Millimeter-Wave Digital Beamforming using Low-Resolution ADCs

Deniz Dosluoglu

A thesis

submitted in partial fulfillment of the

requirements for the degree of

Master of Science

University of Washington

2022

Committee:

Jacques C. Rudell

Visvesh Sathe

Program Authorized to Offer Degree:

Electrical and Computer Engineering

University of Washington

Abstract

Millimeter-Wave Digital Beamforming using Low-Resolution ADCs

Deniz Dosluoglu

Chair of the Supervisory Committee:
Jacques C. Rudell
Electrical and Computer Engineering

There is an on-going pursuit to design wireless systems that achieve ever higher data rates for multiple users. Phased-array beamforming receivers (RXs) have been demonstrated to support multi-gigahertz bandwidths across the millimeter-wave spectrum. The digital beamforming (DB) architecture is attractive in that it can support the most beams per antenna element compared to the hybrid and analog architectures. However a DB RX requires an identical set of analog-to-digital converters (ADCs), downconverters, filters and amplifiers per element, which, relative to other beamforming architectures, can significantly increase the system area, power consumption and complexity, and in turn limit the system scalability if conventional circuit topologies are used.

This thesis studies the utilization of a dithered and oversampled open-loop 1-bit ADC in a DB RX to reduce the front-end power consumption and complexity while maintaining a high signal bandwidth and preserving the system array gain. A behavioral model of a 1-bit ADC is developed and the impact of dithering 1-bit ADCs in a DB RX is analyzed. The circuit design of an oversampled 1-bit ADC in TSMC 28-nm CMOS is described, along with standalone and system-level measurements of the ADC implemented in a 4-element V-band DB RX. The beamforming system achieves a maximum QPSK symbol bandwidth of 400 MS/s with a per element power and area consumption of 96 mW and 0.18 mm^2, respectively.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Table Number Page

ACKNOWLEDGMENTS

I would like to thank my advisor, Professor Rudell, for giving me the opportunity to work on this project. I have learned a great deal the last three years and sincerely appreciate his depth of knowledge and advice. I also give many thanks to Professor Sathe for his insight and availability to brainstorm and dissect different problems. Professor Gupta from Washington State University provided invaluable feedback on the design and testing of the ADC, I owe many thanks to him.

I am grateful to have had Dr. Kun-Da Chu as a mentor, he was instrumental in helping me develop as an IC designer. I thank my fellow FAST lab members Xichen Li, Yi-Hsiang Huang, Ahmed Aboulsaad, Po-Hao Cheng and alumni Ivan Zhao, Dr. Samrat Dey and Dr. John P. Uehlin for their willingness to help, provide technical feedback and interesting discussion.

Thank you all the Psylab members for sharing your resources and answering my endless questions. Diego Pena-Colaiocco deserves an explicit mention for his role in helping me test the chip and fielding hours of my questions at various times.

This work has many industry sponsors who I thank for their support, including Boeing, Qualcomm, NSF and CDADIC. I am particularly grateful to Mark Steeds for coordinating the Boeing design reviews and updates.

Momona Yamagami, Mark Rafferty, Mareldi Paras-Ahumada, Michael Nolan and Chris Yin, thank you for the tireless work you all do for the ECE department, I appreciate all of you. I thank Ary Prasetyowati and Jennifer Huberman for all their help coordinating purchases and providing graduate school-related guidance.

I would not be here without the support of my parents, Mine and Haldun Dosluoglu, I thank you both. I thank Bianca Wilson, who has been with me every step of my grad school journey and continues to motivate me every day. Finally, I dedicate this thesis to my

grandmother, Oya Alpar, who always gifted me books and encouraged my love of reading.

Chapter 1

INTRODUCTION

Supporting data rates on the order of gigabits/s (Gbps) is one of the key features in next-generation 5G wireless systems [5] for applications ranging from high-resolution video streaming to immersive virtual reality simulation. Achieving these high data rates at carrier frequencies in the sub-6GHz bands (used for 4G LTE) is challenging for several reasons, two of which are listed here: 1.) The bands in this region are crowded, which prevents significant increases in available bandwidth to support high data rates and 2.) Designing wide-bandwidth power- and area- efficient receivers integrated on a single chip is challenging due to high linearity requirements for both the analog front-end and ADC to tolerate multiple interferers and recover high-order modulated signals with a sufficiently low error rate.

There is a growing body of research investigating millimeter-wave (mm-Wave) receivers (RX) to support high signal bandwidths and data rates [6, 7, 8, 9, 10, 11]. mm-Wave frequencies are attractive given that for the same fractional bandwidths as used in the sub-6GHz spectrum, the signal bandwidth increases and the size of the passive resonant devices

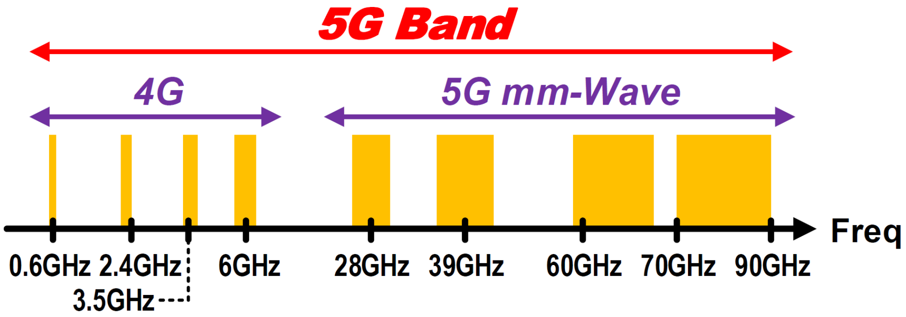

Figure 1.1: 5G band allocation [1]

such as antennas, inductors and capacitors shrink. However, as the carrier frequency ($\mathrm{F_c}$) of the wireless signal increases, the path loss increases proportional to $(\mathrm{F_c})^2$ according to the Friis equation [12]. While the path loss attenuation improves the spatial rejection of mm-Wave interferers and blockers relative to RF frequencies, the desired signal from a single transmitter is also attenuated, which increases its gain requirement while lowering the receiver sensitivity. Furthermore, design at mm-Wave frequencies is more challenging in that the simulations require further complexity, greater care needs to be taken for the physical layout of the chip and expensive equipment is required to test at such high frequencies.

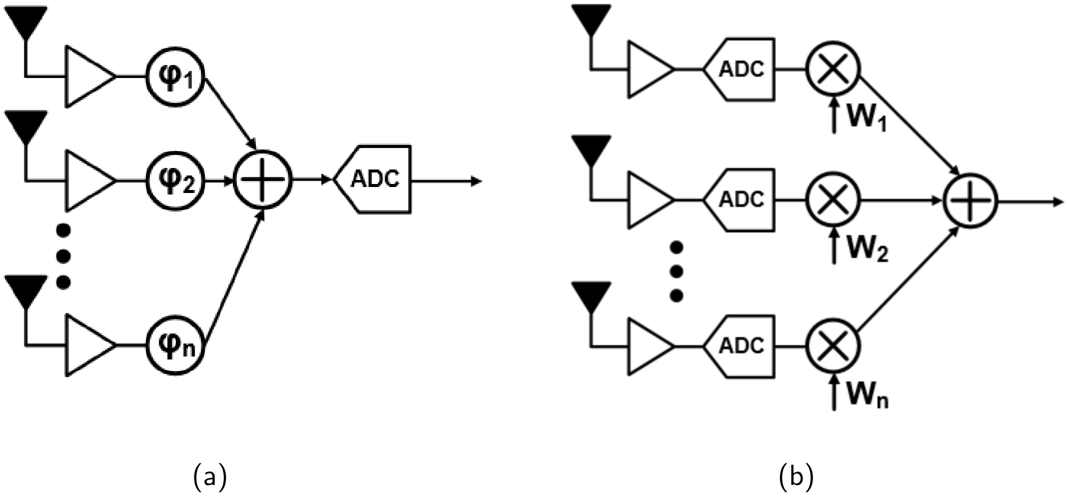

(a) (b)

Figure 1.2: (a) Ideal analog beamformer (b) Ideal digital beamformer

To compensate for the higher path loss, these RXs are typically realized as phased-arrays. Multiple elements are placed in parallel and signals arriving at particular angles can be constructively or destructively summed by adjusting the relative phase shift between each elements. This constructive interference gives rise to the array gain which will be discussed further in Chapter 2.

Prior works have primarily implemented these phased-array systems as analog beamformers, performing the phase shifts in the analog domain prior to quantization as illustrated in Fig. 1.2a. However, this approach comes with several drawbacks: 1.) This architecture can typically only support a single beam, 2.) The phase-shifters can consume a lot of area, as illustrated by the die photo in [13] and 3.) Phase shifting in the analog domain poses more stringent noise and linearity requirements depending on where the shifters are implemented in the RX chain.

An alternative to analog beamforming is digital beamforming (DBF), depicted in Fig. 1.2b. With this architecture, the relative phase shifts are applied in the digital domain. This allows increased flexibility in the beam pattern that can be produced, adaptivity to changing environments and interferers, and independent steering of N beams for N elements. Despite these desirable qualities, the scalability of DBF is limited given that each antenna element requires its own RX and ADC chain, which drastically increases the power and area consumption to support multi-GHz signal bandwidths.

