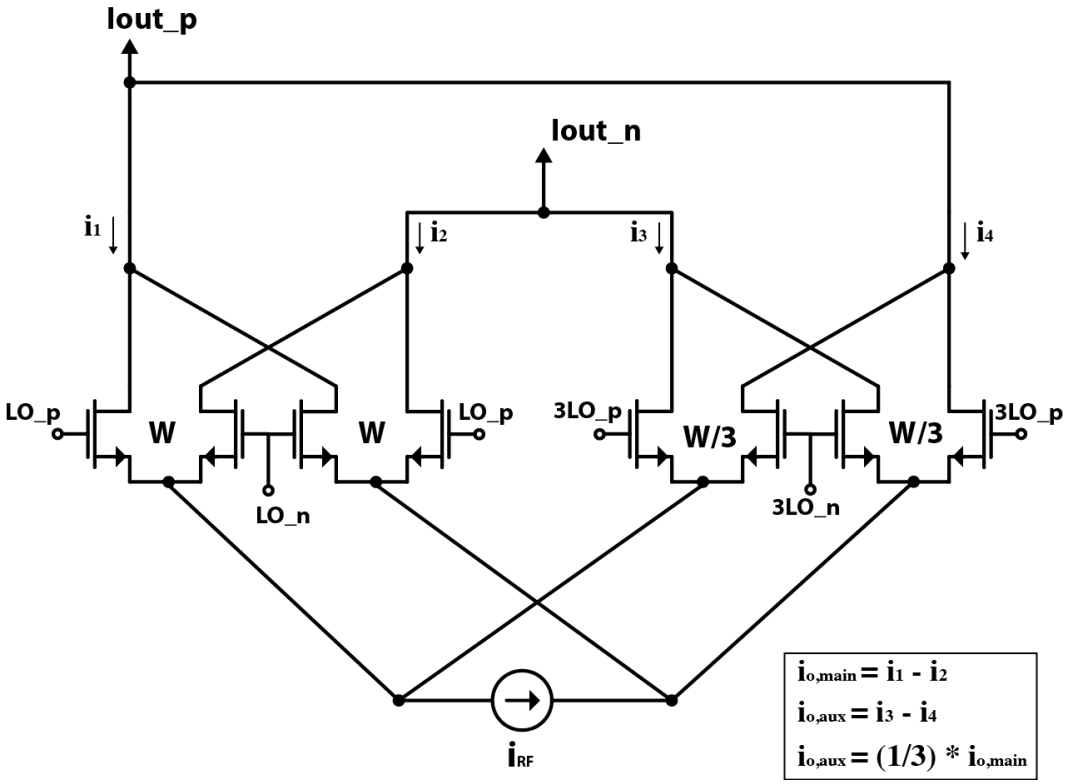

Figure 4.7: Alternative schematic for HRM.

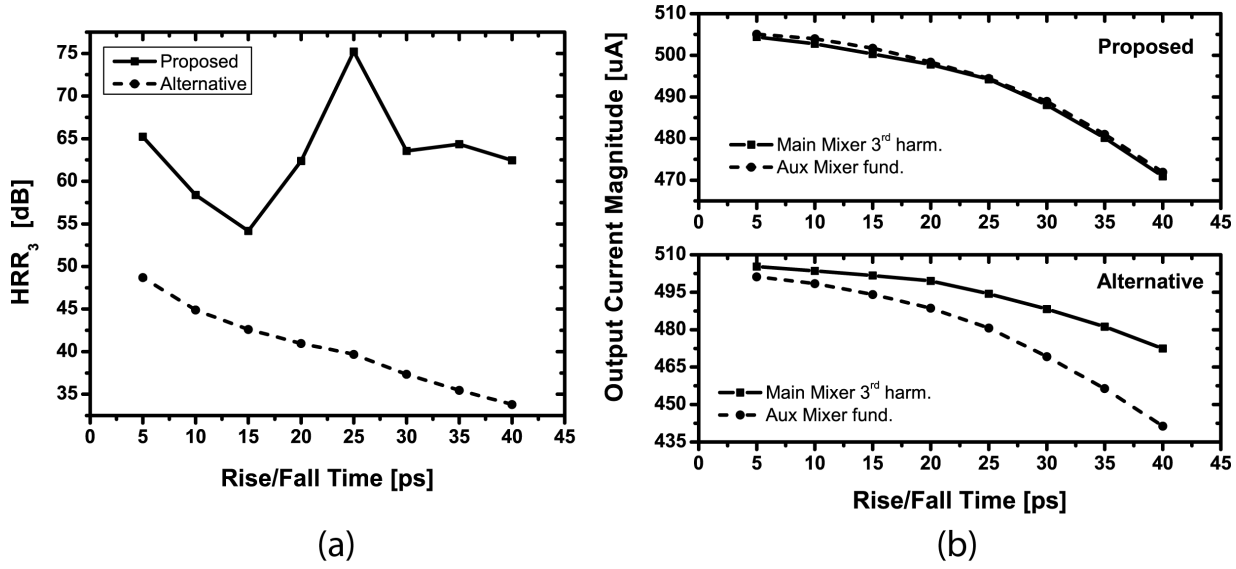

Figure 4.8: Performance comparison of two HRM topologies.

These trends can be explained by Figure 4.8(b) which shows the magnitudes of the harmonic components of interest for the main and auxiliary mixers. For proper harmonic rejection to take place, we require that the magnitude of the 3^{rd} harmonic of the main mixer match the fundamental of the auxiliary mixer. As illustrated in the plot, perfect matching between the two paths is difficult to achieve through device sizing in the alternative approach[4] and causes the attainable HRR to be 20dB lower, on average. Secondly, the magnitudes in the identical path topology exhibit similar behavior as a function of rise/fall time, explaining the previously noted insensitivity. The magnitudes in the alternative topology, however, exhbit poor matching which worsens for increasing rise/fall times. This is because the *non-identical* paths exhibit *non-identical* behaviors as rise/fall time is increased and this causes the attainable HRR to continually degrade.

The other critical advantage of including a gm stage that preceeds the mixing quad is the flexibility to tune the gain of the paths in the presence of process variation. Montecarlo simulations show a mean HRR of 36dB, which means tuning needs to be incorporated in order obtain harmonic rejection levels greater than 40dB. Having a gm stage provides a direct handle on the gain of each path which can be tuned via the tail current source. A high-resolution current DAC can then be used as the tail source in order to achieve the appropriate gain resolution to meet the HRR specification. In the alternative approach, device size would have to be tuned by switching in/out minimum width devices. Since a maximum device width is usually set for keeping the power of the LO driver reasonable, minimum-sized finger widths would likely have to be used in order to obtain the desired resolution. This can lead to a difficult design since large series switches will also be needed in conjunction with the minimum-width devices.

Figure 4.9 shows a schematic of the gm-stage used in the HRM. Two auxiliary stages are used for driving the auxiliary mixers for 3^{rd} and 5^{th} harmonic rejection[5]. An active-load is utilized to increase the output impedance of the gm stage and increase the conversion gain by forcing the signal current to flow into the low-impedance node at the input of the lowpass filter. Since the mixer and lowpass filter chain is operating in current-mode, the signal swing at the output node of the gm stage is small and the common-mode voltage is a (relatively) free parameter that can also be used for optimizing noise/linearity performance. A common-mode feedback (CMFB) amplifier is used to drive the gates of the PMOS active-loads and set the common-mode output voltage. The output common-mode voltage is made tunable by using a votlage DAC to the feed the reference voltage of the CMFB amplifier.

[4]For the alternative topology, the ratio of device sizes was set to yield the maximum HRR at a single rise/fall time of $20ps$. These devices sizes were then used for the various simulation sweeps.

[5]Although channel 2 requires 7^{th} harmonic rejection, it was not implemented in this design

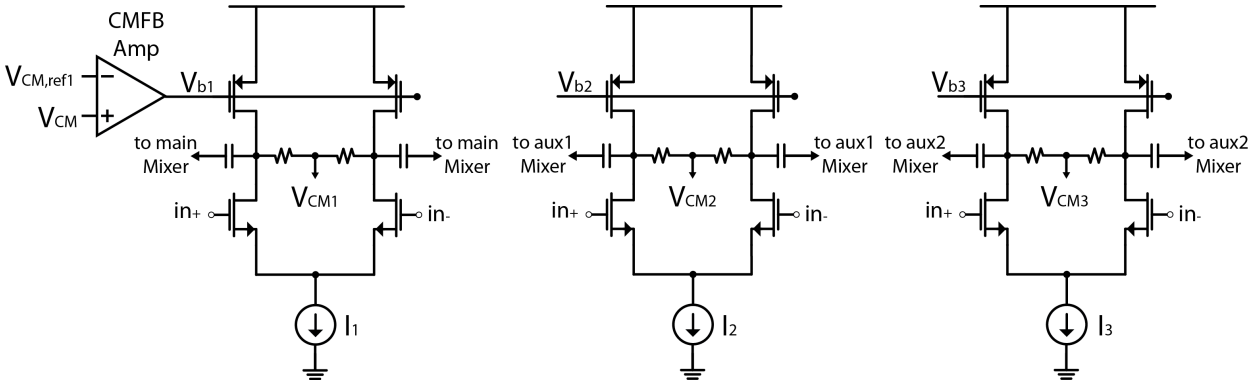

Figure 4.9: Schematic of HRM g_m stage.

Rise/Fall Time Impact on Phase

Once the gain ratios have been appropriately set, the next limiting factor in attainable HRR is the phase match between the two paths. Phase mismatch is primarily caused by non-identical routing on the RF, IF, and LO paths. However, even if very careful layout techniques were used to make all routes nearly identical, there is still another mechanism which limits the attainable HRR – matching of the rise/fall time at the inputs of the main and auxiliary mixers. In the previous simulations (Figure 4.8), the main and auxiliary rise/fall times were set equal. However, as can be seen in Figure 4.10, the attainable HRR is also sensitive to the *matching* of the respective rise/fall times at the input of the main and auxiliary mixers, $t_{r,main}$ and $t_{r,aux}$. In the plot, e_{rise} is the percentage error between $t_{r,main}$ and $t_{r,aux}$ ($e_{rise} = \frac{t_{r,main} - t_{r,aux}}{t_{r,main}}$). The rise time and fall time are equal for each mixer ($t_{rise} = t_{fall}$), and have a nominal value of $30ps$. The fundamental LO frequency is $4GHz$.

The overall sensitivity of the attainable HRR to rise/fall time mismatch can be explained using Fourier analysis. The Fourier series for a square wave with amplitude A, period T, and non-zero rise/fall time, t_r, is given by [35]:

$$C_n \angle \theta_n = A \cdot \frac{sin\left(\frac{n\pi}{2}\right)}{\frac{n\pi}{2}} \cdot \frac{sin\left(\frac{n\pi t_r}{T}\right)}{\frac{n\pi t_r}{T}} \angle - \frac{n\pi}{2} - \frac{n\pi t_r}{T} \tag{4.7}$$

The first sinc expression in Equation 4.7 has values of $\frac{2}{\pi} \cdot [1, \frac{1}{3}, \frac{1}{5}, \frac{1}{7}, ...]$, which are the well known coefficients for a perfect square with zero rise/fall time. The second sinc expression captures the impact that non-zero rise/fall times have on the amplitudes of the fourier coefficients. This sinc shaping of the amplitude as a function of t_r was exhibited in the plots of Figure 4.8(b). Assuming waveforms that have sharp transitions compared to the period of the harmonic ($t_r \ll \frac{T}{n}$), we can write $n\pi t_r \ll T$ and note that $\sin(x) \approx x$ for $x \ll 1$. This sets the second sinc function to unity, and Equation 4.7 simplifies to:

$$C_n \angle \theta_n = A \cdot \frac{sin\left(\frac{n\pi}{2}\right)}{\frac{n\pi}{2}} \angle - \frac{n\pi}{2} - \frac{n\pi t_r}{T} \tag{4.8}$$

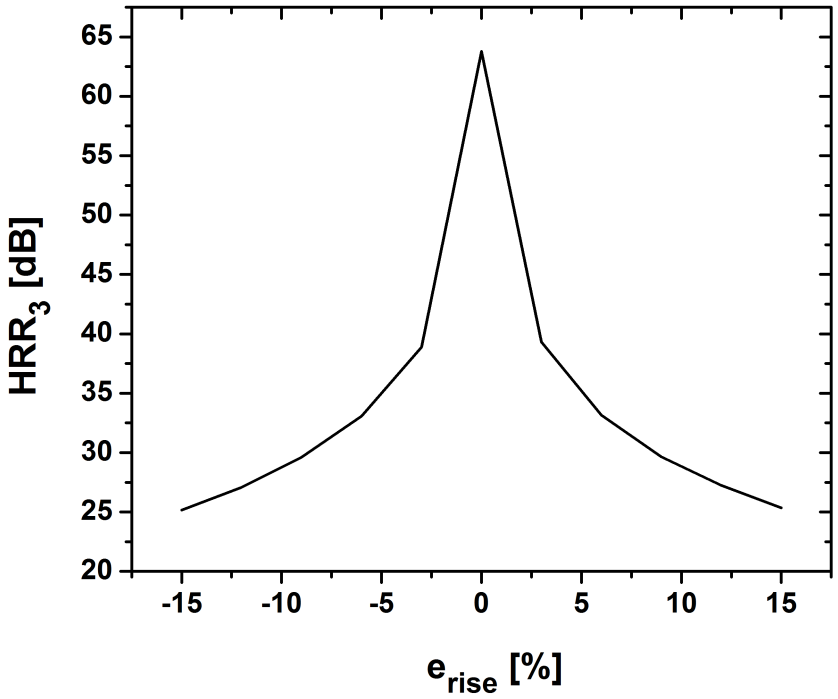

Figure 4.10: Simulated HRR vs. rise/fall time mismatch between main and aux. mixers.

This says the the primary manner in which a mismatch in rise/fall times limits the attainable HRR is by introducing additonal phase shift. Note that the harmonic of interest for the main mixer is the 3^{rd} harmonic ($n = 3$), while the harmonic of interest for the auxiliary mixer is the 1^{st} harmonic ($n = 1$). In both cases, $\frac{n}{T}$ evaluates to the same value since the auxiliary mixer is driven by $3 \cdot LO$ ($T_{main} = 3 \cdot T_{aux}$, where T_{main} and T_{aux} represent the LO period used in Equation 4.8 for the main mixer and auxiliary mixer, respectively). Thus, both harmonics of interest will experience the same phase shift if $t_{r,main} = t_{r,aux}$, resulting in proper harmonic cancellation.