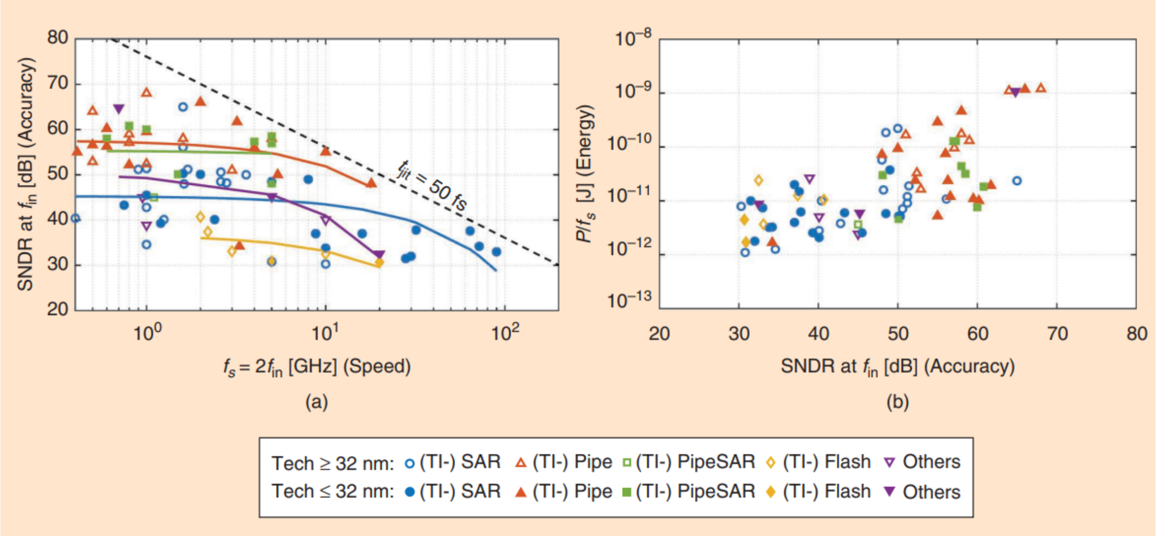

Figure 1.3: (a) State-of-the-art ADC SNDR vs sampling frequency plot from [2]-[3] (b) ADC Energy v. SNDR plot from [2]-[3]

ADCs in particular serve as a bottleneck for the achievable data rate in mm-Wave receivers. Fig. 1.3 from [2] plots the state-of-the-art ADCs published from 1997 to 2021 in ISSCC and VLSI using data from [3]. Fig. 1.3a illustrates that as the sampling frequency of ADCs increases, the resolution is bounded by a given sampling clock jitter.

Successive-approximation register (SAR) ADC architectures have been able to achieve the highest data rates maintaining a moderate effective number of bits (ENOB) of 5-6. The SAR architecture is shown in Fig. 1.4a. It requires B clock cycles to produce B bits. As the input signal level is held, the single-bit quantizer determines the sign of the residue amplitude that is halved every cycle. As shown in Fig. 1.3b, this architecture also has the lowest energy per bit of any other architecture to date. Additionally, it is able to maintain the speed benefits of shrinking feature size since it requires minimal voltage headroom for the single comparator and SAR logic. However, the resolution is in large part limited by the resolution of the digital-to-analog (DAC) converter, and the B-cycle latency may not be attractive for beamforming systems that require rapid adaptivity.

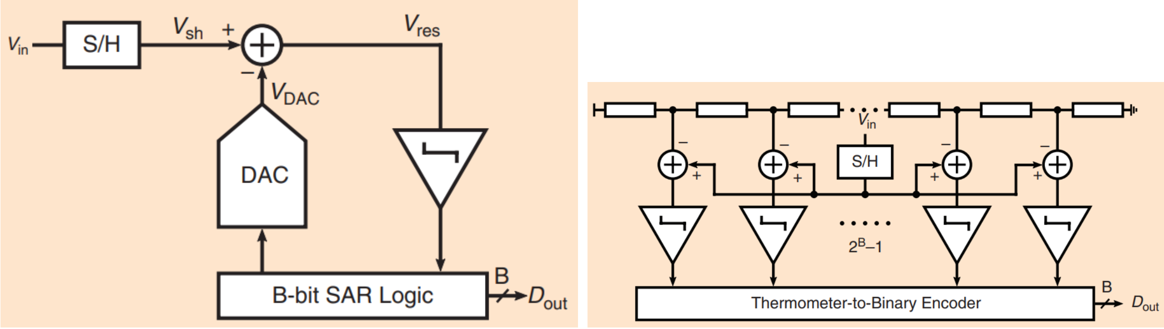

(a) (b)

Figure 1.4: (a) Ideal B-bit SAR ADC from [2] (b) Ideal B-bit FLASH from [2]

Another high-speed, low-resolution and energy-efficient architecture observed in [2] is the FLASH ADC, as illustrated in Fig. 1.4b. The FLASH is the lowest-latency architecture, as it involves 2^B-1 comparators operating in parallel to determine if the input signal is above or below a given threshold level. The comparator outputs are converted to binary bits via a thermometer-to-binary encoder. While this comparator can operate at high speeds similar to the SAR ADC, the power exponentially increases as the resolution is increased (each additional bit requires twice as many comparators). Additionally, the speed of the comparators is limited both due to voltage headroom and the achievable least-significant bit (LSB) while minimizing comparator noise and voltage reference mismatch.

Given that achievable system data rates are proportional to ADC sampling rates (f_s), and the only effective way to reduce the power consumption at high sampling rates appears to be reducing the ADC resolution, two questions arise: (1) What is the minimum ADC resolution required for robust digital beamforming? (2) Can a 1-bit ADC be used in digital beamformers instead of higher resolution ADCs?

This thesis investigates the second question using both MATLAB simulations and silicon measurements. Chapter 2 provides an analysis of 1-bit ADCs in a digital beamforming array and simulation results to determine the achievable ENOB, bandwidth and array gain. Chapter 3 discusses the circuit implementation of the single-bit ADC, as well as blocks in the 4-element mm-Wave DB RX in which it was integrated. Chapter 4 details standalone and system-level measurement results utilizing the 1-bit ADC in the DB RX, and Chapter 5 summarizes the findings and includes possible future directions for this work.

Chapter 2

ADC SYSTEM-LEVEL DESIGN

As discussed in Chapter 1, there is a motivation to use low-resolution ADCs in the DBF Rx architecture to minimize the power and area consumption. A drawback of this approach is the degradation of the phased array gain, which is demonstrated in the following analysis. Assuming every Rx chain is identical:

$$S_{\mathrm{OUT}} = G \cdot n^2 \cdot S_{\mathrm{IN}}$$

$$N_{\mathrm{OUT}} = G \cdot n \cdot (N_{\mathrm{IN}} + N_{\mathrm{Rx}})$$

$$SNR_{\mathrm{Bx}} = n \cdot \frac{S_{\mathrm{IN}}}{N_{\mathrm{IN}} + N_{\mathrm{Rx}}} = n \cdot SNR_{\mathrm{Rx}}$$

where $\mathrm{S_{OUT}}$ is the signal power at the output of the beamformer (Bx), $\mathrm{N_{OUT}}$ is the output noise power, G is the gain of one sub-Rx, n is the number of array elements, $\mathrm{N_{IN}}$ is the input noise power, $\mathrm{N_{Rx}}$ is the noise power of one sub-Rx, $\mathrm{SNR_{Rx}}$ is the SNR at the output of one sub-Rx and $\mathrm{SNR_{Bx}}$ is the SNR at the output of the Bx. The SNR of the Bx is enhanced by N, or the array gain:

$$Array\ Gain\ (dB) = 10 \cdot log_{10}(n)$$

Based on this result, doubling the number of elements improves the peak system SNR by 3 dB. However, the key assumption here is that the noise between all channels is uncorrelated, and the signal is perfectly correlated, so the signal power increases by a factor of n^2 while the noise power only increases by n. This is a reasonable assumption when the ADC quantization noise power is -10 dB or lower than the Rx thermal noise, which consequently only degrades the SNR by $< 0.4\ dB$. However, for a 1-bit ADC, this assumption is no longer applicable since the quantization noise dominates (the quantization noise is ~ 30 dB greater than the Rx thermal noise floor in the example provided in [1] section 4.2). As discussed in [1] Section 4.4, this noise can be correlated with the input signal and between elements. Quantization noise that is perfectly correlated between ADC elements also increases by n^2, and eliminates the array gain.

The addition of a dithering sequence with a rectangular probability distribution function (RPDF) non-subtractively applied to a quantizer has been shown in [14] to de-correlate the quantization noise with the input signal. Linear-feedback shift-registers (LFSRs) are able to generate RPDF sequences, which can be built with exclusive-OR (XOR) gates and flip-flops to implement pseudo-random maximal-length sequences of length $2^{\mathrm{N}} - 1$, where N is the number of flip-flops [15]. Additionally, the auto-correlation $C(\tau)$ of the sequence is $\frac{-1}{2^{\mathrm{N}}-1}$ for $\tau \neq 0$ [15].

Section 2.1 details a study using behavioral MATLAB models injecting minimally-correlated pseudo-random dithering sequences into the reference voltage of a 1-bit ADC to de-correlate the quantization noise between elements and recover the phased-array gain, similar to the idea discussed in [1] Section 4.4. The sequences are based on polynomials that can be implemented as LFSRs (the degree of the polynomial corresponds to the number of flip-flops in the shift register). Section 2.2 briefly discusses the ADC specification for the implemented 4-element DB RX.

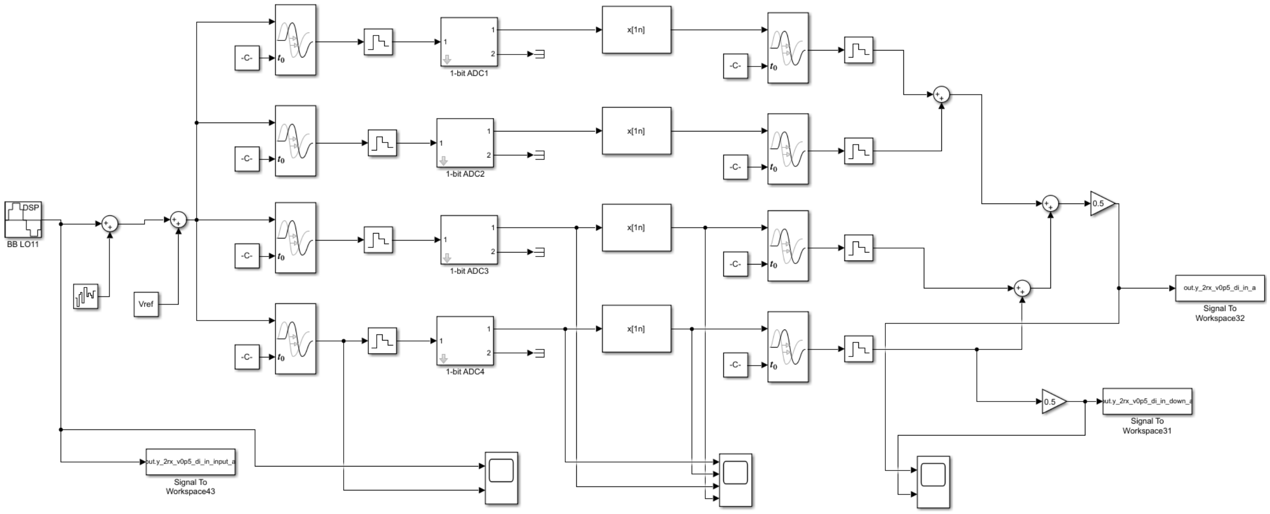

Figure 2.1: Simulkink test bench for 1-bit dithering experiments.