Using Equation 4.8, we can plot the theoretical HRR that results from the phase error introduced when $t_{r,main} \neq t_{r,aux}$ as shown in Figure 4.11. In this figure, the attainable HRR is plotted as function of e_{rise} for $T_{main} = 1$ns ($f_{LO} = 1$GHz) and $T_{main} = 250$ps ($f_{LO} = 4$GHz). The simulated HRR values from Figure 4.10 are also included for comparison. As seen in Figure 4.11, higher sensitivity to rise/fall time mismatch is exhibited by the higher LO fundamental frequency. This is predicted by Equation 4.8 – higher LO frequencies (smaller T_{main}) lead to larger values for $\frac{n}{T}$ and thus a larger phase shift for a given Δt_r. Lastly, Figure 4.11 shows that the phase shift is the dominant non-ideality introduced when $t_{r,main} \neq t_{r,aux}$, as opposed to the amplitude attenuation dictated by the second sinc expression in Equation 4.7. This is exhibited by the close matching between the simulated HRR and the

HRR predicted by Equation 4.8, which only takes phase shift into account.

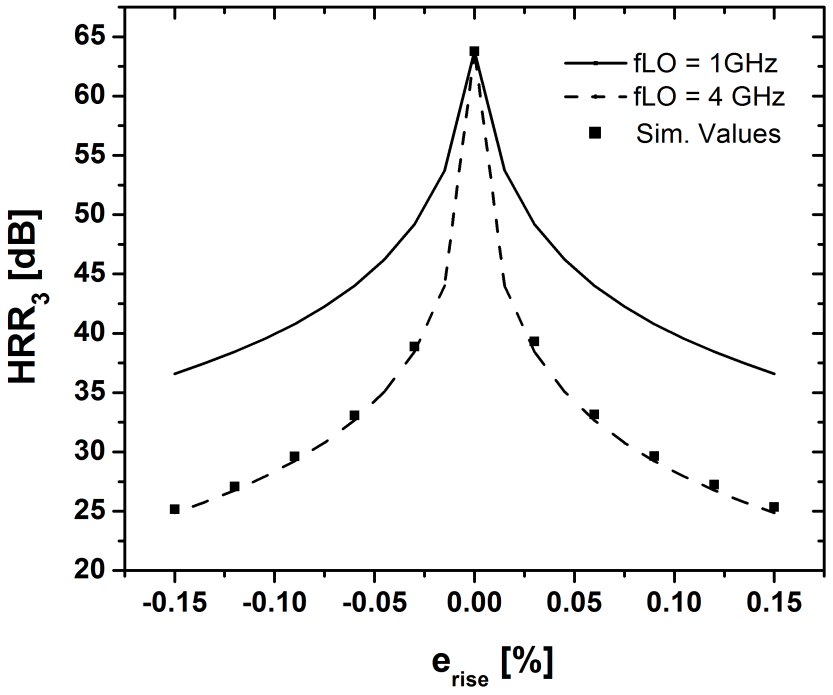

Figure 4.11: HRR vs. e_{rise} as predicted by Equation 4.8.

Figure 4.12 further illustrates how the phase error introduced when $t_{r,main} \neq t_{r,aux}$ is the dominant limitation of attainable HRR. For $+/-15\%$ error, the max amplitude error introduced is 1.6% while the phase error introduced is $+/-10°$. These errors independently translate into max attainable HRRs of 45.5dBand 24.7dB, respectively. Thus, the phase error introduced must be accounted for[6] by using calibration.

In order to calibrate the phase matching of the main and auxiliary paths of the HRM, I/Q phase interpolators [36] are placed on the LO path to the auxiliary mixers as shown in Figure 4.4. A 9b differential current DAC is used for the tail sources of the phase interpolators in order to achieve sufficient resolution over the entire 360° range. With proper tuning, the simulated 3^{rd} and 5^{th} HRRs exceed 40 and 45dB, respectively. The simulated IIP_3 of the HRM was +11dBm. The worst-case NF occurs for the HRM used in channel 2 which utilizes two auxiliary paths (3^{rd} and 5^{th} harmonic rejection) and is equal to 19dB. The simulated NF for the mixer used in channels 4-9 (main mixer only) was 17dB.

[6]In addition to the phase mismatch caused by non-identical routing.

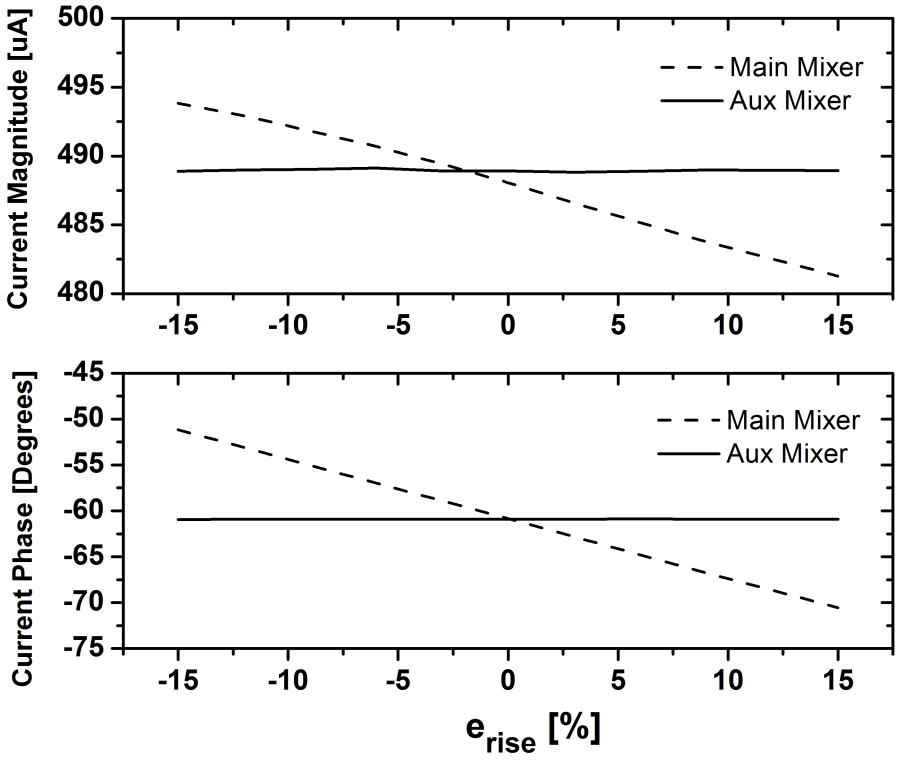

Figure 4.12: Magnitude and phase error caused when $t_{r,main} \neq t_{r,aux}$ ($f_{LO} = 4$GHz).

Multi-Tap Distributed Amplifier

The distributed amplifier (DA) enables wideband signal distribution to the nine channels by absorbing the input capacitance of the channels into the drain transmission line (T-line). A simplified schematic of the 9-stage DA was shown in Figure 4.4. Figure 4.13 shows a more detailed schematic of the DA. In the implemented DA, each gain stage is comprised of a cascoded common-source amplifier for improved stability and reduced resistive loading of the drain T-line. The input of the gain stage is capacitively coupled to the gate T-line. In each stage, a diode-connected mirror device provides dc biasing. The T-lines are formed by discrete inductors and the loading capacitances of each channel and DA gain stage.

Given the noise performance of the HRM, the acceptable gain of the DA was calculated using Equation 4.5 and the cascaded noise factor formula:

$$F_{AFE} = F_{DA} + \frac{F_{HRM} - 1}{G_{DA}} \tag{4.9}$$

where F_{AFE}, F_{DA}, and F_{HRM} are the noise factors of the full AFE, the DA, and the HRM, respectively, and G_{DA} is the power gain of the DA. The IIP_3 of the AFE that appears in Equation 4.5 is also dependent on G_{DA}. Since the DA and LPF operate on a 2.4V supply,

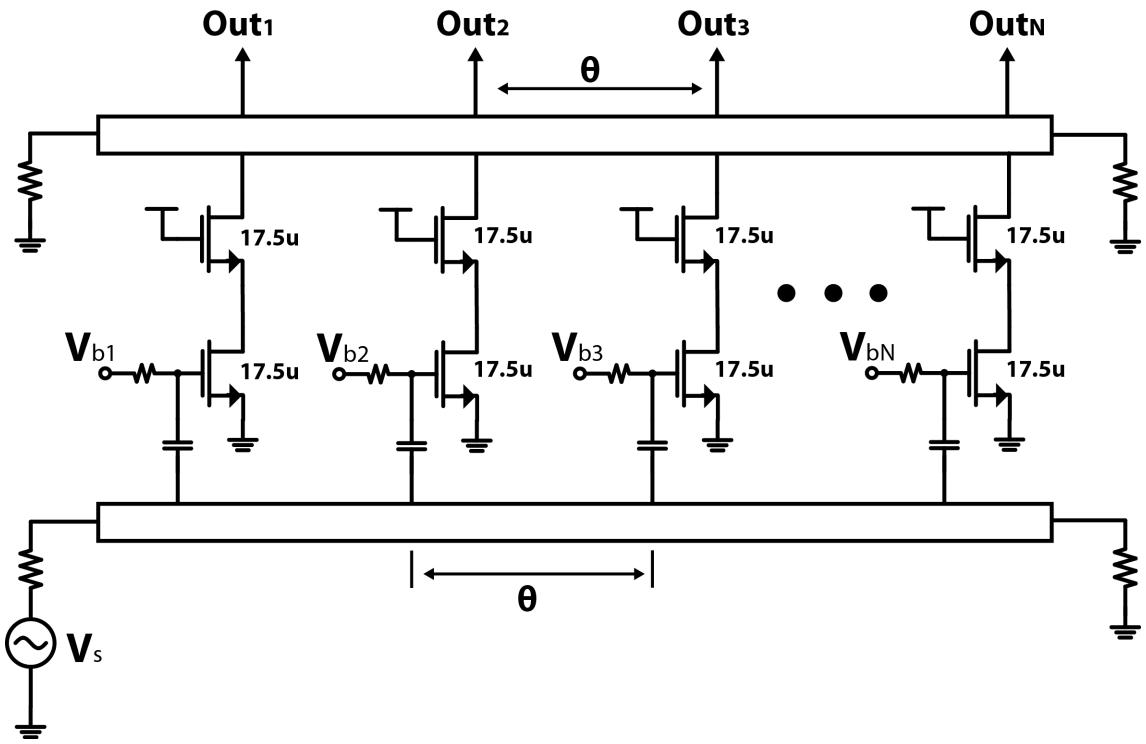

Figure 4.13: Schematic of Distributed Amplifier.

we may assume that the mixer is the dominant linearity-limiting block. Thus, the IIP_3 of the AFE is simply the IIP_3 of the mixer, input referred:

$$IIP_3 = \frac{IIP_{3,HRM}}{G_{DA}} \tag{4.10}$$

Looking at Equations (4.5), (4.9), and (4.10), we see that there may be a maximum allowable G_{DA}. Starting with $G_{DA} = 1$, increasing G_{DA} decreases both NF_{AFE} (Equation 4.9) and IIP_3 (Equation 4.10) which appear on each side of the inequality in Equation 4.5. As G_{DA} is futher increased, NF_{AFE} becomes dominated by the noise performance of the DA, as desired, and becomes less dependent on G_{DA}. At this point, any further increase in G_{DA} only lowers IIP_3 while NF_{AFE} remains almost constant and equal to the noise figure of the DA. This makes it harder to satisfy the inequality in Equation 4.5 and thus, further increasing G_{DA} is undesirable.

Given a targetted DA NF of 6dB, F_{HRM} of 79, $IIP_{3,HRM}$ of +11dBm, and 6b system resolution, we can see from Figure 4.14 that the inequality of Equation 4.5 is met with large margin for G_{DA} values below 18dB. In this figure, the righthand side (RHS) and lefthand side (LHS) of the inequality are plotted as a function of G_{DA} for various targeted system resolutions. According to Equation 4.5, we require that $RHS > LHS$ in order to meet specifications.

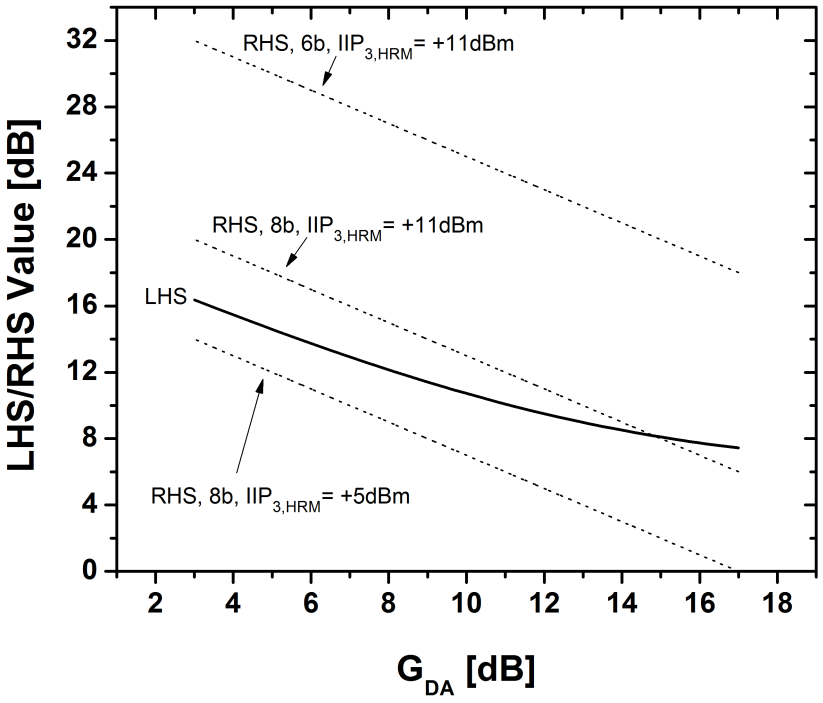

Figure 4.14: Plot of RHS and LHS of Equation 4.5 for various system specifications.

Recall that Equation 4.5 was derived assuming a 6b system. If we were designing an 8b system, the maximum value for G_{DA} becomes 15dB.

The implication of Equation 4.6 can be seen by lowering $IIP_{3,HRM}$ to +5dBm for an 8b system. As shown in Figure 4.14, there aren't any sufficient values of G_{DA} that can satisfy the inequality which means that the linearity/noise performance of the AFE is insuffcient for an 8b system. Thus, maximizing the linearity performance of the mixer is critical when designing higher resolution systems.

In order to allow for the AFE to be adapatable for potential usage in an 8b system, we need to set $G_{DA} < 15$dB. As such, 10dB was the targeted gain for the DA with a NF of 5-6dB. Figures 4.15 and 4.16 show the simulated gain and group delay response of the designed DA. Figure 4.15 plots the response exhibited at each tap. The DC gain is 7.8dB and simulated NF is 4dB. The gain is below the targeted 10dB value due to T-line losses. As seen in the figure, the bandwidths of all taps are not equivalent – earlier taps (taps closest to the source) exhibit lower bandwidths than later taps with the final tap achieving the full 24GHz bandwidth. This behavior is explained in more detail in the following section.

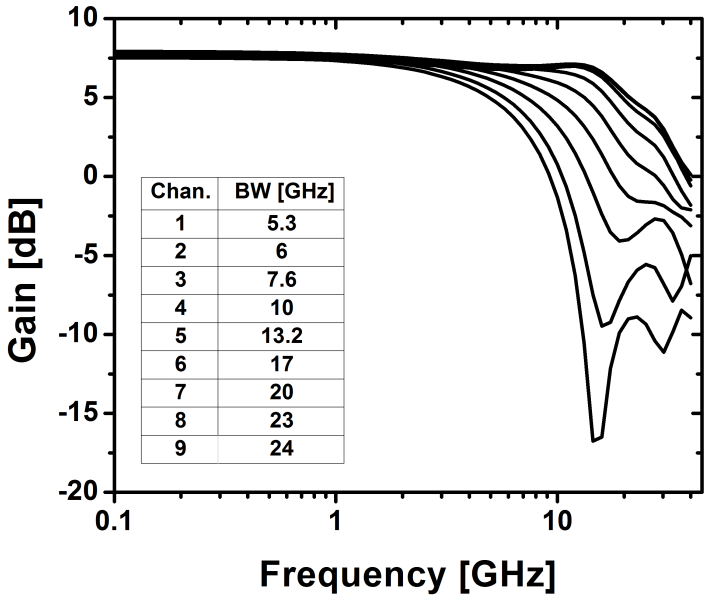

Figure 4.15: Simulated frequency response of the DA.

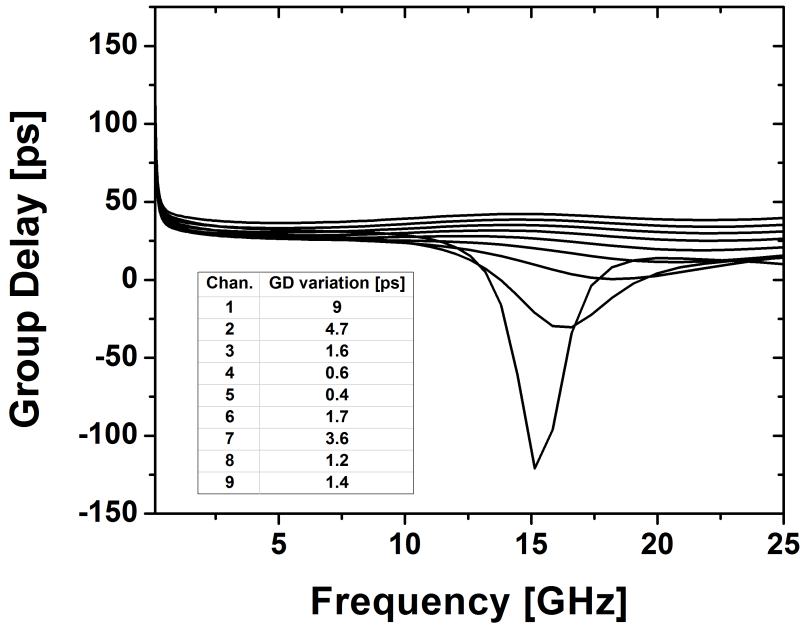

Figure 4.16: Simulated group delay response of the DA.

Tap Bandwidth

Traditional DAs have a single output taken at the final stage. The forward traveling wave on the gate line is coupled to the drain line via the gain stages. A forward and reverse traveling wave is then generated on the drain line. If the propagation velocities of the gate and drain T-lines are equal, then the forward traveling waves on the drain line add coherently, yielding a voltage gain proportional to the number of stages [33]:

$$V_o = Z_d \cdot I_d \tag{4.11a}$$

$$= Z_d \cdot -\frac{g_m}{4} v_s e^{-(N-1)j\theta} \cdot N \tag{4.11b}$$

where V_o is the voltage at the load, Z_d is the load impedance of the drain line which is equal to the characteristic impedance, I_d is the drain current at the load, g_m is the transconductance of each gain stage, v_s is the input voltage amplitude, N is number of stages, and θ is the phase delay between each stage[7]. From Equation 4.11b, we see that the output voltage is proportional to N and g_m, while the exponential term captures the phase delay from input to output.

Equation 4.11b is valid for the last stage of the DA. Unlike the signal at the last stage (load) which is a summation of only forward traveling waves, the signal at each intermediate tap of the DA is a superposition of the forward traveling waves *and* the reverse traveling waves created by the latter stages. Due to the varying path lengths, the reverse traveling wave will experience a different phase delay than the forward traveling wave. This creates a low pass filtering effect – the phase difference is larger at higher frequencies ($\theta \propto \omega$) and results in higher attenuation when summing the forward and reverse waves. This low pass filtering effect is more pronounced for the earlier stages due to the larger number of reverse traveling waves, with successively increasing bandwidth at each tap as you move towards the final output.

To illustrate this effect, we can derive the magnitude response at each tap. The *total* drain current at tap m is a summation of the m forward traveling currents created by the gain stages preceeding (and including) that stage and the $N - m$ reverse traveling waves created by succceeding gain stages:

$$I_{d,m} = I_{d,forward} + I_{d,reverse} \tag{4.12a}$$

$$= \frac{1}{2} \sum_{k=1}^{m} i_{d,k} e^{-(m-k)j\theta} + \frac{1}{2} \sum_{k=m+1}^{N} i_{d,k} e^{-(k-m)j\theta} \tag{4.12b}$$

where $i_{d,k}$ is the drain current contributed by stage k, and is given by:

[7]Note that θ is proportional to the frequency of the signal, $\theta \propto \omega$.

$$i_{d,k} = -g_m v_{gs,k} \tag{4.13a}$$
$$= -g_m \frac{v_s}{2} e^{-(k-1)j\theta} \tag{4.13b}$$

where v_s is the amplitude of the input signal. Using Equation 4.13b, $I_{d,forward}$ becomes:

$$I_{d,forward} = -\frac{g_m}{4} v_s \sum_{k=1}^{m} e^{-(k-1)j\theta} e^{-(m-k)j\theta} \tag{4.14a}$$
$$= -\frac{g_m}{4} v_s \sum_{k=1}^{m} e^{-(m-1)j\theta} \tag{4.14b}$$
$$= -\frac{g_m}{4} v_s e^{-(m-1)j\theta} \cdot m \tag{4.14c}$$
$$\tag{4.14d}$$

and $I_{d,reverse}$ becomes:

$$I_{d,reverse} = -\frac{g_m}{4} v_s \sum_{k=m+1}^{N} e^{-(k-1)j\theta} e^{-(k-m)j\theta} \tag{4.15a}$$
$$= -\frac{g_m}{4} v_s e^{(m+1)j\theta} \sum_{k=m+1}^{N} e^{-(2k)j\theta} \tag{4.15b}$$

With Equations (4.14d) and (4.15b), Equation (4.12b) now becomes:

$$I_{d,m} = -\frac{g_m}{4} v_s \left[m e^{-(m-1)j\theta} + e^{(m+1)j\theta} \sum_{k=m+1}^{N} e^{-(2k)j\theta} \right] \tag{4.16a}$$
$$= -\frac{g_m}{4} v_s e^{j\theta} \left[m e^{-mj\theta} + e^{mj\theta} \sum_{k=m+1}^{N} e^{-(2k)j\theta} \right] \tag{4.16b}$$

The expression in brackets in Equation 4.16b is a complex number that captures the frequency dependent amplitude (and additional phase shift). The frequency dependence is captured by θ. Thus, we are interested in the magnitude of the expression in brackets. We can first simplify the summation term using the identity:

$$\sum_{k=0}^{N} e^{jkx} = \frac{\sin(\frac{1}{2}(N+1)x)}{\sin(\frac{1}{2}x)} e^{jxN/2} \tag{4.17}$$

The summation in Equation 4.16b can now be written as:

$$\sum_{k=m+1}^{N} e^{-(2k)j\theta} = \sum_{k=0}^{N} e^{-(2k)j\theta} - \sum_{k=0}^{m} e^{-(2k)j\theta} \tag{4.18a}$$

$$= \frac{\sin((N+1)\theta)}{\sin(\theta)} e^{-jN\theta} - \frac{\sin((m+1)\theta)}{\sin(\theta)} e^{-jm\theta} \tag{4.18b}$$

$$= S_N e^{-jN\theta} - S_m e^{-jm\theta} \tag{4.18c}$$

where $S_N = \frac{\sin((N+1)\theta)}{\sin(\theta)}$ and $S_m = \frac{\sin((m+1)\theta)}{\sin(\theta)}$. Plugging Equation 4.18c back into Equation 4.16b, we get (after simplification):

$$I_{d,m} = -\frac{g_m}{4} v_s e^{-j(m-1)\theta} \left[m + S_N e^{-j(N-2m)\theta} - S_m e^{jm\theta} \right] \tag{4.19}$$

Equation 4.19 is a general formula for the drain current at any given tap in the DA. We can verify that it agrees with the traditional formula by plugging in $m = N$:

$$I_{d,N} = -\frac{g_m}{4} v_s e^{-j(N-1)\theta} \left[N + S_N e^{-j(N-2N)\theta} - S_N e^{jN\theta} \right] \tag{4.20a}$$

$$= -\frac{g_m}{4} v_s e^{-j(N-1)\theta} \left[N + S_N e^{jN\theta} - S_N e^{jN\theta} \right] \tag{4.20b}$$

$$= -\frac{g_m}{4} v_s e^{-j(N-1)\theta} \cdot N \tag{4.20c}$$

which agrees with the formula for drain current at the load that was given in Equation 4.11b.

We can now easily identify the real and imaginary parts of the complex number contained in brackets in Equation 4.19:

$$\Re[I_{d,m}] = m + S_N \cos\left((N-2m)\theta\right) - S_m \cos(m\theta) \tag{4.21a}$$

$$\Im[I_{d,m}] = -\left(S_N \sin\left((N-2m)\theta\right) + S_m \sin(m\theta)\right) \tag{4.21b}$$

The magnitude response as a function of θ (and hence, frequency) can be written as:

$$|I_{d,m}| = \frac{g_m}{4} v_s \cdot \left[\left(m + S_N \cos\left((N-2m)\theta\right) - S_m \cos(m\theta)\right)^2 + \left(S_N \sin\left((N-2m)\theta\right) + S_m \sin(m\theta)\right)^2 \right]^{\frac{1}{2}} \tag{4.22}$$

Equation 4.22 captures the frequency dependence of the drain current magnitude as affected by the summation of forward and reverse traveling waves along a *lossless* transmission line. We can plot the magnitude response as a function of frequency by substituting in the long expression for θ [37][38]:

$$\theta = \beta \cdot l \tag{4.23a}$$
$$= \omega\sqrt{LC} \cdot l \tag{4.23b}$$

where β is the imaginary part of the T-line propogation constant, l is the distance (measured along the T-line) between each gain stage, and L and C are the inductance and capacitance per unit distance, respectively. Figures 4.17 and 4.18 plot the normalized frequency response according to Equation 4.22 and the group delay response as derived from Equations 4.21a and 4.21b. The T-line parameters (L, C, l) are chosen to be on the same order as the designed T-line. However, the plots are not meant to model the true responses of the designed DA, but only capture the general behavior.

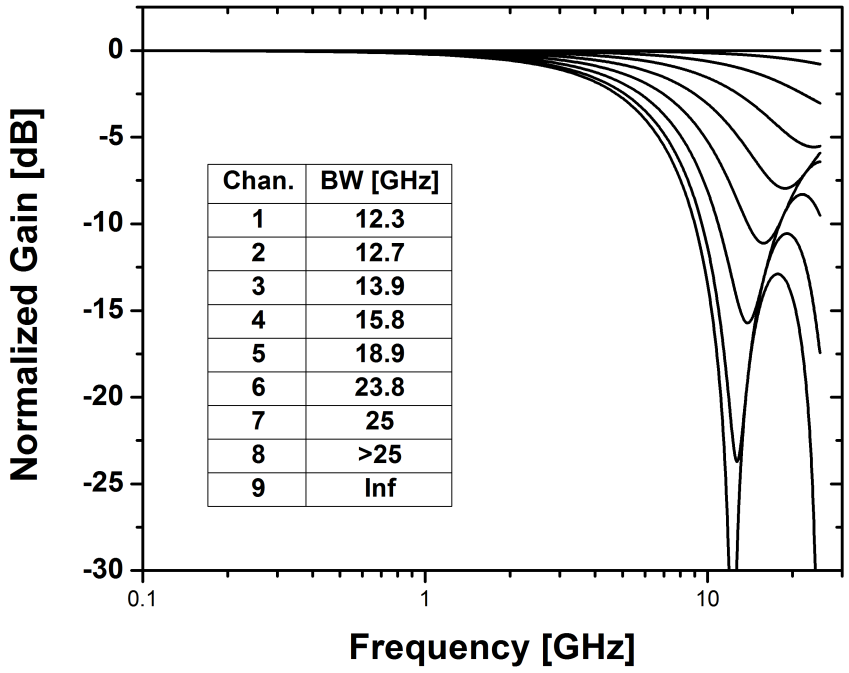

Figure 4.17: DA frequency response of according to Equation 4.22.

Comparing Figure 4.17 to Figure 4.15, we see that it captures the bandwidth limitation experienced by earlier stages in the DA. The final stage's response experiences no attentuation (has infinite bandwidth) since there are no reverse traveling waves to destructively interefere with the forward traveling waves. In the real design, this bandwidth is finite and set by the losses of the T-line as well as the cutoff frequency of the synthesized T-line [37][38].