2.1 1-bit ADC behavioral testbench

The Simulink testbench for the 1-bit ADC dithering experiments is shown in Fig. 2.1. An incident baseband signal with additional thermal noise is applied to four 1-bit ADCs. The signals are sampled with a zero-order hold prior to quantization by the ideal single-ended comparator (Fig. 2.2). The signals are then downsampled and filtered by ideal decimation filters and summed and scaled based on an incident angle of 0°. The dither is injected to the reference voltage by a pseudo-random noise sequence generator (Fig. 2.2).

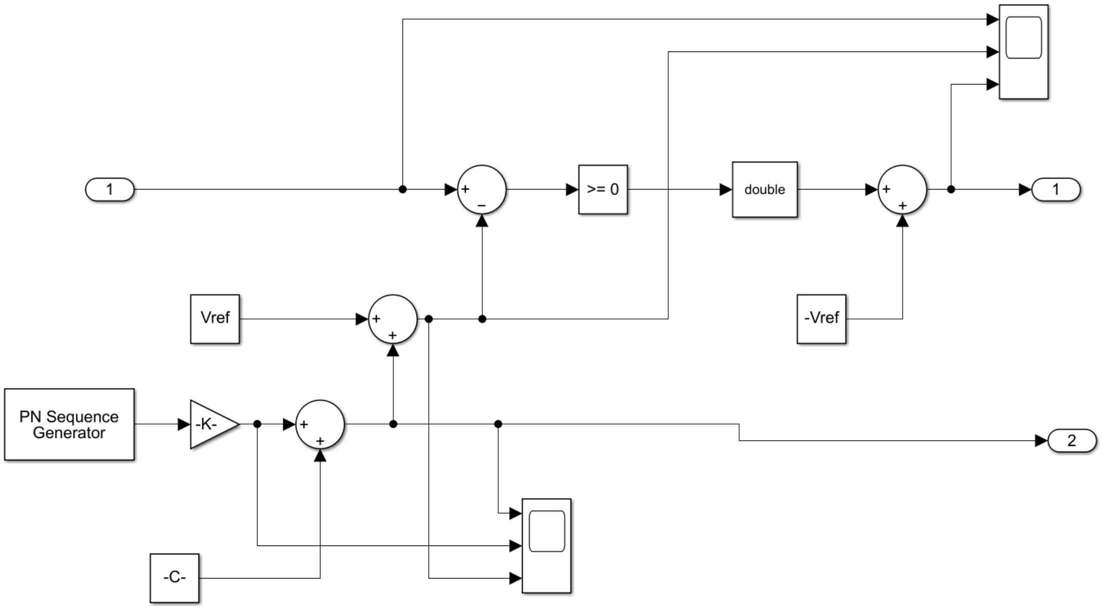

Figure 2.2: Behavioral Model of 1-bit ADC with Vref dithering [1]

Initially, the dithering amplitude was set to zero, and a 31-Hz 0.5 V amplitude tone was applied to the 1-bit ADCs sampled at a 2-kHz Nyquist rate, as shown in Fig. 2.3a. $SNDR_{\text{Bx}} = SNDR_{\text{Rx}} = 6.29\ dB$ was calculated from the samples were collected over a 1-second interval. Then, the dither amplitude was set to 0.5 V, identical to the signal amplitude. The dithering sequence was applied at the sampling rate, and a degree of 11 was selected to prevent the sequence from repeating within the collected sample. The dither resolution was set to be the same as the polynomial degree, 11 bits. Figs. 2.3b and 2.3c illustrate the de-correlation of the noise floor when choosing identical and unique initial conditions, or seeds, for each dithering sequence, respectively. Application of the dither only degraded the SNDR of a single 1-bit ADC to 4.6 dB, without improving the system SNDR, but when seeding each sequence uniquely, the SNDR improved to 7.9 dB at the summed output of four elements. Since a maximal-length LFSR can produce $2^{\mathrm{N-1}}$ unique combination, the following conjecture was formed:

***Conjecture* 1**

If an identical RPDF dither sequence with a unique seed is applied to all $1 - bit$ *quantizers in a DB RX*,

$$N > 2^{\mathrm{n}} - 1$$

where n is the number of elements in the DB RX and N is the degree of the polynomial.

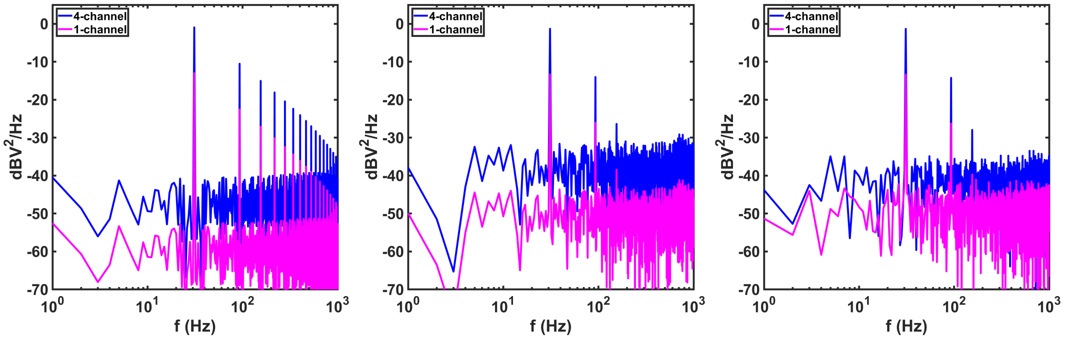

(a) (b) (c)

Figure 2.3: FFT of the (a) ideal 1-bit ADC before and after beamforming with a Full-Scale 31-MHz tone, $\mathrm{F_s} = 2$ kHz, (b) 1-bit ADC with 11-bit full-scale identically seeded dither (c) 1-bit ADC with 11-bit full-scale uniquely seeded dither

Observing Figs. 2.3a and 2.3c, applying dither appears to reduce the signal distortion power while increasing the noise floor. This is confirmed quantitatively with an SDR increase from 6.9 dB to 11.58 dB and an SNR reduction from 15.12 dB to 5.62 dB, respectively.

To make the advantages of the SDR increase more apparent, we must first quickly examine the SNR for an ideal N-bit ADC. The famous equation tells us that $SNR\ (dB) = 6.02 \cdot B + 1.76 + 10 \cdot log_{10}(OSR)$, where B is the number of bits, and OSR is the oversampling ratio [16]. The derivation assumes the quantization noise is uniform, and thus as sampling rate doubles only half of the noise is integrated after filtering, improving the SNR by 3 dB. However, the effective number of bits (ENOB) for a physical ADC is determined by the SNDR, and since the 1-bit ADC is inherently nonlinear, the SDR and thus the SNDR can not be improved by oversampling. If the distortion can be reduced so that the integrated noise power dominates the SNDR denominator, then oversampling can provide some benefit.

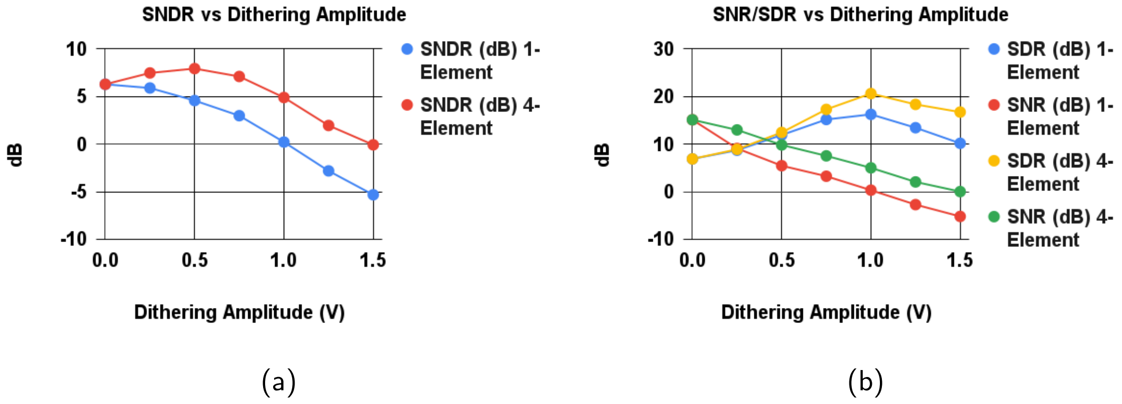

(a) (b)

Figure 2.4: (a) SNDR of the ideal 1-bit ADC before and after beamforming with a 31-MHz tone varying the 11-bit dithering amplitude, $F_s = 2$ kHz, (b) SDR and SNR breakdown

For the second experiment, the dithering amplitude was varied from 0 V to 1.5 V. Initially, it appears in Fig. 2.4a that choosing the dither amplitude to equal the signal amplitude provides the optimum result, as the beamformed SNDR is the highest at this point and increasing the dithering amplitude further steadily reduces the single-element and combined SNDR. The SDR on the other hand, appears to peak at 20.6 dB when the dither amplitude is 1V (Fig. 2.4b), or 2x the signal amplitude, and the array gain at this point is 1.3 dB greater than at 0.5 V, illustrated by the difference between the red and blue curves in Fig. 2.4a. Thus, a 2:1 ratio for the dithering and signal amplitudes was selected for the remaining experiments.

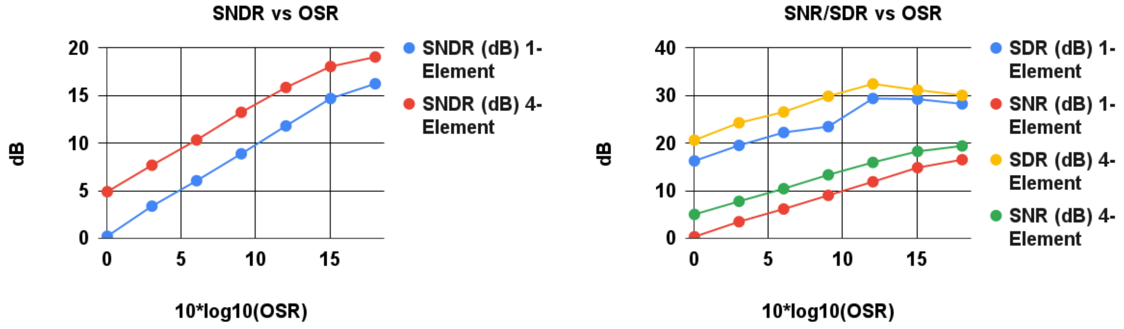

(a) (b)

Figure 2.5: (a) SNDR of the ideal 1-bit ADC before and after beamforming with a 31-MHz tone varying the OSR, (b) SDR and SNR breakdown

To improve the SNR, the OSR was swept from 1 to 64 as shown in Fig. 2.5a. As the OSR increases to 8, the array gain appears to be consistent at roughly 4.3 dB, and gradually shrinking to 2.8 at an OSR of 64. The SDR peaks at 32.4 dB with an OSR of 16, while the SNR linearly increases. However, increasing the OSR decreases the signal bandwidth for a fixed sampling rate, which is undesirable for 5G communication systems that require the bandwidth of a single channel to be 400 MHz [5], so an OSR of 8 is chosen as a compromise.

Observing the FFT of the 1-bit ADC oversampled eight times, additional spurs appear in the noise floor (Fig. 2.6a). This is due to the periodicity of the dithering sequence in the collected sample. As the OSR increases and the post-filtered sample size remains fixed, the period of the dithering sequence relative to the sample time decreases. Therefore, Conjecture 1 needs to be modified:

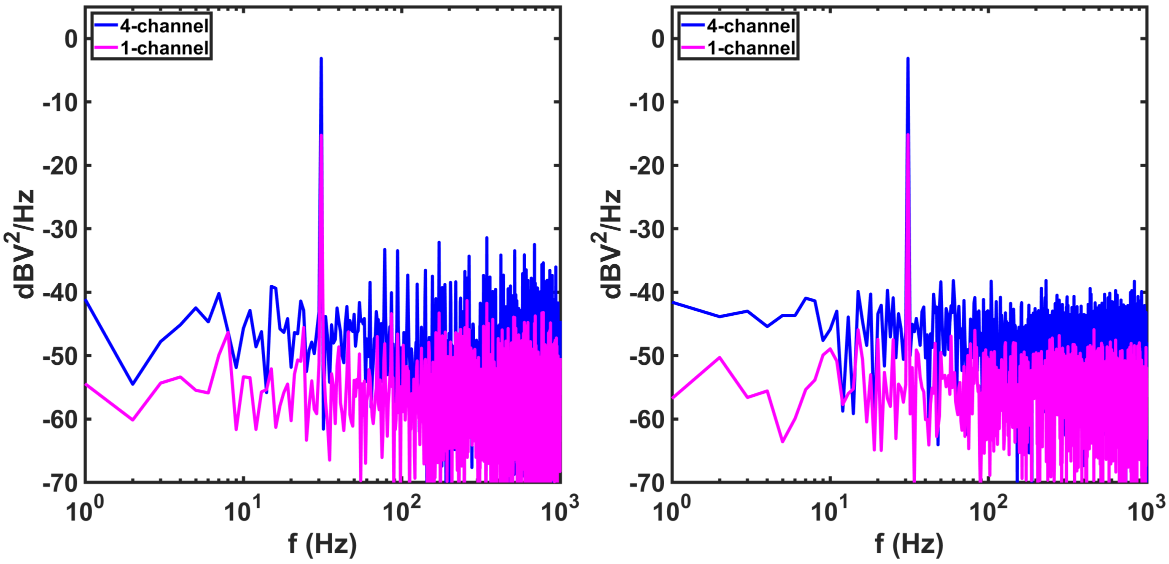

(a) (b)

Figure 2.6: (a) FFT of the 1-bit ADC with 11-bit 2*FS dither, $F_s = 16$ kHz (b) FFT of the 1-bit ADC with 18-bit 2*FS dither, $F_s = 16$ kHz

Revised Conjecture **1**

If an identical RPDF dither sequence with a unique seed is applied to all $1 - bit$ *quantizers in a DB RX every sample,*

$$N > \ max(2^{\mathrm{n}} - 1,\ S \cdot OSR)$$

where S is the decimated sample size. In Fig. 2.6b, applying an 18-bit dither sequence appears to eliminate the spurs and flatten the noise floor while maintaining the beamformed SNDR (13.2 dB) and slightly decreasing the array gain by 0.14 dB. Furthermore, as seen in Fig. 2.7, the resolution does not have to equal the order of the dithering sequence, which would significantly reduce the design complexity of the binary-weighted DAC. Tapping just 3 most significant bits from the generated 18-bit dither resulted in a < 10% SNDR deviation. Also, since the SNDR slightly decreased using a 9-bit to 18-bit dither resolution, this suggests the highest achievable resolution may not be optimal. An intuitive explanation could be that the resolution of an 18-bit dither sequence is too fine to be resolved by the ADC in the given sampling time, resulting in some additional errors. A dithering resolution of 6-bits was selected, sacrificing 0.4 dB of array gain to reduce the DAC complexity.

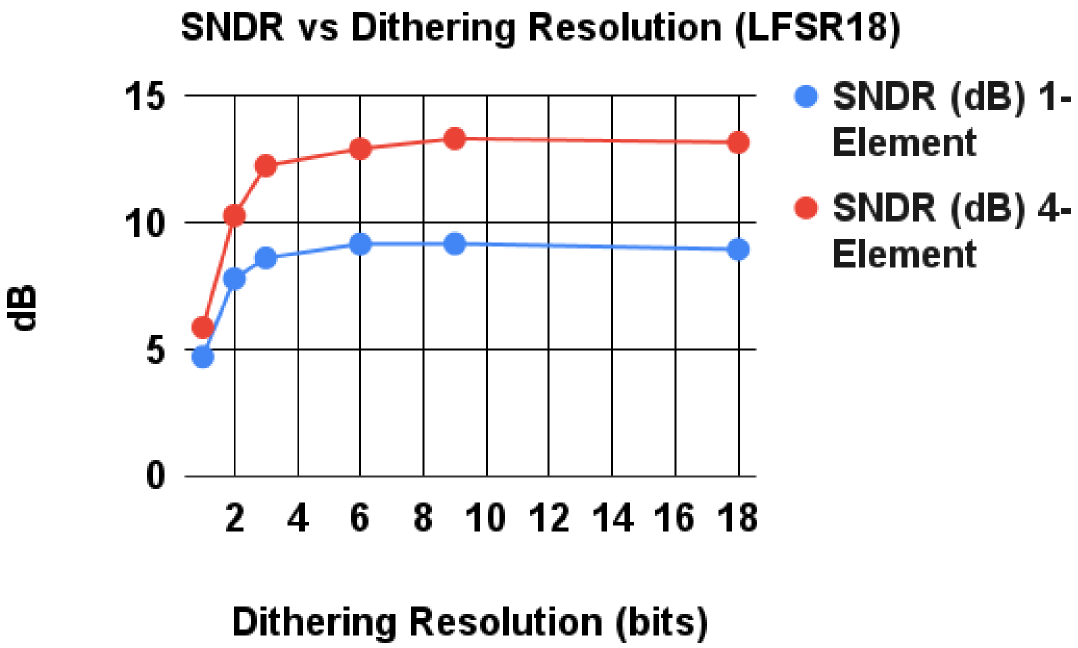

Figure 2.7: SNDR of the ideal 1-bit ADC before and after beamforming with a 31-MHz tone varying the resolution of the 18-bit 2*FS dither sequence, $F_s = 16$ kHz

The dynamic range (DR) of an 8x oversampled and dithered ADC was found to be 25.7 dB at an input frequency of 31 Hz (Figs. 2.8a, respectively). A 2:1 ratio of the dithering and signal amplitudes was maintained as the signal amplitude was reduced. Initially, the dithering resolution of 6-bits was applied at all amplitudes. To further reduce the DAC design complexity, another experiment was run where, as the dithering amplitude was halved, the dither resolution decreased by 1 bit, so the same 6-bit DAC could be used to cover the DR. As seen in (Fig. 2.8a), the "variable res" curves closely track the 6-bit dithered curves. The DR was also measured using a 677-Hz input tone and variable dithering resolution, which is close to identical to the 31-Hz input tone DR. This further demonstrates that the distortion at this design point is negligible, since no harmonics of the 677-Hz tone are present in the half-Nyquist bandwidth, and that the ENOB in the signal bandwidth is approximately constant.

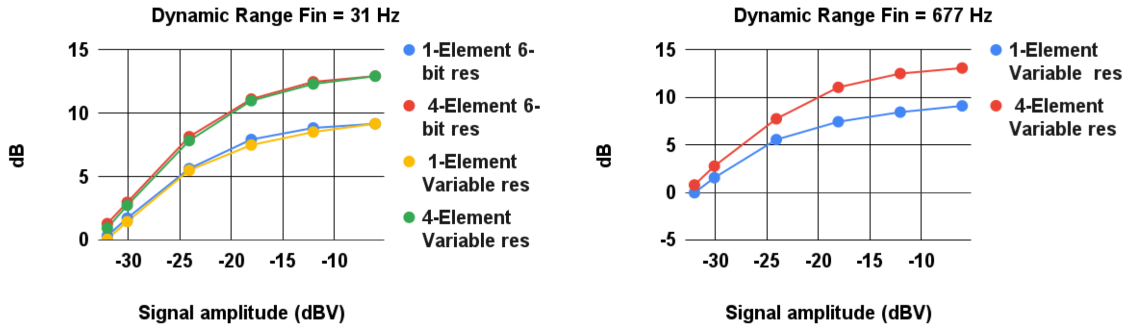

(a) (b)

Figure 2.8: (a) DR of the ideal 1-bit ADC before and after beamforming with a 31-MHz tone, $F_s = 16$ kHz, preserving a 2:1 dither-to-signal amplitude ratio (b) DR with a 677-Hz tone

Finally, this model was expanded to include 16 and 64 elements to determine if the dithering benefits could still be observed. The 1-bit array gain appears to monotonically increase as more elements are added based on Fig. 2.9a, although it deviates further from the ideal gain. The 1-bit array gain is also inconsistent, ranging from 4.39 dB with 16 elements to 2.95 dB with 64 elements. The beamformed SNDR closely matches the SNR (Fig. 2.9b), although a further analysis is required to determine if there is an SNR limit of the dithered 1-bit beamformed system after which adding more elements does not improve the array gain.