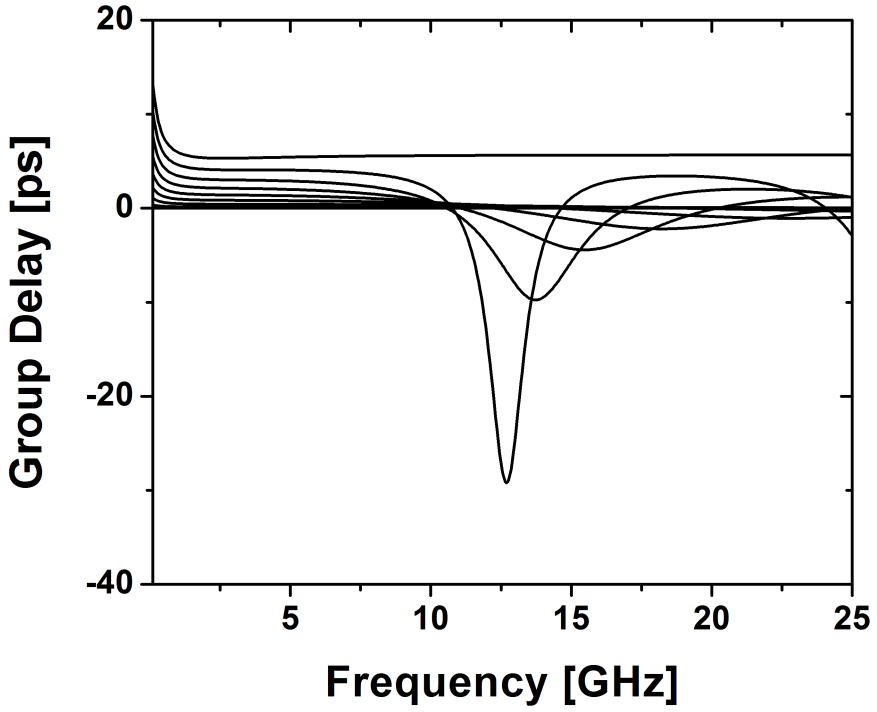

Figure 4.18: DA group delay response of according to Equations 4.21a and 4.21b.

The low-pass frequency response of each tap is a convenient effect as it attenuates out-of-band signals. The only constraint is to make sure that the inherent bandwidth of each DA tap is greater than or equal to the highest frequency that channel is responsible for processing.

Current-Mode LPF

The high-frequency lowpass filter is implemented using a current-mode topology as presented in [39]. The schematic of the lowpass filter is shown in Figure 4.19. As discussed in [39], the input impedance looking up into the drains of M_1 and M_2 resemble that of an inductor in shunt with a resistor (Figure 4.20). The output current being fed to the load resistor is the current through the inductor of the equivalent circuit. When combined with the input capacitor, C_1, the transfer function formed from i_{in} to i_{out} is a second-order lowpass filter with two complex-conjugate poles. The lowpass behavior can be understood by observing Figure 4.20. At DC, the inductor presents ideally zero impedance to the source and passes all of the input current to the load. As the frequency is increased, the inductor impedance increases, creating a less desirable path for the signal current to flow while the capacitor, C_1, shunts increasingly more of the signal current to ground.

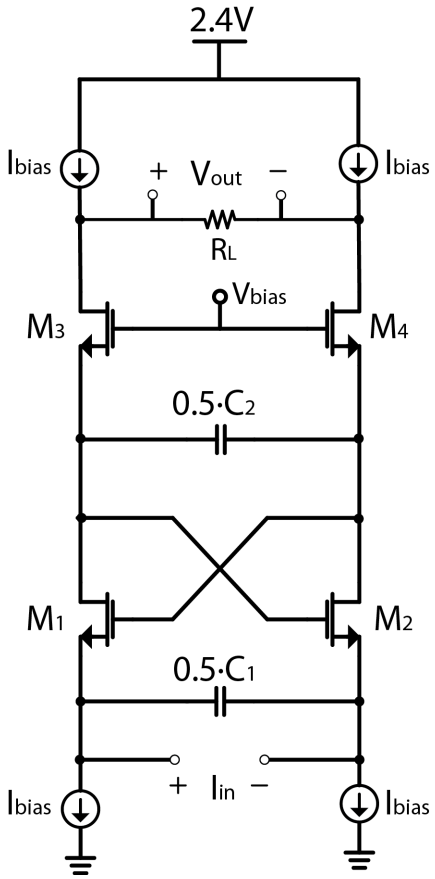

Figure 4.19: Schematic of current-mode lowpass filter.

The equivalent inductance and shunt resistance placed in parallel with C_1 are given by:

$$L_{eq} = \frac{C_2}{g_m^2} \tag{4.24}$$

$$R_{eq} = \frac{1}{g_m} \tag{4.25}$$

where g_m is the transconductance of transistors M_1 and M_2. The transfer function of the filter is then given by:

$$\frac{i_{out}}{i_{in}} = \frac{g_m^2/C_1C_2}{s^2 + s \cdot \frac{g_m}{C_1} + \frac{g_m^2}{C_1C_2}} \tag{4.26}$$

Equation 4.26 is a second-order lowpass response with a corner frequency, ω_o, and Q factor of:

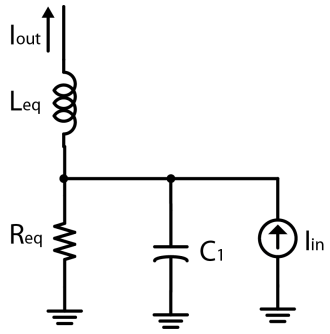

Figure 4.20: Equivalent circuit for current-mode lowpass filter.

$$\omega_o = \frac{g_m}{\sqrt{C_1 C_2}} \tag{4.27}$$

$$Q = \sqrt{\frac{C_1}{C_2}} \tag{4.28}$$

Equations (4.27) and (4.28), are two equations with three unknowns. With the addition of one more constraint, there is enough information to perform the design of the filter. The last constraint is derived from the desire to maximize the conversion gain of the entire AFE. The maximum input impedance of the filter is $\frac{1}{g_m}$ and occurs at ω_o [39]. This input impedance needs to be much smaller than the output impedance of the HRM in order to draw all of the signal current into the filter. The caveat is that we would like to do this within a reasonable power budget since the AFE will have 17 instances of the lowpass filter. Given this, $\frac{1}{g_m}$ was set to 100Ω. This value yielded acceptable tradeoff between the degradation in the AFE conversion gain and the power consumption of the LPF. The filter quality factor was set to $Q \approx 2$ since low Q poles yield improved group delay response (e.g. Bessel filters) and the minimum allowable Q is $\frac{1}{2}$ in order to guranatee complex conjugate poles.

The corner frequency of the LPF was designed to be larger than the channel processing bandwidth for improved in-band group delay and noise performance. The simulated magnitude, group delay, and noise responses are shown in Figures 4.21, 4.22 and 4.23, respectively. The filter has a cutoff frequency of 4.5GHz and exhbits less than 2ps of group delay variation within the 1.5GHz bandwidth. The noise response exhbits the behavior predicted by [39]. The excessive bandwidth of the filter was needed to push the noise bump exhibited in Figure 4.23 far enough out so as to not affect the in-band noise performance; due to the lack of excessive front-end gain, the noise performance of latter blocks has a non-negligible impact on the noise figure of the AFE.

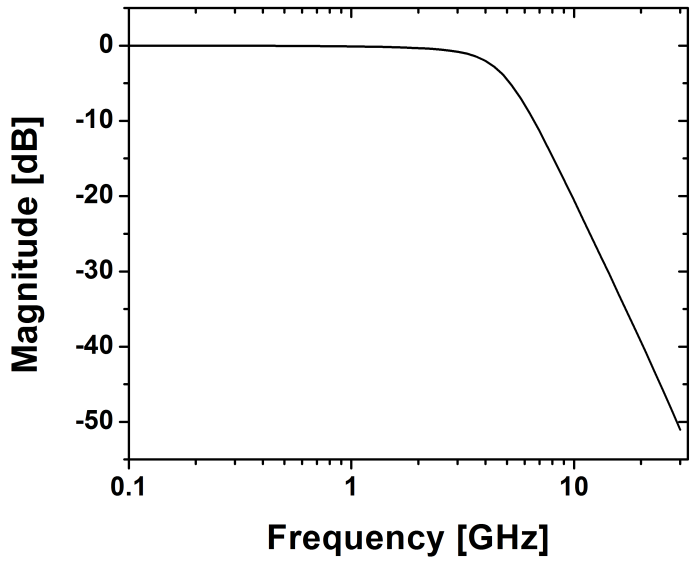

Figure 4.21: Simualted magnitude response of lowpass filter.

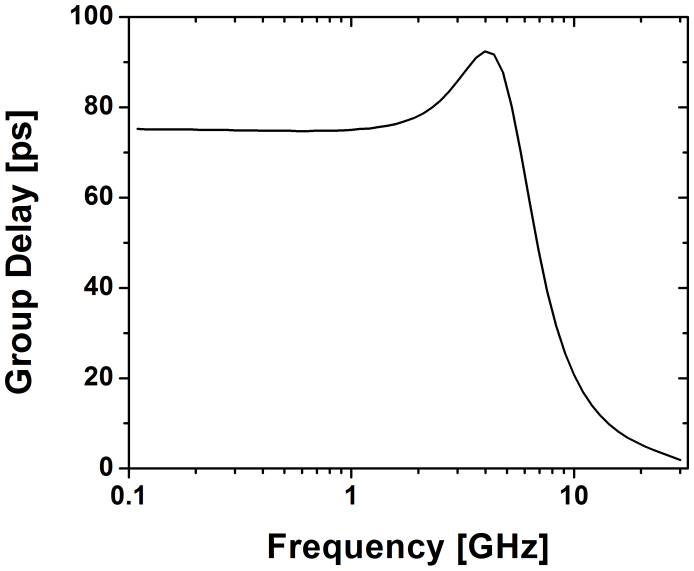

Figure 4.22: Simulated group delay response of lowpass filter.

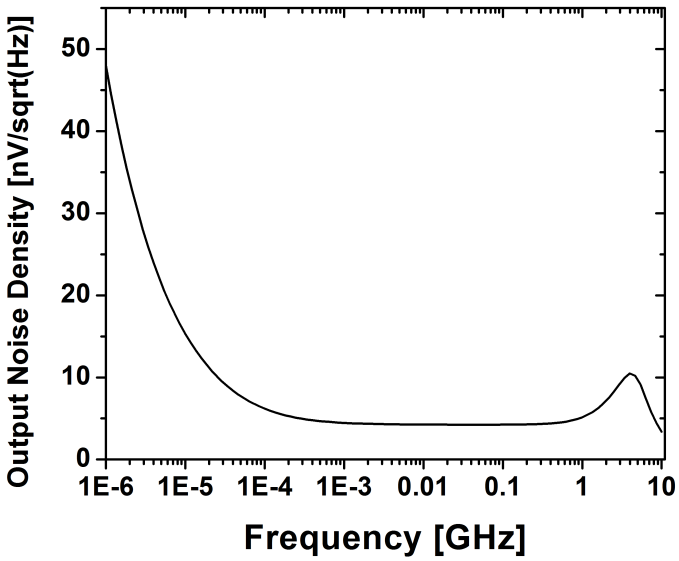

Figure 4.23: Noise response lowpass filter.

For this design, the lowpass filter was included in order to perform wideband current-to-voltage conversion for driving the 50Ω output buffers while also providing some level of attenuation of out-of-band signals. Proper filtering in the analog domain is handled by off-chip filters that precede the sub-ADCs, making the DA, harmonic rejction mixers, and LO generation/distribution, the focus of the design.

LO Generation

In order to provide the eight LO frequencies needed for downconversion and harmonic rejection, two chains of frequency dividers are used (Figure 4.4). The first chain is driven by a 48 GHz input, and generates the 24, 12, 6, and 3GHz LOs. The second chain is driven by a 36GHz input, and generates 18 and 9GHz. SSB mixing is performed between the 3 and 18 GHz LOs to generate 15 and 21GHz. CML latch-based dividers are utilized to perform frequency division at high frequencies [40]. Figure 4.24 shows a schematic of the CML dividers.

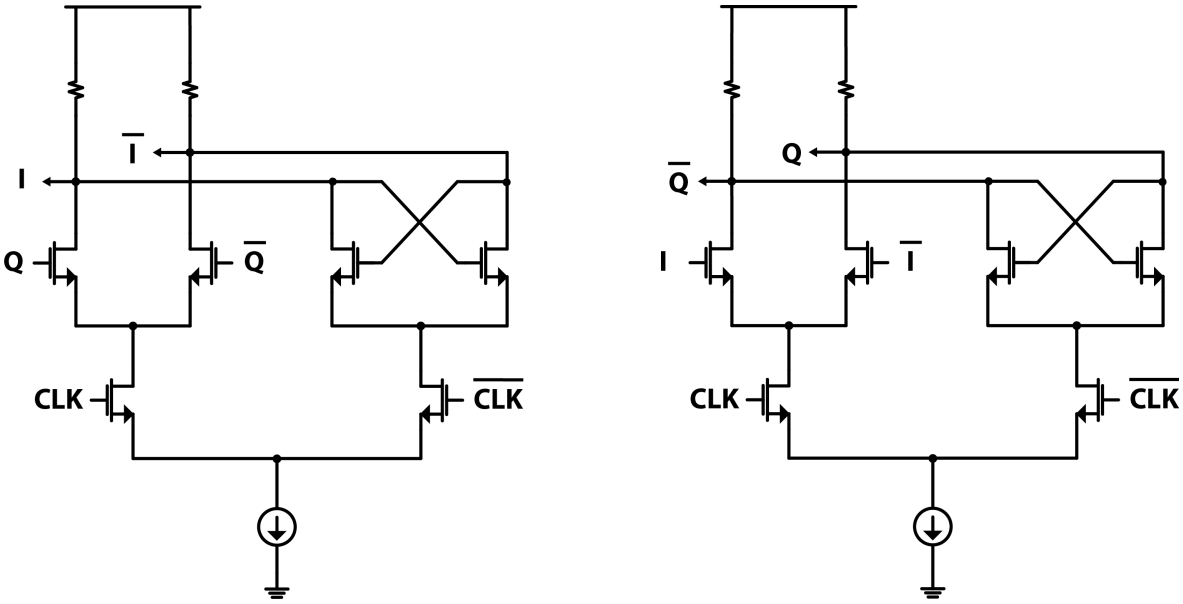

Figure 4.24: Schematic of high-speed CML latch-based divider.

The same divider topology is used for all frequencies. The device sizes and bias currents are progressively scaled down as the input frequency to each successive divider is cut in half.

SSB Mixing

In the case of SSB mixing (generation of 15 and 21GHz), care must be taken to ensure that the image rejection is sufficient for the targeted resolution. Furthermore, any nonlinear products introduced during the SSB mixing process must also be below the required noise floor. Figure 4.25 shows a block diagram of the architecture of the SSB mixer.