2.2 ADC Specifications

As part of a wireless communication channel an ADC should have a sufficiently high SNDR to recover a modulated signal with an acceptable error rate. The minimum error vector magnitude (EVM) specified by [5] for user equipment operating in mm-Wave bands is -

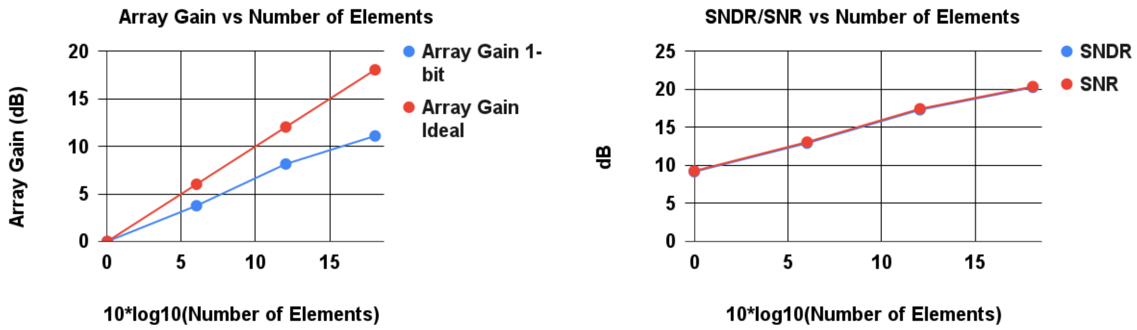

(a) (b)

Figure 2.9: (a) The array gain of an ideal 1-bit ADC applying a 2*FS 6-bit dithering sequence, $F_s = 16$ kHz (b) Beamformed SNDR/SNR versus number of elements

15 dB and -18 dB for QPSK and 16-QAM, respectively. From the derivation in [17] we can determine the minimum system SNR required from the approximation $SNR\ (dB) \approx -EVM\ (dB)$.

As a proof-of-concept to demonstrate digital beamforming with low-resolution ADCs, a 4-element RX was proposed in [1]. Based on Fig. 2.5a a four-element digital beamformer using 1-bit ADCs can achieve an SNDR >18 dB with an OSR of 64, but this is unattractive for a number of reasons:

- To support a maximum channel bandwidth of 400 MHz [5], this would require the sampling rate to be at minimum 51.2 GHz. To date, only 5 ADCs with sampling rates $\geq$ 51.2 GHz have been published in ISSCC, all of which consume $\geq$ 200 mW of power [3]. The high power consumption per ADC prohibits the scalability of the digital beamforming architecture, which was the original motivation of pursuing a 1-bit ADC.

- There are ADCs that achieve bandwidths >400 MHz and SNDRs >27 dB operating at sampling rates <1 GHz and consuming less area and power [3], so in this context a 64x oversampled 1-bit ADC doesn't appear to have any advantages.

Selecting a lower OSR of 8, based on the behavioral study, will not achieve a system SNR of >15 dB. For 16-QAM, the ADC will need additional margin to cover the non-unity amplitude peak-to-average power ratio (PAPR). Thus, to maintain the signal bandwidth the sampling frequency will need to be increased further, which is undesirable. Based on [1] Section 5.1.2, the minimum dynamic range (DR) required for 16-QAM is 31 dB, including 5 dB for peak-to-average power variation and 8 dB of backoff so the ADC noise floor is lower than the receiver thermal noise floor and only degrades the system SNR by 0.6 dB. This specification assumes DR = SNDR.

An architecture that can achieve this SNDR using a 1-bit quantizer and a relatively low oversampling rate is a continuous-time delta-sigma modulator (CTDSM), which has an SNR that can be expressed as follows:

$$SNR_{\text{dB}} = 3.01n(2L+1) - 9.36L - 2.76$$

where L is the loop order and n is $\log_2$(OSR) [18]. According to Table 2.1 generated from this equation, the theoretical limit of a third-order CTDSM oversampled 8 times is 32.4 dB.

Thus, a third-order CTDSM was designed in parallel to meet the 16-QAM specification, the details of which can be found in [19]. The proposed reconfigurable ADC for the 4-element DB RX consists of both the CTDSM and a 1-bit ADC to verify the behavioral simulation, and is described in detail in the following chapter.

Table 2.1: CTDSM ADC SNR as a function of OSR and loop order [1]

		OSR (2^n)				
		2^1	**2^2**	**2^3**	**2^4**	**2^5**
Loop Order	**1**	-3.09	5.94	14.97	24	33.03
	2	-6.43	8.62	23.67	38.72	53.77
	3	-9.77	11.3	**32.37**	53.44	74.51
	4	-13.11	13.98	41.07	68.16	95.25
	5	-16.45	16.66	49.77	82.88	115.99

Chapter 3

CIRCUIT IMPLEMENTATION

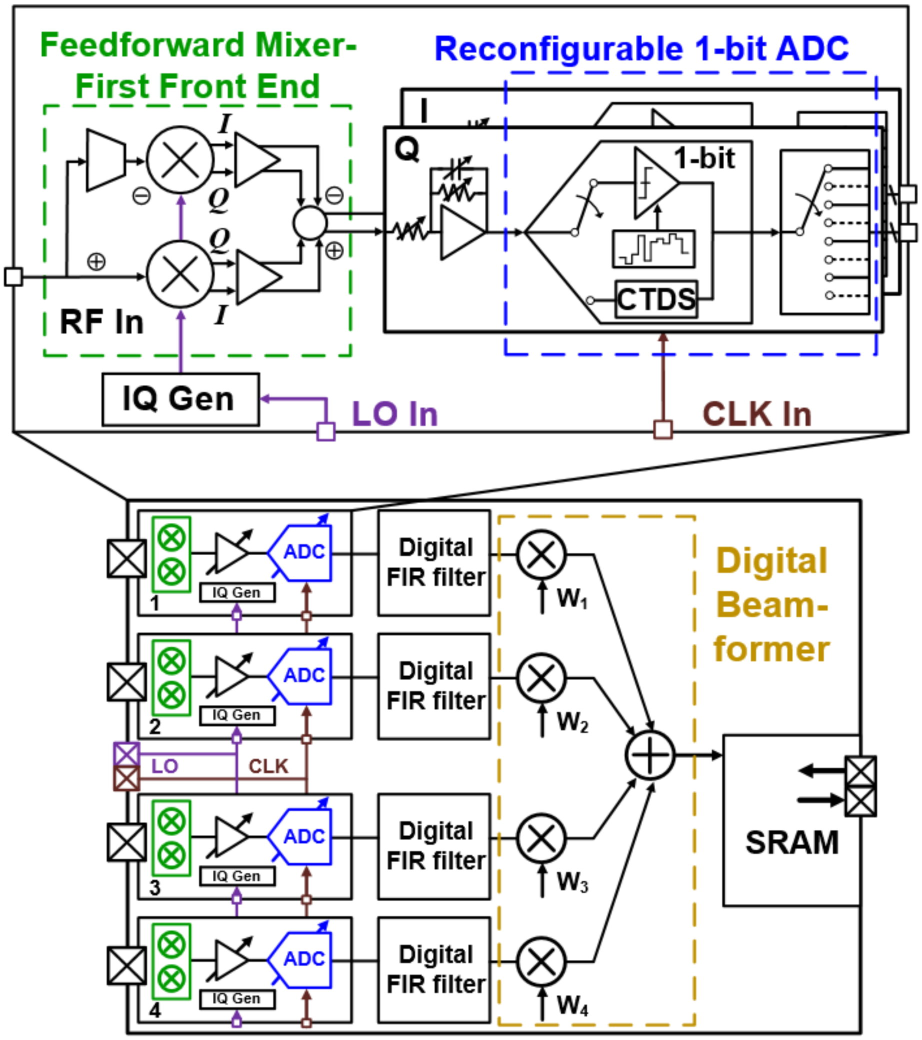

Figure 3.1: Block-level diagram of the digital beamforming mixer-first RX.

The block diagram of the designed 4-element mm-Wave DB direct-conversion (DC) RX with reconfigurable 1-bit ADCs is shown in Fig. 3.1. Each element includes a feedforward mixer-first front-end (MFFE) (described in [1, 20]), baseband amplification with low-pass filtering, a reconfigurable 8x oversampled (OS) 1-bit ADC, and custom digital timing blocks to interface with the synthesized on-chip digital beamformer (BX) and data demodulator. Both the local oscillator (LO) and high-speed ADC clock (CLK) are supplied from off-chip, then routed to each of the four RX elements using an H-tree distribution strategy. This RX was designed to operate with a 45-70 GHz LO and an 8-GHz CLK to achieve a maximum data rate of 666 MS/s.

CLK is frequency-divided down to Nyquist rate and provided to the digital back end (DBE). Data from the output of the decimator, finite-impulse response (FIR) filter and beamformer can be stored on a 0.5 Mb SRAM. This work uses the beamformer with 10-bit weights described in [21].

3.1 Mixer-First Front-End

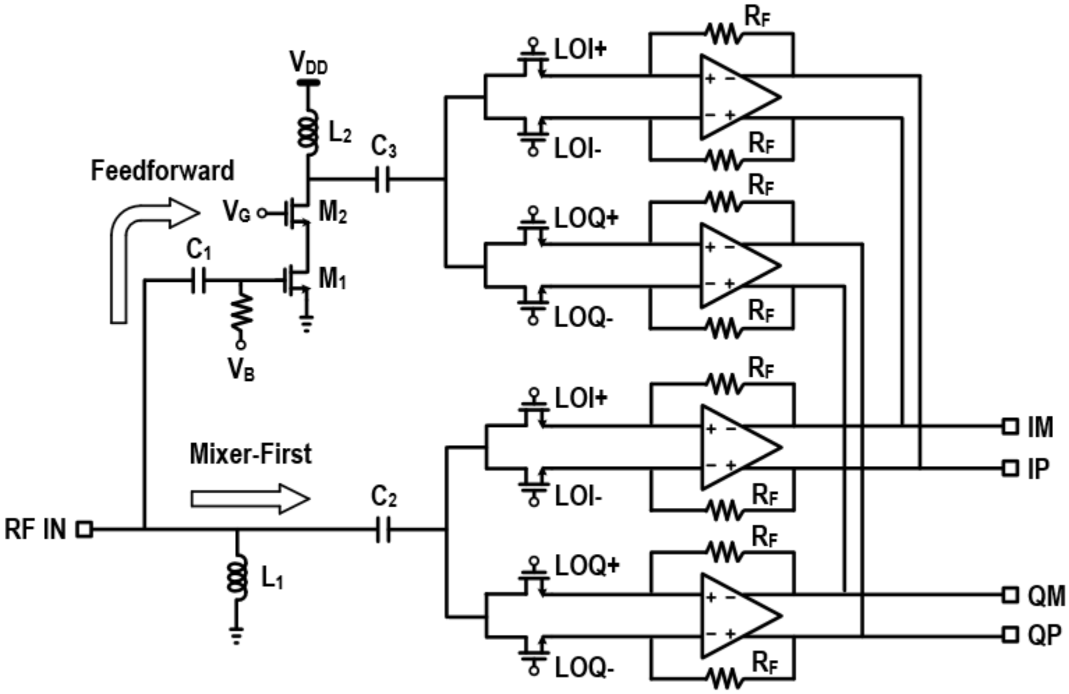

Figure 3.2: Mixer-first Direct-Conversion Front-End with an Auxilliary Gm path.

The MFFE is employed to improve both the power and area efficiency while maintaining an acceptable linearity across a wide bandwidth compared to the conventional DC RX architecture [22]. All of these attributes are desired to implement an analog front-end in a massive DB array that both achieves high data rates and performs spatial filtering in the digital domain. A feedforward path is utilized to improve the front-end (FE) gain and compensate for the relatively poor noise performance of MFFEs. This RX FE has two paths to amplify the signal: the main mixer-first (MF) path and the auxiliary feedforward Gm path. The mixer-first path is designed to perform the input match while the Gm path provides additional gain. The signal in the main MF path is summed constructively with its amplified replica at the BB output (Fig. 3.2). In contrast, the noise contributed by the equivalent 50-ohm impedance from the MF path presented at the input has the opposite phase at the Gm input, and is thus suppressed at the transimpedance amplifier output after subtraction. This method is similar to the RF noise-canceling approach described in [23].

3.2 Reconfigurable ADC

The reconfigurable ADC can operate in two modes (see Fig. 3.3 and Fig. 3.6): (1) as an open-loop oversampled 1-bit digitizer in scenarios where low resolution is permissible (e.g. when the input signal is QPSK-modulated [24]) and (2) as a 3rd-order continuous-time delta-sigma modulator (CTDSM) which can achieve a higher resolution as needed. The samplers in both ADC modes are 2-way time-interleaved (Fig. 3.3) and clocked at 4x the Nyquist rate to achieve an oversampling rate (OSR) of 8. The following sub-sections provide greater detail of the circuit design of the illustrated blocks in Figs. 3.3 and 3.6

3.2.1 Sample-and-Hold

The sample-and-hold holds the analog input sampled at a given time instant. The time constant during the track phase should be small enough to converge to the input signal level within half the sampling clock period, while output sampling capacitor should be large enough to hold the value during the other half period. The analysis in [4] finds the relationship between the normalized dominant pole P_1 and transconductance of the amplifier

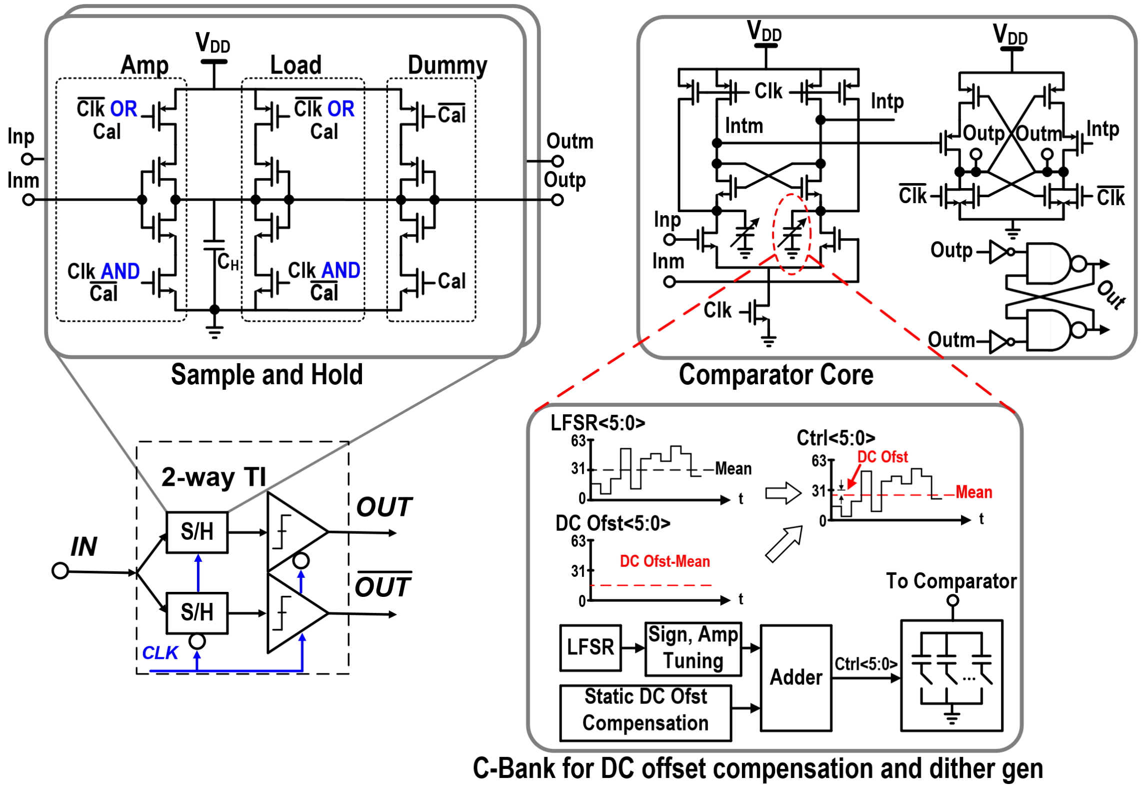

Figure 3.3: High-speed comparator implementation with offset cancellation and dither injection circuitry used for the 1-bit open-loop ADC mode.

g_{m1} in a conventional buffered switch-based track-and-hold circuit (Fig. 3.4a) expressed by the following equation:

$$\frac{P_1}{2\pi f_T} = \frac{1}{\frac{15}{4} + \frac{2\pi f_T C_L}{g_{m1}} + 2\sqrt{\frac{3}{2} \cdot \frac{2\pi f_T C_L}{g_{m1}}}}$$

where f_T is the unity-gain current frequency of the NMOS transistor. Plotting this equation in Fig. 3.5a, [4] visually demonstrates the steep increase in power (proportional to g_m) required to increase the dominant pole. Intuitively, as the widths of both the buffer and switch in series increase, so does the capacitance on both the input and output switch nets, requiring greater power consumption to drive the load presented by the buffer and

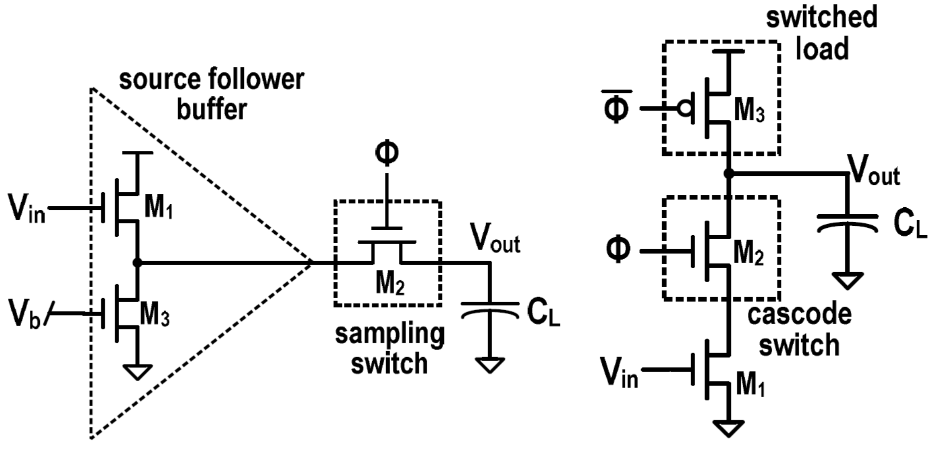

(a) (b)

Figure 3.4: (a) Conventional sample-and-hold buffered sampling switch [4] (b) Cascode sample-and-hold proposed in [4]

switch in series.

The authors in [4] propose a cascode-based track-and-hold to reduce the power-consumption at high bandwidths, shown in Fig. 3.4b. With this topology, the dominant pole is set by the triode PMOS resistance and the output load capacitance, with the size of the NMOS sampling switch negligibly impacting the output impedance and P_1 (based on a first-order approximation). The derived equation for this cascode sampler is as follows:

$$\frac{P_1}{2\pi f_{\mathrm{T}}} = \frac{1}{\frac{5}{2} + \frac{2\pi f_{\mathrm{T}} C_{\mathrm{L}}}{g_{\mathrm{m1}}}}$$

The g_{m1} and, consequently, power increase as shown in Fig. 3.5a is approximately linear for a significant fraction of the device f_{T}, and appears much more attractive for high-bandwidth sampling applications. However, as later discussed in [4], the non-linearity of the input device g_{m} along with the limited headroom presents a trade-off between the achievable linearity and the input swing of the sampler. This is mitigated by replacing the PMOS triode load with an NMOS in saturation, but in doing so the supply voltage must be raised to 1.6V so all 28-nm stacked devices are biased in saturation.

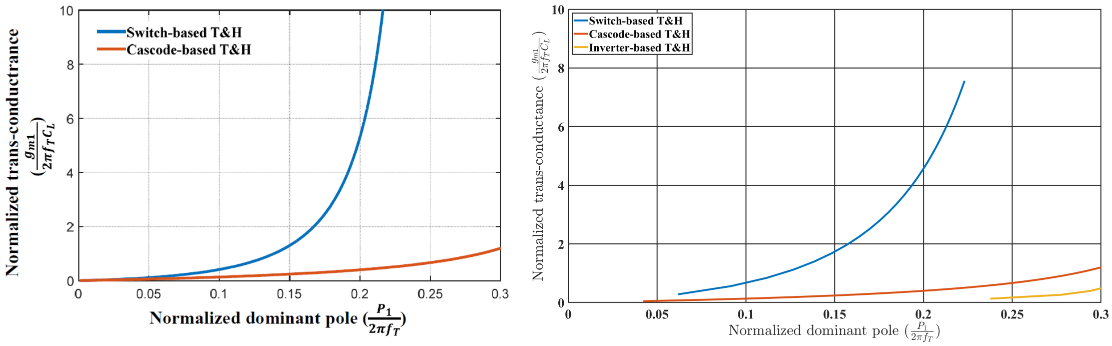

(a)

(b)

Figure 3.5: (a) Normalized $\mathrm{g_{m1}}$ vs. $\mathrm{P_1}$ from [4] for switch-based and cascode-based track-and-hold (b) Replicated plot from (a) including inverter-based track-and-hold

As an alternative to the cascode topology, an inverter-based track-and-hold circuit described in [25] was implemented, as shown in Fig. 3.3. In this topology, the output gain is set by the parallel $\mathrm{g_m}$ ($\mathrm{g_{m,eff}}$) of the PMOS and NMOS devices in the amplification branch multiplied by the output impedance which is dominated by the $\frac{1}{g_{\mathrm{m,eff}}}$ impedance of the load branch with diode-connected devices. When Clk is low, the PMOS and NMOS switches in both branches are turned off, so the charge on the capacitor $\mathrm{C_H}$ is held as it has no discharge path. The load branch also determines the output voltage bias point, obviating the need for current-biasing, and only requires two stacked devices to be in saturation instead of four in the cascode (including the three shown in Fig. 3.4b plus an additional tail device to provide a current bias), thus eliminating the headroom-linearity tradeoff. In addition to the amplifier and load branches, a dummy branch was added to present a stable impedance for the input of the comparator during offset calibration while the track-and-hold is turned ”off” via clock gating.