A passive switching quad is used to perform the mixing between the two frequencies with minimal added noise. The RF port of the switching quad is driven by the 3GHz LO via a g_m stage, while the LO port is driven by the 18GHz LO. The 3GHz LO is fed to the SSB mixer prior to the full rail-to-rail amplification that occurs in LO distribution chain. This maintains a "small-signal" on the RF path so as to reduce non-linearites created by the g_m stage. The choice of port allocation (i.e. 3GHz to RF port, 18GHz to LO port) reduces the amount of nonlinearites introduced at the output of the mixer caused by harmonics of the LO. A current buffer follows the passive mixer in order to perform summation in the current domain while preventing unwanted loading/crosstalk between the two paths. The current buffer is implemented using a psuedo-differential topology as shown in Figure 4.26. The tuned load on the current buffer is used to filter out the image frequency and any other undesired tones.

In order enable tuning of the image rejection to within specification, varactors are placed on the tuned 18GHz LO buffers that feed the SSB mixer. By making small adjustments

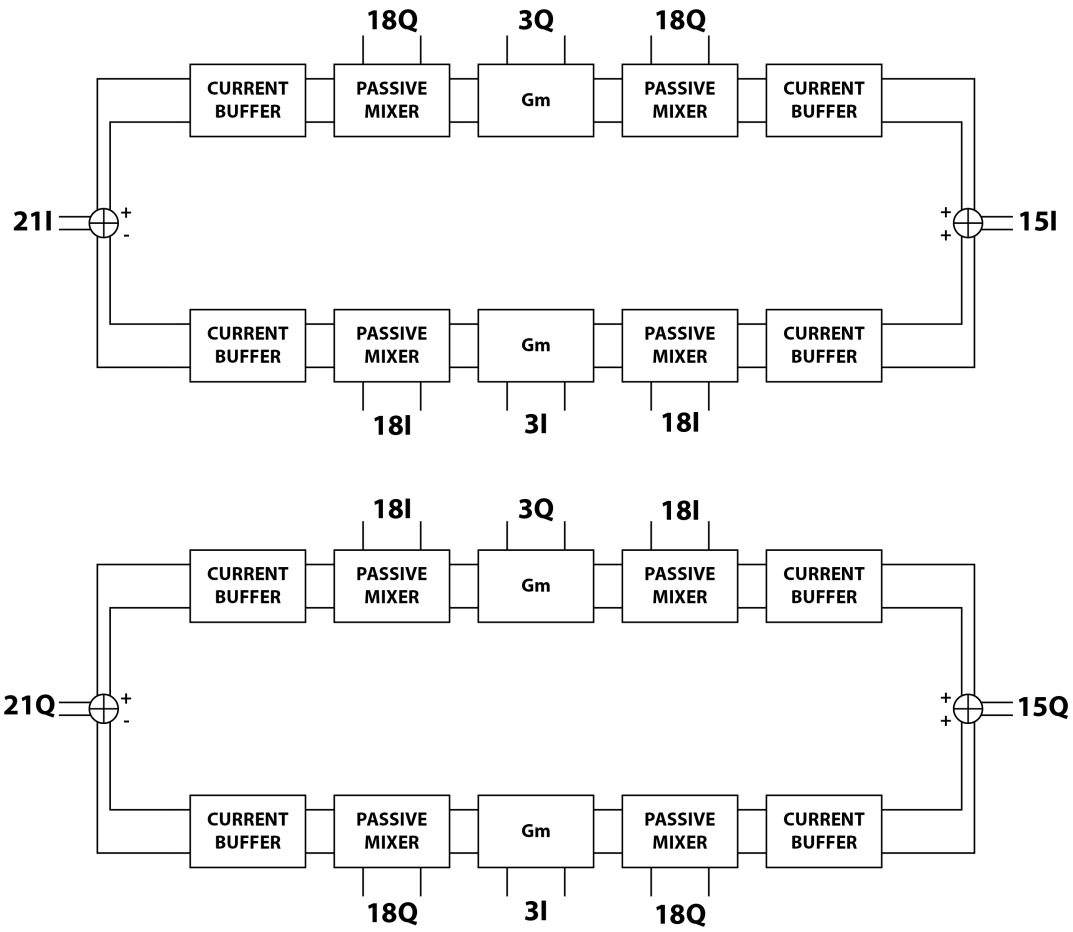

Figure 4.25: Architecture used for SSB mixing.

to the resonant frequency of the tank, varactors can introduce sufficient phase shift on the 18GHz LO path while maintaing sufficient LO amplitude to drive the passive switching quad. The varactor is tuned using a 4b voltage DAC. In addition, the gain of the g_m stage is also tunable via a programmable tail current.

LO Buffering/Distribution

The passive switching quad of the AFE downconversion mixers require large amplitude LO signals for sufficient conversion gain and noise performance. In order to avoid significant degradation in mixer gain and noise figure, a $0.8 - 1\mathrm{V_{pp}}$ signal needs to delivered to the LO port of the mixer. Given the $200 - 300\mathrm{mV_{pp}}$ swing at the output of the CML dividers, LO buffers are utilized to amplify the LO signals to the appropriate level. For the high LO frequencies (12GHz and above), a two-stage buffering approach is used. The first stage, referred to as the *PreAmp*, is an intermediary buffer between the CML divider and the

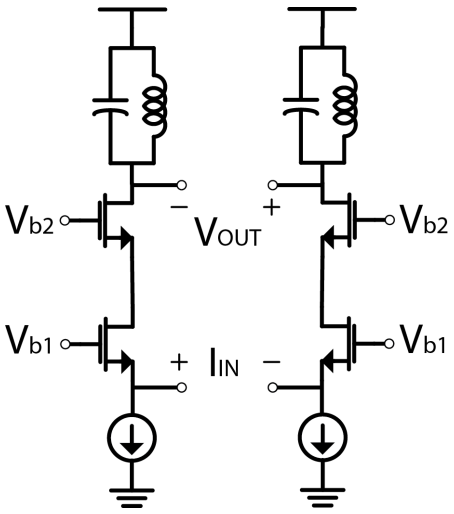

Figure 4.26: Current buffer for SSB mixer.

main LO driver, and presents a minimal amount of capacitive loading to the high-frequency dividers while providing a gain of ~ 2. The second stage is the *LO driver* which amplifies the signal to the full $1V_{pp}$ nominal swing and drives the 100fF capacitive load presented by the passive mixer quad in addition to the parasitic capcitance associated with the routing. For the lower LO frequencies (9GHz and below), a static CMOS inverter chain is used as shown in Figure 4.27. The output of the CML dividers are AC-coupled to the first stage of self-biased inverters. The resistive feedback places the input and output bias points at mid-rail and the pMOS and nMOS are sized such that the threshold voltage, V_M, of the inverter is also equal to the mid-rail voltage. This yields maximum sensitivity and minimal duty-cycle distortion when performing the level conversion. Differential signaling is reinforced by using weak cross-coupled inverters across the two single-ended paths.

A majority of the LO distribution routes have lengths greater than 500μm. As a result, the routing capacitance is very significant. To make matters worse, high-frequency channels are subject to transmission-line effects which transform the load capacitance to a larger value dependent on the length of the line [37]. A differential transmission line ($Z_{o,d} = 100\Omega$) is used for the routing of the LOs in order to eliminate the parasitic capacitance to ground associated with the routing, leaving only the transformed load to be driven. The 9, 15 and 18GHz LOs are impacted the most since they are high frequency signals with the longest trace lengths [8]. For the 15 and 18GHz LOs, it is critical to properly model and simulate these load transformations in order to accurately design the resonance frequency of the LO drivers. Extensive EM simulations were performed using Integrand EMX®.

Increased load capacitance due to load transformations can lead to significant increase in power consumption of the LO drivers connected to these loads. In order to reduce the power

[8]The 9, 15 and 18GHz LOs travel across the full length of the DA in order to reach the phase interpolators used for harmonic rejection, and thus have the longest trace lengths ($>$ 1mm).

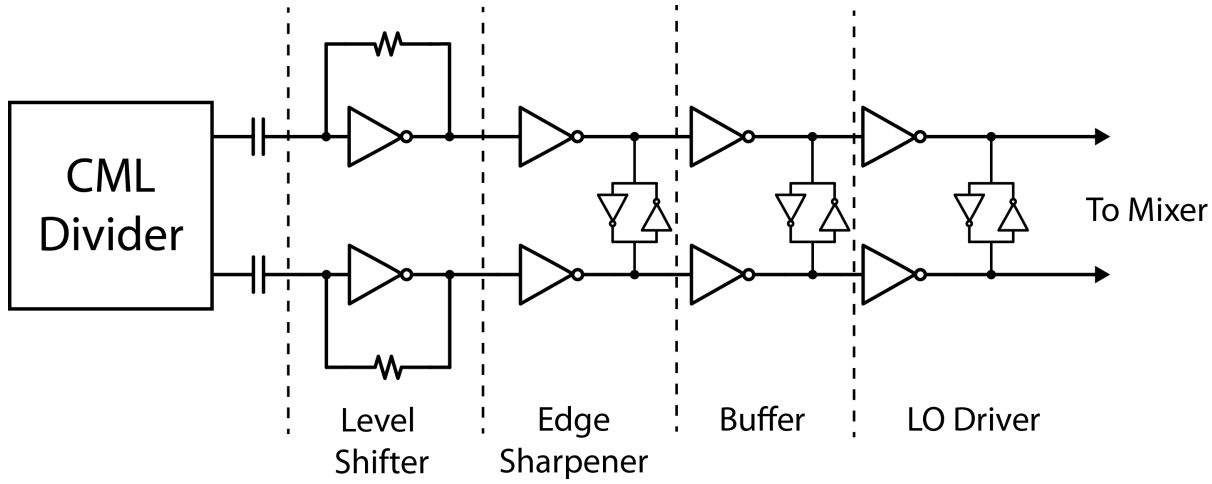

Figure 4.27: CMOS inverter chain used for 3, 6, and 9 GHz LO buffering.

consumption of the 9, 15, and 18GHz LO drivers, intermediary buffers are used to segment the lengthy routing which load the drivers. As shown in Figure 4.28, the traces were divided into two segments: LO driver → channel mixer and channel mixer → phase interpolator. By segmenting the traces, the load transformations are mitigated (reduced trace lengths) and the power consumed by the LO drivers is greatly reduced.

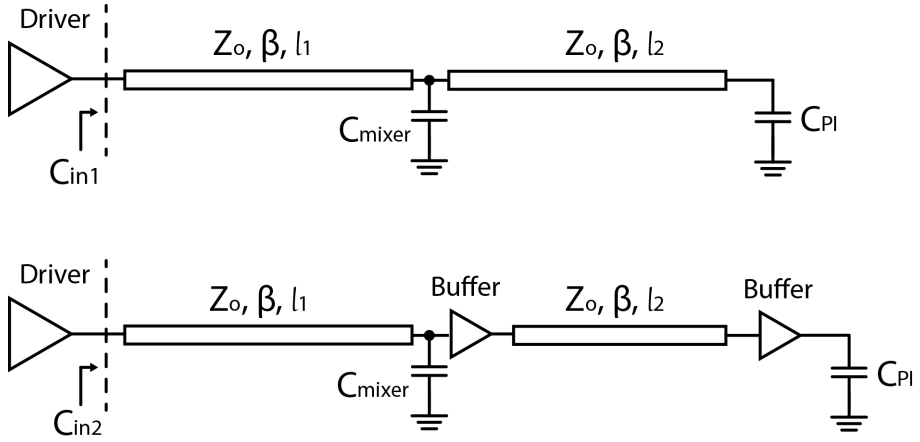

Figure 4.28: Placement of intermediary buffers to reduce load transformations.

LO Routing Technique

With eight LOs being routed within the same vicinity of one another, coupling/crosstalk between the LO routing is a concern. Crosstalk on the LO routing leads to undesired down-conversion of out-of-band signals due to the presence of LO spurs located at other channel's

LO frequencies. In order to reduce the amount of crosstalk, a planar twisted-pair routing technique is used to reject magnetic coupling between the the lines. The twisted-pair is a very common and useful RF technique used in communication systems (e.g. Ethernet) that effectively reduces magnetic coupling by alternating the polarity of the induced *emf* on the conducting path [38]. Figure 4.29 illustrates how the planar twisted-pair was implemented using the two top metal layers.

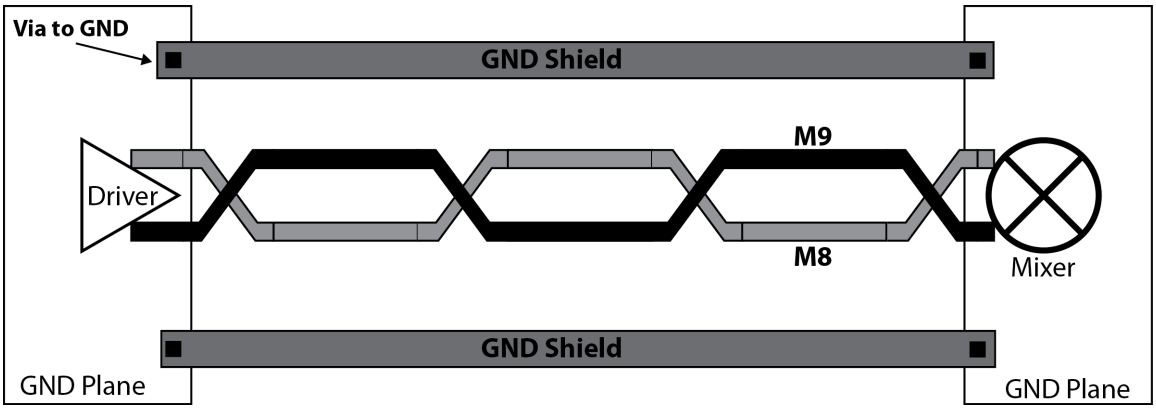

Figure 4.29: Planar twisted pair technique used for LO routing.

The planar twisted-pair rejects *magnetic* coupling onto the LO distribution lines. However, the routing is still susceptible to *electric* coupling between nearby traces. The electric coupling can also be viewed as capacitance that is accumulated bewtween two closely-spaced parallel metal wires and which provides a path for high-frequency signals to couple from one wire to the other. This capacitance is formed when the two metal wires are at different potentials, causing the electric field lines of one wire to terminate on the other. If the induced charge caused by electrical coupling does not have a low-resistive path to ground then it will induce an undersired change in potential on the affected line; this is the nature of the electrical crosstalk. Shielding is the common solution to reducing electrical crosstalk [41]. As shown in Figure 4.29, planar ground shields are placed on both sides of the twisted pair. The ground shields are implemented on the top metal layer. Since the shield lines are nominally at zero potential, most of electric field lines eminating from the LO traces will terminate on the shield lines which provide a path to ground for any induced charge.