The switch-, cascode- and inverter-based topologies were simulated on the schematic-level in Cadence Virtuoso using ultra-low $\mathrm{V_t}$ RF devices from the TSMC 28-nm RF+HPC+ process design kit (PDK) to determine each normalized $\mathrm{g_{m1}}$ vs normalized $\mathrm{P_1}$. Multiple points were collected for the plots by sweeping the amplifier and switch sizes for the same $\mathrm{C_L}$. Based on the derivation in [25], the normalized $\mathrm{P_1}$ for the unity-gain track-and-hold can be written as:

$$\frac{P_1}{2\pi f_{\mathrm{T}}} = \frac{C_{\mathrm{gI}}}{3C_{\mathrm{gI}} + C_{\mathrm{L}}}$$

As seen in Fig. 3.5b, as compared to the switch- and cascode-based track-and-hold topologies, the inverter-based track-and-hold achieves a lower power consumption than both up to $\frac{3}{10}f_{\mathrm{T}}$. Thus, this topology was selected to ensure it would not be the bandwidth bottleneck in the 1-bit oversampled ADC.

Based on post-extracted simulation results, the -3dB bandwidth of the track-and-hold in "track" mode is 19 GHz. The worst-case gain when applying a 16-GHz clock is -2.2 dB. It achieves a worst-case IP1dB of -2.2 dBm and IIP3 of 10.2 dBm. The equivalent kT/C noise is 0.63 $\mathrm{V_{RMS}}$ and the worst-case power consumption is 1.9 mW at $\mathrm{V_{DD}} = 0.95$ V.

3.2.2 Comparator

The topology of the implemented two-stage latched comparator is shown in Fig. 3.3. This architecture utilizes stages described in [26] and [27] for the pre-amplification and latch stages, respectively. Both stages consume no static power. A cross-coupled pair was stacked on top of the two input devices in order to further amplify the voltage difference $\mathrm{V_{XY0}}$ to improve the regeneration time t of the 2nd-stage latch. This relationship is described by the following equation from [28]:

$$t = \tau_{\mathrm{reg}} ln \frac{V_{\mathrm{DD}}}{V_{\mathrm{XY0}}}$$

where τ_{reg} is the regeneration time constant.

Additionally, including the cross-coupled pair creates low-impedance nets at the drains of the input devices. This allows an additional tunable C-bank to be added for static offset tuning while minimizing the additional delay for an acceptable tuning range. This C-bank was also used to inject the dither sequence described in 3.2.3 and was simulated to achieve a maximum and minimum offset tuning range of 14.6 mV and 0.65 mV with 6-bits of resolution, respectively. It should be noted that increasing the dither amplitude involves increasing the capacitive load of the comparator pre-amplfier which, in turn, reduces the maximum F_s and increases the minimum comparator offset. Finally, the top PMOS switches are used to reset the internal nets of the pre-amplifier so the input pair is biased in saturation at the start of the comparison and the input devices to the second-stage latch are turned off.

The modified latch from [27] is used in order to ensure the latch does not start drawing current until the intermediate nets from the pre-amp stage are sufficiently discharged, reducing the power consumption compared to the conventional double-tail described in [29]. The NMOS devices in parallel to the cross-coupled NMOS pair are enabled during reset so the latch output returns to ground. The set-reset (SR) latch is used to hold the output of the comparator while the internal nets are being set. The inverters preceding the SR latch are implemented using High-V_t devices to reduce the probability of a false decision.

Based on the track-and-hold simulations, the full-scale comparator input was set at 200-mV peak amplitude. While the least-significant bit (LSB) of a Nyquist-rate 1-bit ADC is equal to the full-scale amplitude A_{FS}, an oversampling 1-bit ADC is required to resolve finer amplitude inputs around the zero-point of the maximum input sine wave. Assuming the sampling frequency is 16 times the input tone frequency (OSR = 8) and the two sampling time instances are $\pm\frac{\pi}{2OSR}$, then $A_{\text{FS}}sin(\frac{\pi}{2OSR}) = 39mV$. Since the sampling instances will not be perfectly symmetrical due to clock jitter, the LSB was reduced to 35 mV. Given the LSB and V_{FS}, an overdrive recovery test (ORT) as described in [26] was performed to determine the decision time and reset time. The worst-case post-extracted decision time and reset time were 52.9 ps and an 18.9 ps, respectively. Given that the output needs to be regenerated within half a clock cycle, the maximum comparator clock frequency was determined to be 8 GHz. In order to achieve the sampling rate of 16 GHz, the comparators were two-way time interleaved.

The worst-case post-extracted power consumption for the comparator core (not including the SR latch and inverters) was 2.8 mW at $V_{\text{DD}} = 0.95V$, with a maximum input-referred noise of 125 uV_{RMS}. At 8 GHz, the probability of error was determined to be $1.65 \cdot 10^{-4}$, which exceeds the acceptable bit-error rate required for communication systems [28]. Based on the previous latch equation, substituting t for $\frac{T_{\text{CLK}}}{2}$ where T_{CLK} is the clock period, we can write:

$$T_{\text{CLK}} = 2\tau_{\text{reg}} ln \frac{V_{\text{DD}}}{V_{\text{XY0}}}$$

In order to reduce the probability of metastability within a given half clock cycle, we can increase the clock period at the expense of reducing the achievable signal bandwidth.

3.2.3 Dither Generator

Based on the results of Section 2.1, a 6-bit LFSR sequence was applied to the comparator C-bank, shown in 3.3. The length of the LFSR is adjustable to 6 and 18, while combinational logic was inserted to adjust the amplitude as well as the drain net of the comparator input device to control the sign. A 6-bit adder prior to the C-bank was used to adjust the dither offset based on the comparator static offset.The dither sequence is clocked at the same rate as the comparator (maximum 8 GHz), and consumes 9.3 mW at a V_{DD} of 0.95 V based on simulation.

3.2.4 3rd-Order Continuous-Time Delta-Sigma ADC

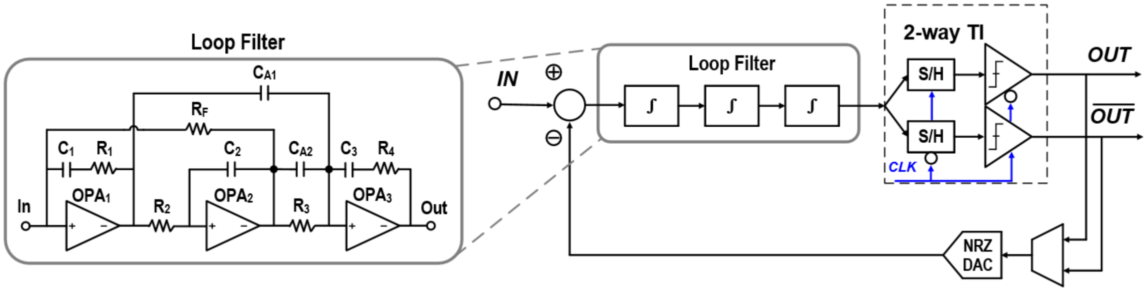

Figure 3.6: ADC in CTDSM Mode.

The 3rd-order CTDSM improves the resolution of a 1-bit quantizer utilizing feedback to shape the in-band quantization noise and is described in further detail in [19]. A capacitive feedforward architecture was implemented for the loop filter. The excess loop delay is compensated by tuning resistors R_1, R_4 and R_f (Fig. 3.6).

Chapter 4

MEASUREMENTS

The 4-element DB chip was fabricated in TSMC 28HPC+RF CMOS process and is shown in Fig. 4.1. Each sub-RX consumes an area of 0.18 mm^2, including I/O pads. In addition to the four receivers, DBE and LO and clock routing, the chip includes pads and test structures to characterize the MFFE and ADC.

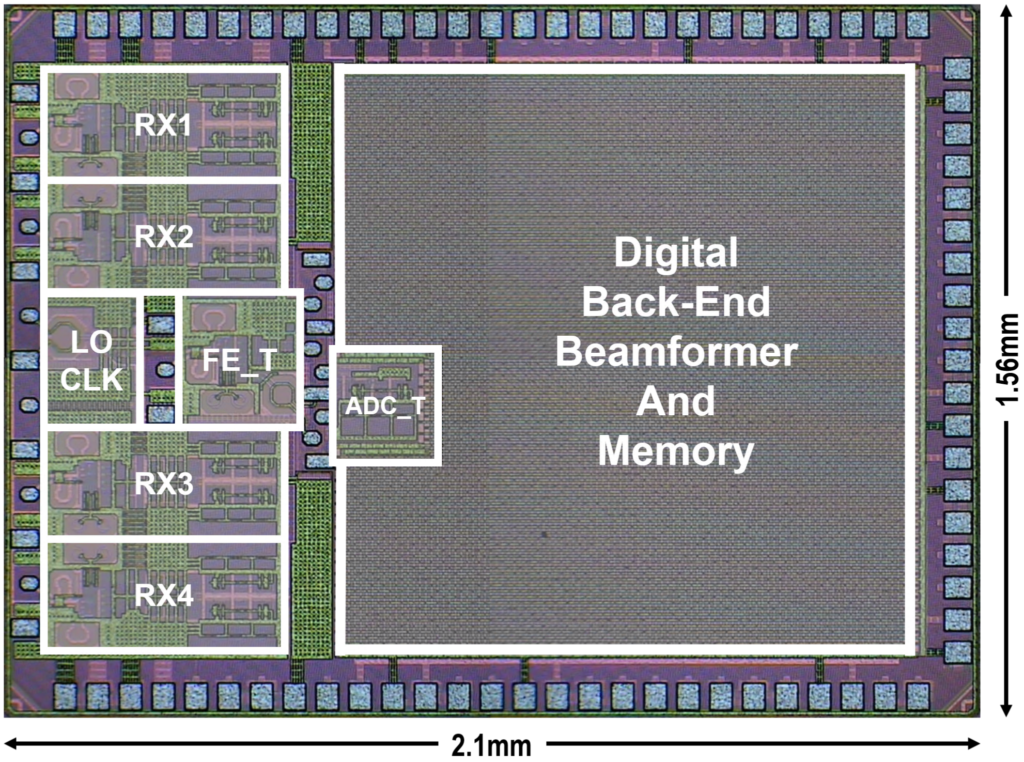

Figure 4.1: Die Photo of the 28nm CMOS 4-element DB, total area 3.3 mm^2.