In order to retain a high common-mode rejection along the differential LO routes, the ground shields are only shorted to ground at the source and load. This is because the ground shield is the return path for any common-mode signals, and thus, it is desirable to have a a high impedance on the common-mode return path in order to impedede any common-mode current flow. Of course, this is in contrast with the desired low-resistance needed on the ground line for mitigating electric coupling, so a reasonable tradeoff must be made.

Chapter 5

Measurements

5.1 FI-ADC Analog Front-End

As mentioned in the Chapter 4, the first phase of this project focused on the design of the analog front-end. The AFE is responsible for the distrbution and channelization of the wideband input signal while sustaining sufficient linearity and noise (SNDR) performance for the specified resolution. A design of the FI-ADC AFE was taped out in a TSMC 65nm GP+ process. The die photo of the microchip is shown in Figure 5.1. As seen in Figure 4.4, the design includes a 9-stage DA which feeds the nine processing channels. Each processing channel contains a g_m stage, I/Q passive mixer, and 2^{nd} order current-mode lowpass filter. In order to facilitate measurements, 50Ω differential output buffers are used to interface each channel with external equipment. The LO generation and distribution is comprised of frequency dividers, SSB mixers, LO buffers and line drivers. Harmonic rejection is performed for Channel 2 (3^{rd} and 5^{th}) and Channel 3 (3^{rd}). In order to calibrate HRR performance, 9-bit current DACs are utilized in the phase interpolators and gm-stages.

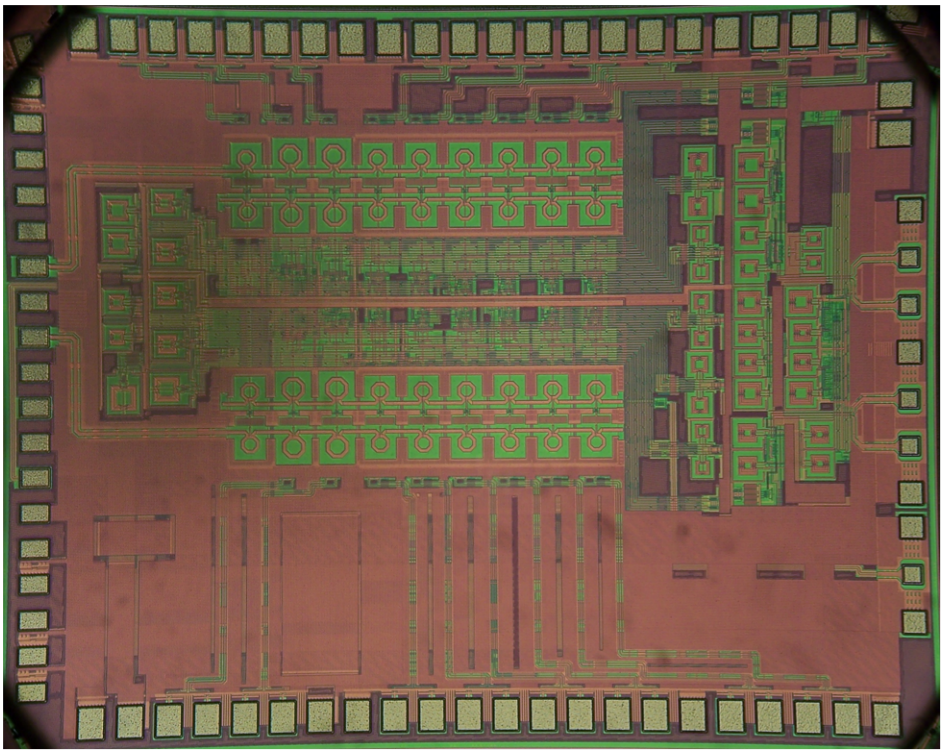

Figure 5.1: Die photo of FI-ADC AFE.

5.2 Measurement Setup

The focus of the initial measurements of the FI-ADC was to assess the linearity and noise performance. The measurement setup used for testing is shown in Figure 5.2. A chip-on-board (COB) assembly was used for wirebonding the baseband outputs of each channel to a host board with SMA connectors. The host board is then connected to a custom high-speed ADC board used for digitization of the analog outputs. The high-speed ADC board passes the digital data of each channel to a Virtex-7 FPGA via FMC connectors. The FPGA is programmed to offload the data from the ADCs and pass it to the PC for backend processing and signal reconstruction.

The initial implementation of the custom high-speed board contains two 3GS/s, 8-bit ADCs (PN: ADC08B3000)[1]. As such, the I and Q paths of a single channel could be processed simultaneously in the backend. The high-speed ADCs have integrated 4kB buffers which allow for the offloading of the data at a user-defined rate of 100MHz. On-board lowpass filters ($\omega_c = 900$MHz) are included to provide sufficient out-of-band attenuation of aliased signals. As discussed in Chapter 4, the fully integrated system would require an 6GS/s ADC in conjunction with a 6^{th} order Inverse-Chebyshev filter. Since 6GS/s ADCs

[1]Generously donated by Texas Instruments.

are not commercially available, 3GS/s ADCs had to be used along with appropriately chosen filters. This means that only 900MHz of the full 1.5GHz bandwidth of each I/Q channel could be utilized in the current measurement setup.

Die probing is utilized for connecting the high-frequency inputs. The left-side probe contains the RF input, 1.2V (LO GEN and Analog Core) and 2.4V power rails, input reference current, and scan chain signals. The right-side probe brings in the differential 36GHz and 48GHz signals. Two Keysight 50GHz signal generators are used for generating the differential signals.

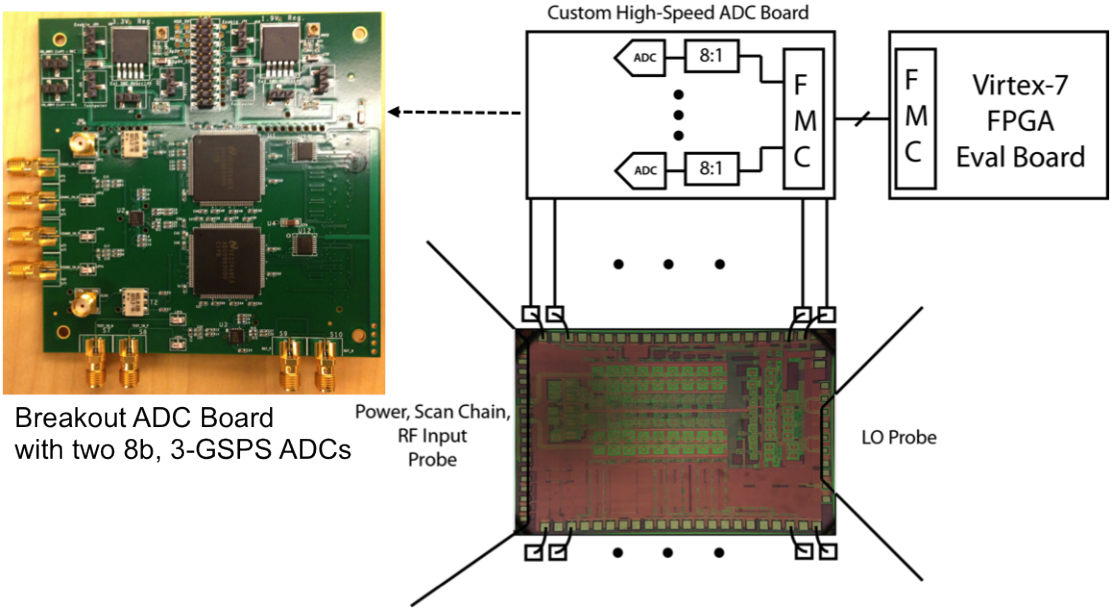

Figure 5.2: Measurement setup.

Due to limited equipment availability, the 36GHz and 48GHz signals could not be generated simultaneously. The initial method for generating the phase-coherent 36GHz and 48GHz signals is illustrated in Figure 5.3. This method utilizes multiplication of a single 3GHz source up to the desired frequencies. Filters are used to filter out any undesired harmonics generated during the multiplication process. The signal is fed single-endedly to the chip while the other terminal is tied to the input bias voltage. The one issue with this method was the amount of signal loss incurred after the last stage of quaruplers. The measured output power of the quadruplers was +18dBm while the amount of signal loss due to the filters, bias-tees, cables, and the probe itself was 19.5dB resulting in a delivered power of -1.5dBm, which was insufficient to achieve locking with a single-ended drive. As such, the two 50GHz signal generator approach had to be adopted.

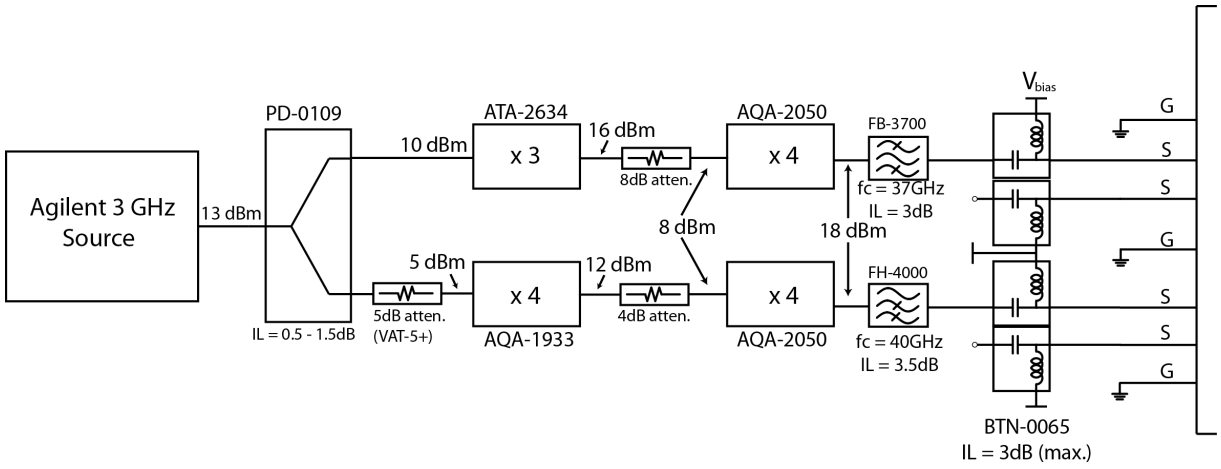

Figure 5.3: Generation of phase-coherent 36GHz and 48GHz input.

5.3 Measurement Results

The first set of measurements focused on measuring the linearity of the channels. Figure 5.4 and Figure 5.5 show the IIP_2, IIP_3, and P_{1dB} measurements. The measured IIP_2 and IIP_3 are +5dBm and +30dBm, respectively. Once the intercept points were identified, subsequent testing was conducted at 16dB backoff from these points in order to yield intermodulation products that are in compliance with SNDR requirements for the targeted resolution, in this case, 32dB. This corresponds to a Pin of -11dBm.

An FFT was peformed on the output samples from the ADC to allow for extraction of SNDR performance. The SNDR measurements for the 18GHz and 24GHz channel are shown in Figure 5.6. From the plots we see that the SNDR at the output is sufficient for a 5b system, with a peak SNDR of 34.6dB. The output power levels were slightly lower than what we expected from simulations. The cause of this reduced output power is partially due to high on-chip temperatures– entire chip consumes 900-950mW. Simple cooling of the chip using a condensed aircan shows up to 5dB improvement in the output power, which placed the measured ENOB much closer to the 6b target resolution. Each channel maintains greater than 5b resolution over a bandwidth of 1.8GHz which is the bandwidth enforced by the lowpass filters on the ADC board. Beyond this bandwidth, degradation in the ENOB is caused by attenuation of the output power due to the lowpass filter. Within the 1.8GHz bandwidth, there is still a bandpass shaping present. This is likely caused by the upconversion of the baseband parasitic capacitors in the passive mixer.

The 18GHz and 24GHz channels were the only passband channels that could be measured since the latter dividers did not lock without externally cooling the chip. Locking of the 18GHz-to-9GHz divider was observed when the chip was cooled (using a simple air spray can) which is a clear indicator that on-chip temperature is an issue.

A sample set of the corresponding FFTs from which the SNDR was extracted is shown in

Figure 5.7. The FFT plots exhibit spectral widening around the main signal. This widening is not due to spectral spreading since the input frequencies were chosen to ensure coherent sampling with respect to the sampling frequency. Instead, this non-ideality was present in the analog output spectrum as shown in Figure 5.8. It is suspected that these are noise sidebands caused by coupling between the input signal and the supply network resulting in undesired modulation. These noise sidebands are proportional to the input signal power level and are also present in the baseband channel where no mixing is performed. The sidebands are contained within a 20MHz bandwidth around the fundamental tone.

As demonstrated by Figure 5.9, these noise sidebands limit the measured ENOB. The figure shows a plot of the theoretical ENOB (ENOB_ideal), raw ENOB (ENOB), and a modified ENOB (ENOB_mod). The ENOB_ideal response exhibits the 1 bit per 6dB slope as predicted by theory. The raw ENOB is extracted from the unaltered FFT plots. As seen in the figure, the raw ENOB only follows the theoretical trend for low input powers, after which it begins to degrade. This degradation is due to the increasing power levels of the noise sidebands which become higher than the thermal and quantization noise floor at higher input power levels. In order to get a measure of the ENOB due to the true noise floor, and not the noise sidebands, the noise sidebands around the fundamental tone and its harmonics[2] were set equal to the neighboring noise floor level. This modified ENOB measurement, ENOB_mod, exhibits behavior much closer to what is expected when input power is increased. However, there is still degradation in ENOB_mod at higher input power levels. The deviation above $P_{in} = -15$dBm is partly due to the onset of compression. Note that ENOB_mod is the reported ENOB for this design and what is reported in Figure 5.6.

Due to the locking issue with latter dividers, direct harmonic rejection measurements could not be performed. In order to assess the potential harmonic rejection performance, indirect measurements of the phase interpolator were taken. Figure 5.10 shows the measured phase response of the phase interpolator. Because of the non-ideal nature of the measurement, the output signal was noisy and reliable phase measurements could only be performed for 6 bits of resolution on the interpolator; the full resolution of the interpolator was 9 bits. With 6 bits of resolution, the worst-case phase step was measured to be 5.4° (Figure 5.11). Given this phase resolution, we performed a circuit-level simulation of the mixer to assess the attainable harmonic rejection ratio subject to the worst-case resolution. From the simulation, the attainable HRR is 37dB.

The measured SNDR performance and indirectly measured HRR performance illustrate that the AFE is suitable to be used in a 5.5b ADC system. What's left to be demonstrated is the reconstruction of a wideband signal (e.g. pulse-modulated carrier). Unfortunately, reconstruction of a multi-channel signal was infeasible with the current measurement setup for two reasons. First, the custom high-speed ADC board only allowed simultaneous sampling of two outputs at a time, thus the best case would be sampling of I/Q of a single channel. Second, the frequency divider locking issue caused by on-chip heating prevents a majority

[2]These are harmonics of the baseband signal (post-downconversion) that are created by the final 50Ω output buffer

of the channels from being measured consistently and reliably. Measurement setups that provide better heat removal from the die are being investigated.

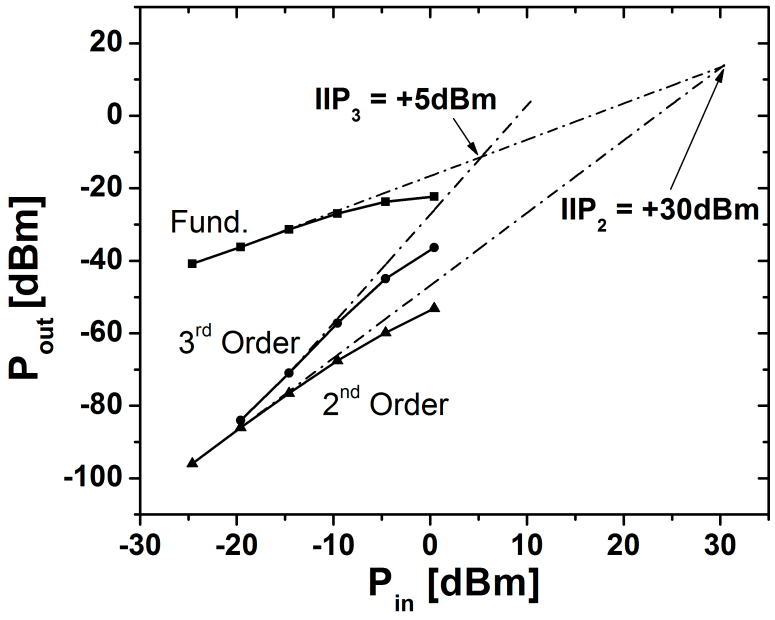

Figure 5.4: $IIP_{2/3}$ measurements.

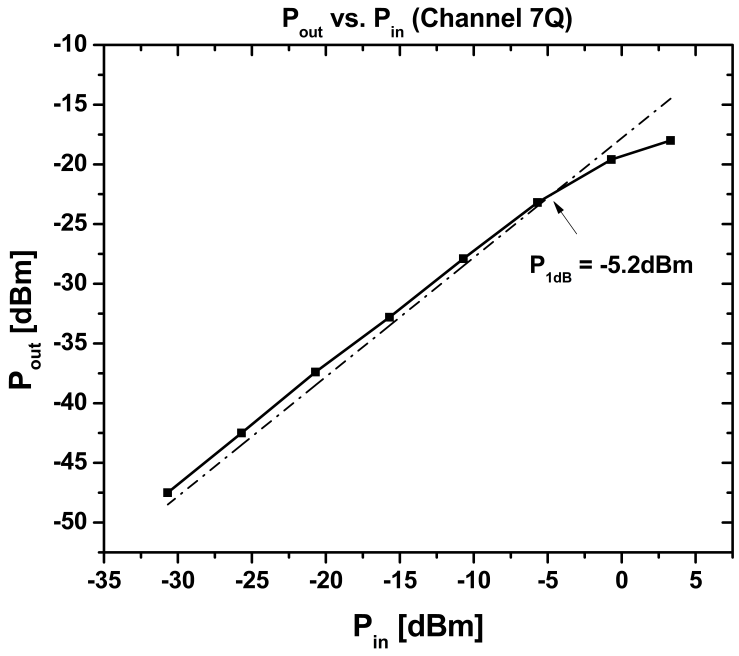

Figure 5.5: P_{1dB} measurements.

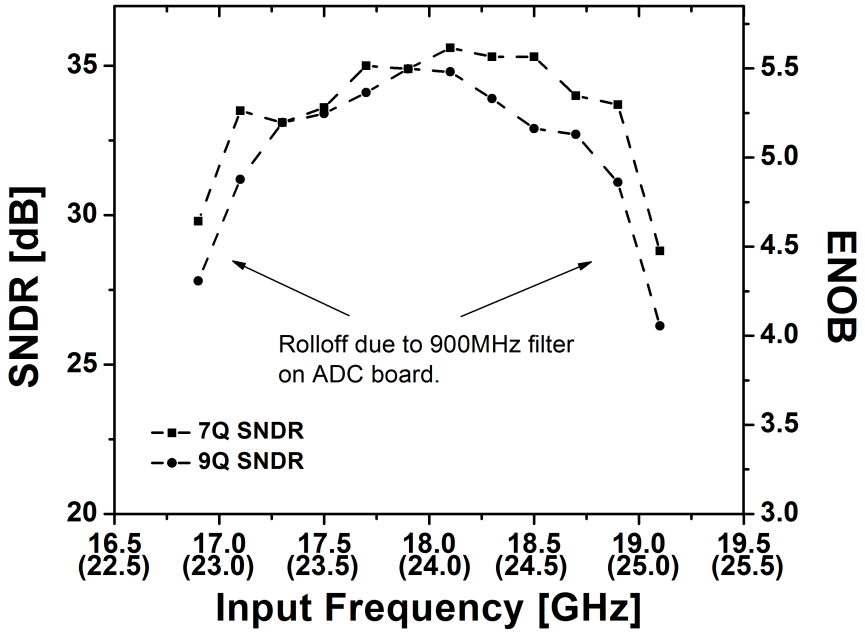

Figure 5.6: SNDR measurements.

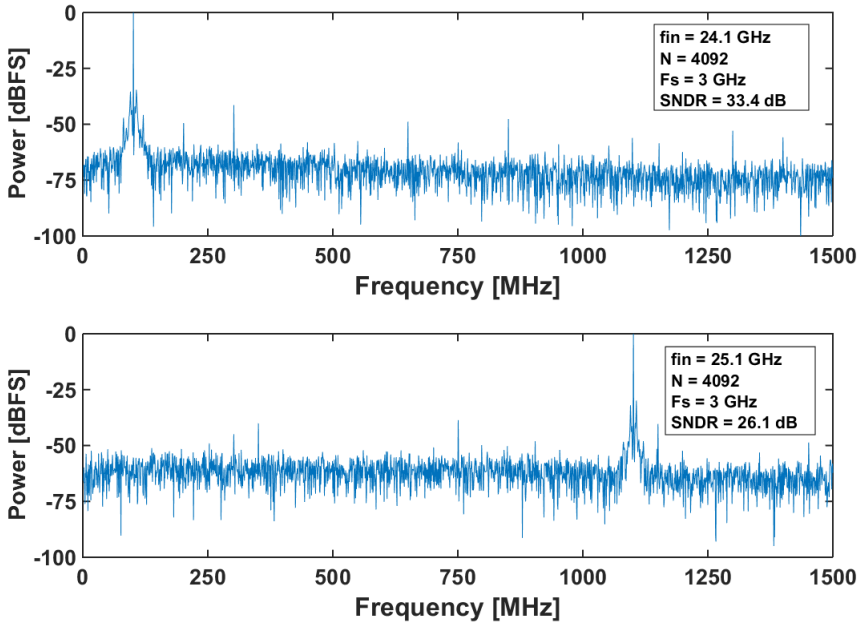

Figure 5.7: FFT of baseband signal for 24.1GHz and 25.1GHz input signals.

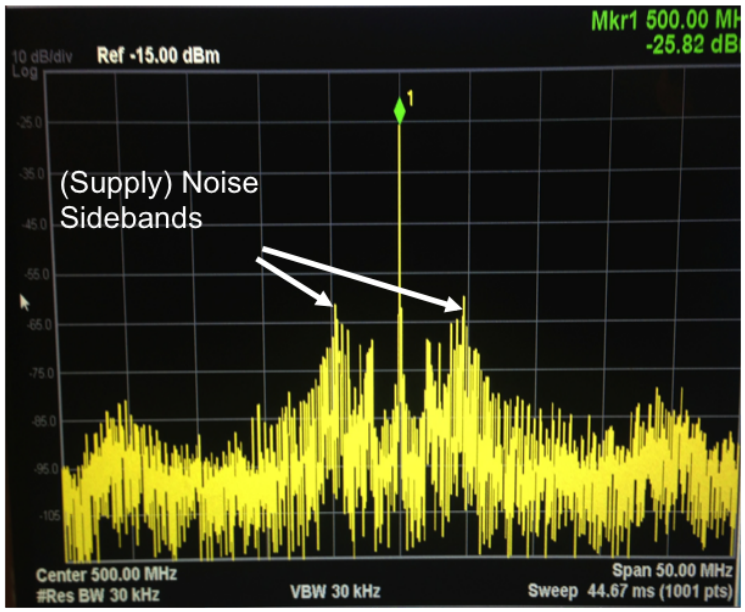

Figure 5.8: Spectrum of analog output.

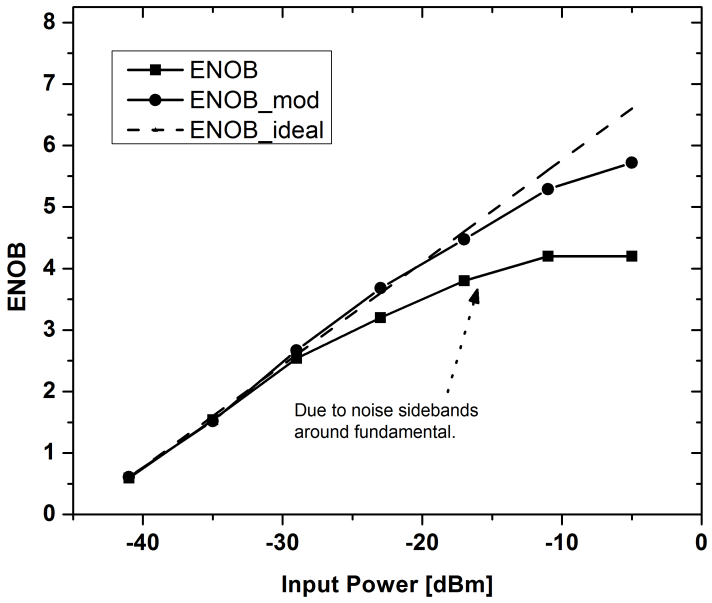

Figure 5.9: Impact of noise sideband on measured ENOB.

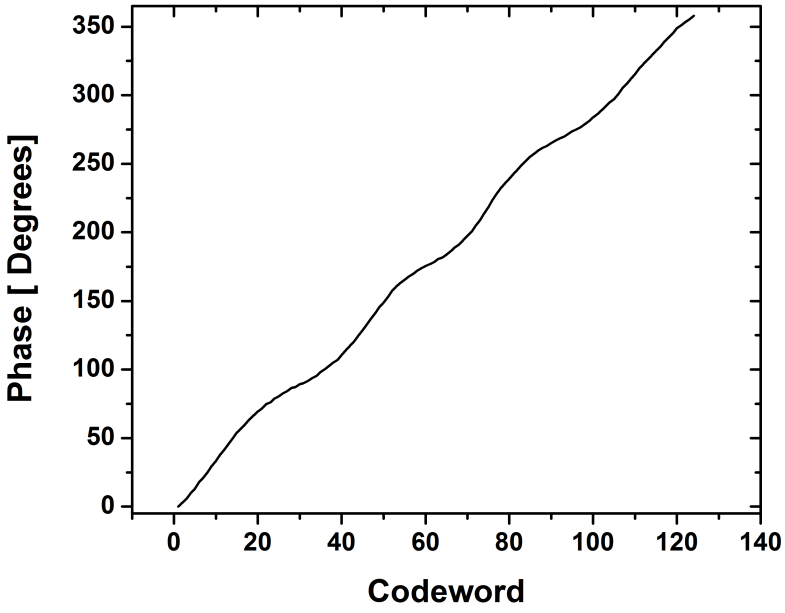

Figure 5.10: Measured phase response of phase interpolator.

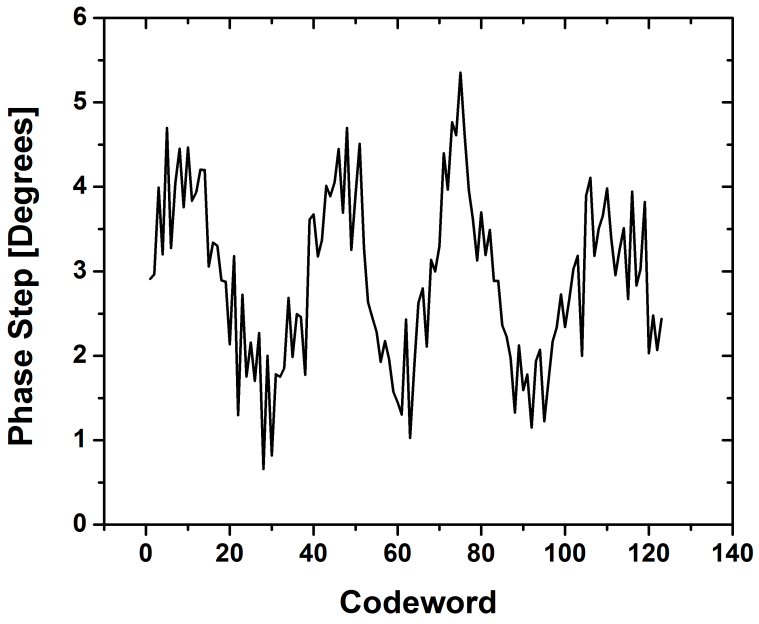

Figure 5.11: Measured phase steps of phase interpolator.

Chapter 6

Conclusion

Summary of Thesis

The frequency-interleaved ADC (FI-ADC) has great potential to provide improved resolution over the time-interleaved ADC (TI-ADC) at very high sample rates. The resolution of conventional high-speed ADC architectures is ultimately limited by the aperture jitter in the sampler. The majority of these ADCs can only offer resolutions up to 4b at sample rates greater than 25GS/s[1]. Proposed as an alternative to the TI-ADC, the FI-ADC architecture has been shown to have a decreased sensitivity to jitter due to the decreased bandwidth presented to the sampler. However, up until now, there hasn't been a detailed analysis which provides insight into the impact that LO phase noise has on the FI-ADC architectures that utilize downconversion mixers. The focus of this work was to answer this question, first and foremost, while also providing insight into designing energy-efficient wideband FI-ADCs.

In this thesis, a comparison between the two ADC architectures has been presented. System-level simulations have shown that the FI-ADC can improve ADC resolution by 1-2b at the Nyquist frequency. It has also been shown that LO phase noise plays a critical role in determining how much of a net benefit is gained by switching to a FI architecture. As such, proper care must be taken to ensure that the LO phase noise is properly designed so as to not mask the inherent SNR gain offered by channelizing the signal bandwidth.

Design techniques for optimizing front-end design while lowering power consumption have also been presented. The FI-ADC is a complex system and a high-level view must be taken at the onset of design in order to properly manage the tradeoffs between the analog front-end, sub-ADCs, and the digital backend. An overview of the design tradeoffs between various blocks in the design has been provided to help guide future designs.

Lastly, the design of an AFE for usage in a 5-6b 50GS/s FI-ADC system was presented. By utilizing a front-end distributed amplifier and high-frequency HRMs, the AFE is capable of distributing and channelizing a 25GHz signal with sufficient SNR and linearity for a 5-6b system. The AFE achieves an IIP3 of +5dBm and peak output SNDR of 34dB for a 25GHz input signal, which equates to 5.3b ENOB.

[1]This is for published ADCs and the resolutions are measured at the Nyqusit frequency.

Future Work and Directions

While this work has focused on some key fundamental questions regarding the FI-ADC, there are a lot more questions to be answered. For one, a measurement setup for comparing the direct sampling architecture to the mix-then-sample setup would be informative and useful for backing the presented theory. This measurement setup would parallel the simulation setup used in Chapter 3 and feature only the front-end sampling network of TI-ADCs and the AFE of the FI-ADC. To simplify the comparison, the two architectures could be compared only at the Nyqusit frequency since it yields the worst-case performance for both architectures. This will simplify the FI-ADC design since only a single channel needs to be designed.

Another key question to answer is whether or not the FI-ADC is a contender to the TI-ADC at lower sampling frequencies. Although the TI-ADC can provide 8-10b ENOB for sampling frequencies of 1-5GS/s, there are still emerging applications that need to push the performance up to 12 bits – a very difficult task in the presence of the jitter barrier (Figure 1.1) – and the FI-ADC may pave the way to those additional two bits. Now, at very high sampling frequencies (such as the targeted design for this work), a drawback of the FI-ADC is the intimidating overhead of generating and distributing the multiple LO frequencies. However, at lower frequencies, this may become less of an issue since distribution becomes (relatively) easier and the generation of multiple LOs can be done with advanced digital techniques as was done in [20]. Building on the techniques of [20], further work can be done to push the attaianble ENOB of the FI architecture at sampling frequencies of 1-5GS/s.

Lastly, another concern for the FI-ADC, which wasn't discussed in this thesis, is the power consumption of the digital backend. Methods for handling and processing multiple high-speed data channels in an energy-efficient manner are essential to keeping the FI-ADC figure-of-merit comparable to that of the TI-ADC. These methods will likely require innovations at both the circuit and system level.

General Thoughts and Comments

Throughout this project, I have had countless discussions with professors and industry members about whether the FI-ADC is truly better than the TI-ADC. The response is generally mixed. On one hand, there are potential believers in the FI-ADC who recognize the severity of the jitter problem and who believe that there is some potential to alleviate this issue with the FI architecture. On the other hand, there are skeptics who immediately, and justifiably, raise the question of phase noise introduced during the mixing process, claiming that the jitter problem has simply been transferred to the mixer and therefore, there is no performance benefit in moving to a FI architecture. My answer to this is always: it depends. As I have hopefully demonstrated in this thesis, the phase noise of the LO plays a critical role in determining the performance gain, if any, afforded by moving to an FI architecture. As seen from the simulation results, there are scenarios where the FI-ADC can outperform the TI-ADC depending on how the LO is designed and distributed. My hope is that these

results may help curb the skepticism and instead prompt a more detailed look into whether or not the FI architecture is suitable for a particular application.

I would be remiss if I did not mention my own doubts and struggles with the FI architecture and its performance in comparison to the TI architecture. It is a very subtle problem. However, after countless hours mulling over the equations, the theory, and the simulations, my belief is that the FI-ADC is an answer to the jitter problem faced by current state-of-the-art high-speed ADCs, provided it is designed properly. I guess, in a sense, my answer is still: it depends. The main question is whether the sacrificed ENOB at lower frequencies is justified by the gains at the very high-frequencies (Figure 3.13). This is dependent on the application and a decision that needs to be made by the engineering architect. If filtering is utilized in the LO distribution (via tuned LO drivers), then the answer may be trivial since the FI architecture may outperform the TI architecture across the board as it did in Figure 3.16. However, depending on the LO generation/distribution employed, the usage of tuned buffers may not be an option, especially considering the large area consumed by inductors, and the performance comparison may be closer to that of Figure 3.18. Still, as bandwidth demands increase and new applications emerge, the FI architecture may be utilized more frequently for the accurate digitization of very high frequencies,. This is already happening in the oscilloscope industry where an increasing number of high-end oscilloscopes use the FI architecture (e.g. Teledyne LeCroy). Perhaps this trend will spread to the many other high-speed applications.

Bibliography

[1] C. Wu and B. Nikolic, "A 0.4 GHz - 4 GHz direct RF-to-digital $\Sigma\Delta$ multi-mode receiver," in *2013 Proceedings of the ESSCIRC (ESSCIRC)*, Sept 2013, pp. 275–278.

[2] A. Arbabian, "Time-Domain Ultra-Wideband Synthetic Imager in Silicon," Ph.D. dissertation, University of California, Berkeley, 2011.

[3] A. Arbabian, S. Callender, S. Kang, B. Afshar, J.-C. Chien, and A. M. Niknejad, "A 90 GHz Hybrid Switching Pulsed-Transmitter for Medical Imaging," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 12, pp. 2667 –2681, Dec. 2010.

[4] A. Arbabian, S. Callender, S. Kang, M. Rangwala, and A. M. Niknejad, "A 94 GHz mm-Wave-to-Baseband Pulsed-Radar Transceiver with Applications in Imaging and Gesture Recognition," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 4, pp. 1055–1071, Apr. 2013.

[5] S. Velazquez, "Hybrid Filter Banks for Analog/Digital Conversion," Ph.D. dissertation, Massachusetts Institute of Technology, 1996.

[6] J. Black, W.C. and D. Hodges, "Time interleaved converter arrays," *IEEE Journal of Solid-State Circuits*, vol. 15, no. 6, pp. 1022–1029, Dec 1980.

[7] R. Walden, "Analog-to-Digital Converter Survey and Analysis," *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 4, pp. 539–550, Apr. 1999.

[8] B. Murmann, *ADC Performance Survey 1997-2015*. [Online]. Available: http://web.stanford.edu/~murmann/adcsurvey.html

[9] E. Alon, *EE240 Class Notes Spring 2009*.

[10] L. Kull, T. Toifl, M. Schmatz, P. Francese, C. Menolfi, M. Braendli, M. Kossel, T. Morf, T. Andersen, and Y. Leblebici, "A 90GS/s 8b 667mW 64x Interleaved SAR ADC in 32nm digital SOI CMOS," in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, Feb 2014, pp. 378–379.

[11] N. Kurosawa, H. Kobayashi, K. Maruyama, H. Sugawara, and K. Kobayashi, "Explicit analysis of channel mismatch effects in time-interleaved ADC systems," *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 48, no. 3, pp. 261–271, Mar 2001.

[12] D. Stepanovic and B. Nikolic, "A 2.8 GS/s 44.6 mW Time-Interleaved ADC Achieving 50.9 dB SNDR and 3 dB Effective Resolution Bandwidth of 1.5 GHz in 65 nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 4, pp. 971–982, April 2013.

[13] N. Le Dortz, J.-P. Blanc, T. Simon, S. Verhaeren, E. Rouat, P. Urard, S. Le Tual, D. Goguet, C. Lelandais-Perrault, and P. Benabes, "A 1.62GS/s time-interleaved SAR ADC with digital background mismatch calibration achieving interleaving spurs below 70dBFS," in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, Feb 2014, pp. 386–388.

[14] X. Yi, C. Boon, H. Liu, J. Lin, J. Ong, and W. Lim, "A 57.9-to-68.3GHz 24.6mW frequency synthesizer with in-phase injection-coupled QVCO in 65nm CMOS," in *2013 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, Feb 2013, pp. 354–355.

[15] V. Szortyka, Q. Shi, K. Raczkowski, B. Parvais, M. Kuijk, and P. Wambacq, "A 42mW 230fs-jitter sub-sampling 60GHz PLL in 40nm CMOS," in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, Feb 2014, pp. 366–367.

[16] K. Lee and W. Namgoong, "A 0.25 um CMOS 3b 12.5 GS/s frequency channelized receiver for serial-links," in *2005 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, Feb 2005, pp. 336–337.

[17] S. Jalali Mazlouman, S. Sheikhaei, and S. Mirabbasi, "A prototype implementation of a two-channel frequency-translating hybrid ADC," in *50th Midwest Symposium on Circuits and Systems (MWSCAS)*, Aug 2007, pp. 144–147.

[18] Jalali Mazlouman, S. and Sheikhaei, S. and Mirabbasi, S., "4- and 6-GS/s 4-bit frequency-translating hybrid ADCs in 90-nm CMOS," *Analog Integrated Circuits and Signal Processing*, vol. 68, no. 3, pp. 329–340, 2011. [Online]. Available: http://dx.doi.org/10.1007/s10470-011-9604-y

[19] S. Hoyos, S. Pentakota, Z. Yu, E. Ghany, X. Chen, R. Saad, S. Palermo, and J. Silva-Martinez, "Clock-Jitter-Tolerant Wideband Receivers: An Optimized Multichannel Filter-Bank Approach," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 58, no. 2, pp. 253–263, Feb 2011.

[20] T. Forbes and R. Gharpurey, "A 2 GS/s Frequency-Folded ADC-Based Broadband Sampling Receiver," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 9, pp. 1971–1983, Sept 2014.

[21] G. Ding, C. Dehollain, M. Declercq, and K. Azadet, "Frequency-Interleaving Technique for High-Speed A/D Conversion," in *Proceedings of the 2003 International Symposium on Circuits and Systems (ISCAS '03)*, May 2003, pp. I–857 – I–860.

[22] S. Jalali Mazlouman, S. Sheikhaei, and S. Mirabbasi, "Digital Compensation Techniques for Frequency-Translating Hybrid Analog-to-Digital Converters," *IEEE Transactions on Instrumentation and Measurement*, vol. 60, no. 3, pp. 758–767, March 2011.

[23] P. J. Pupalaikis and D. C. Graef, "High Bandwidth Real-Time Oscilloscope," U.S. Patent 7 058 548, Jun 6, 2006.

[24] S. Hoyos, B. Sadler, and G. Arce, "Ultra-Wideband Multicarrier Communication Receiver Based on Analog to Digital Conversion in the Frequency Domain," in *2005 IEEE Wireless Communications and Networking Conference*, Mar 2005, pp. 782–787.

[25] J. Weldon, R. Narayanaswami, J. Rudell, L. Lin, M. Otsuka, S. Dedieu, L. Tee, K.-C. Tsai, C.-W. Lee, and P. Gray, "A 1.75-GHz highly integrated narrow-band CMOS transmitter with harmonic-rejection mixers," *IEEE Journal of Solid-State Circuits*, vol. 36, no. 12, pp. 2003–2015, Dec 2001.

[26] A. Rafi, A. Piovaccari, P. Vancorenland, and T. Tuttle, "A harmonic rejection mixer robust to RF device mismatches," in *2011 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, Feb 2011, pp. 66–68.

[27] H. Zhang, T. Gao, S. Tan, and O. Shana'A, "A harmonic-rejection mixer with improved design algorithm for broadband TV tuners," in *2012 IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, June 2012, pp. 163–166.

[28] J. Weldon, "High Performance CMOS Transmitters for Wireless Communication," Ph.D. dissertation, University of California, Berkeley, 2005.

[29] C. Azeredo-Leme, "Clock Jitter Effects on Sampling: A Tutorial," *IEEE Circuits and Systems Magazine*, vol. 11, no. 3, pp. 26–37, 2011.

[30] B. Lathi, *Signal Processing and Linear Systems*. Oxford University Press, 2000.

[31] W. Kester, *Converting Oscillator Phase Noise to Time Jitter*, Tutorial MT-008, Analog Devices, 2009. [Online]. Available: http://www.analog.com/media/en/training-seminars/tutorials/MT-008.pdf

[32] W. Egan, *Frequency Synthesis by Phase Lock*. Wiley-Interscience, 1999.

[33] J. Beyer, S. Prasad, R. Becker, J. Nordman, and G. Hohenwarter, "MESFET Distributed Amplifier Design Guidelines," *IEEE Transactions on Microwave Theory and Techniques*, vol. 32, no. 3, pp. 268–275, Mar. 1984.

[34] S. Jalali Mazlouman and S. Mirabbasi, "A Frequency-Translating Hybrid Architecture for Wide-Band Analog-to-Digital Converters," *IEEE Transactions on Circuits and Systems-II*, vol. 54, no. 7, pp. 576–580, 2007.

[35] C. Paul, "Bandwidth of Digital Waveforms," *IEEE Practical Papers, Articles, and Application Notes*, pp. 58–64, 2009.

[36] L. Ye, J. Chen, L. Kong, E. Alon, and A. M. Niknejad, "Design Considerations for a Direct Digitally Modulated WLAN Transmitter with Integrated Phase Path and Dynamic Impedance Modulation," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 12, pp. 3160–3177, Dec. 2013.

[37] D. Pozar, *Microwave Engineering*. Wiley, 2004.

[38] A. Niknejad, *Electromagnetics for High-Speed Analog and Digital Communication Circuits*. Cambridge, 2007.

[39] A. Pirola, A. Liscidini, and R. Castello, "Current-Mode, WCDMA Channel Filter With In-Band Noise Shaping," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 9, pp. 1770–1780, Sep. 2010.

[40] C. Kromer, G. von Buren, G. Sialm, T. Morf, F. Ellinger, and H. Jackel, "A 40-GHz static frequency divider with quadrature outputs in 80-nm CMOS," *IEEE Microwave and Wireless Components Letters*, vol. 16, no. 10, pp. 564–566, Oct 2006.

[41] H. Ott, *Noise Reduction Techniques in Electronic Systems*. John Wiley & Sons Inc., 1976.