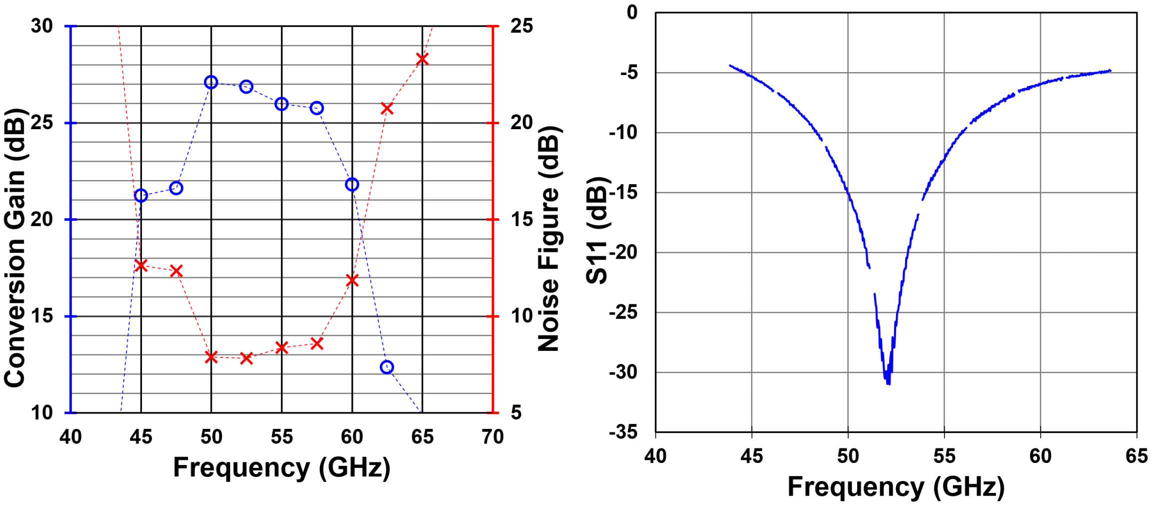

(a) (b)

Figure 4.2: (a) Measured NF and gain from a single-element front-end RX AFE (b) Single-element measured S_{11}

4.1 *Front-End Test Structure*

The RX FE operates from 48 GHz to 60 GHz achieving a maximum gain of 27.1 dB and a minimum noise figure (NF) of 7.8dB (Fig. 4.2a). An $S_{11} < -10dB$ was measured over an 8-GHz bandwidth (Fig. 4.2b). An input $IP3 > -18dBm$ was measured over the bands of interest.

4.2 *ADC Test Structure*

Fig. 4.3a illustrates the effect dithering has on a 1-bit quantizer in a four-element array. A full-scale input 10-MHz tone was sampled at 2.4 GHz. Without dithering, the distortion noise dominates. The signal distortion is reduced after applying the maximum amplitude 6-bit LFSR6-generated dither sequence at the cost of raising the noise floor (see Table 4.1). However, since the period of this dither sequence is less than the pre-decimated sample size dithering spurs can be observed in the spectrum. To ensure the injected noise spectrum is flat, LFSR18 is utilized to generate a sequence with a period greater than the maximum sample that can be stored on the chip. Only six of the 18 bits are tapped and applied to the C-bank.

The measured magnitude of the frequency response for a 100-MHz full-scale input tone is shown in Fig. 4.3b for the 3rd order CTDSM. A maximum signal-to-noise and distortion ratio (SNDR) of 21 dB was achieved.

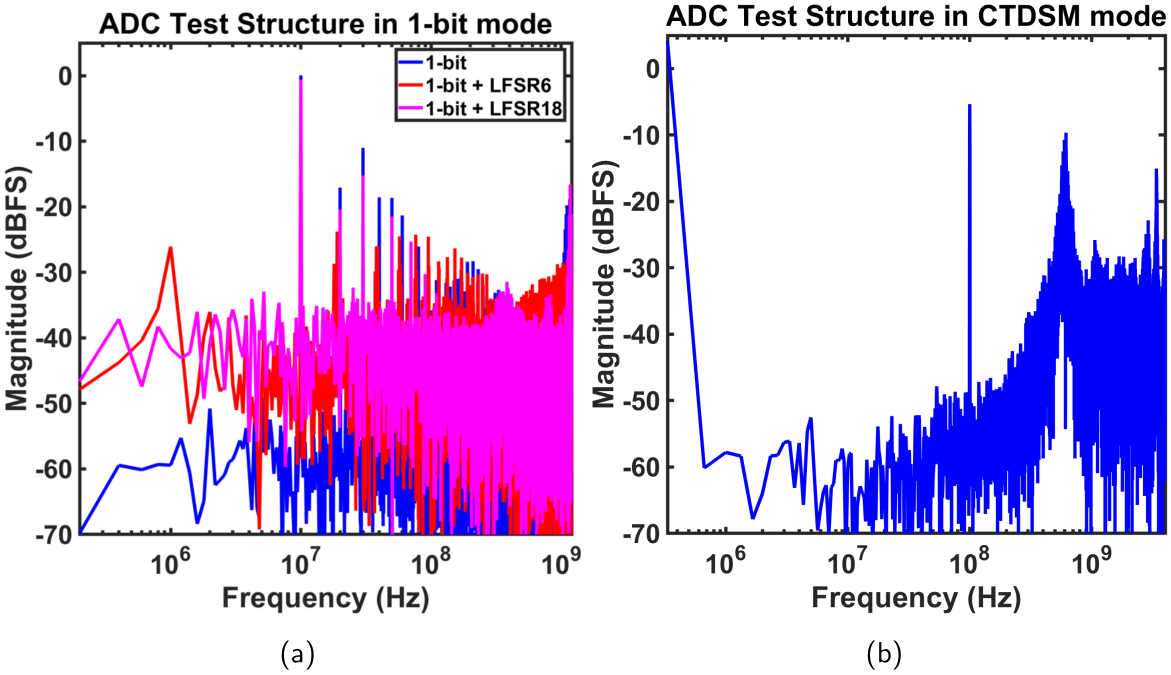

(a)

(b)

Figure 4.3: (a) Measured Output Spectrum of ADC test structure in 1-bit mode with a Full-Scale 10-MHz tone, $F_s = 2.4$ GHz, (b) Measured Output Spectrum of ADC Test Structure in CTDSM Mode with Full-Scale 100-MHz tone, $F_s = 8$ GHz

4.3 System-Level Measurements

The setup for the phased-array measurement is shown in Fig. 4.4. An AWG was used to generate modulated signals with an intermediate frequency (IF) of 3 GHz. The phase differences between the RX inputs were generated using programmable delay lines to emulate the different delays associated with the incident angles of incoming signals, with respect to

Table 4.1: SNDR and Power Breakdown of Fig. 4.3a

Dither Mode	SNDR (dB)	Signal Bin Power (dBm)	Total Distortion Power (dBm)	Total Noise Power (dBm)
No Dither	9.1	-24.9	-34.2	-50
LFSR6	9.5	-25.7	-39	-37.5
LFSR18	10.2	-25.7	-38.8	-39.2

the 3-GHz IF. The IF signals were then up-converted to the carrier frequency (52.5 GHz) and supplied to the left side of the chip (Fig. 4.1) along with the RX LO and CLK using a custom 13-contact probe.

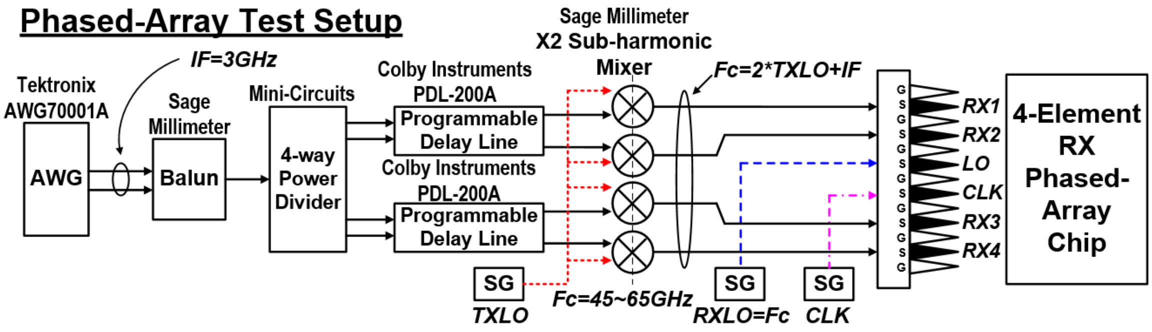

Figure 4.4: Test Setup for System Level Phased-Array Measurements

The impact of dithering the ADCs in 1-bit mode is illustrated in Fig. 4.5 using data collected from the on-chip memory post-beamforming. Without dithering, the measured array gain for the open loop ADCs is less than 6 dB for I and Q. Applying the maximum amplitude LFSR18 dither sequence restores the array gain, as illustrated by the +6-dB SNDR boost post-beamforming at 0°, Table 4.2. Each LFSR is seeded differently so the dither sequences applied to each ADC are uncorrelated.

Two beam patterns steered at 0° and 45° (Fig. 4.6) were measured using a 200-MHz input tone in 5° steps and using post-beamformed chip data in both the 1-bit + LFSR18 and CTDSM ADC modes. A 400-MS/s QPSK signal applied at 0° was recovered at the output of the ADC decimation filters. After an unsigned to signed conversion off-chip, the

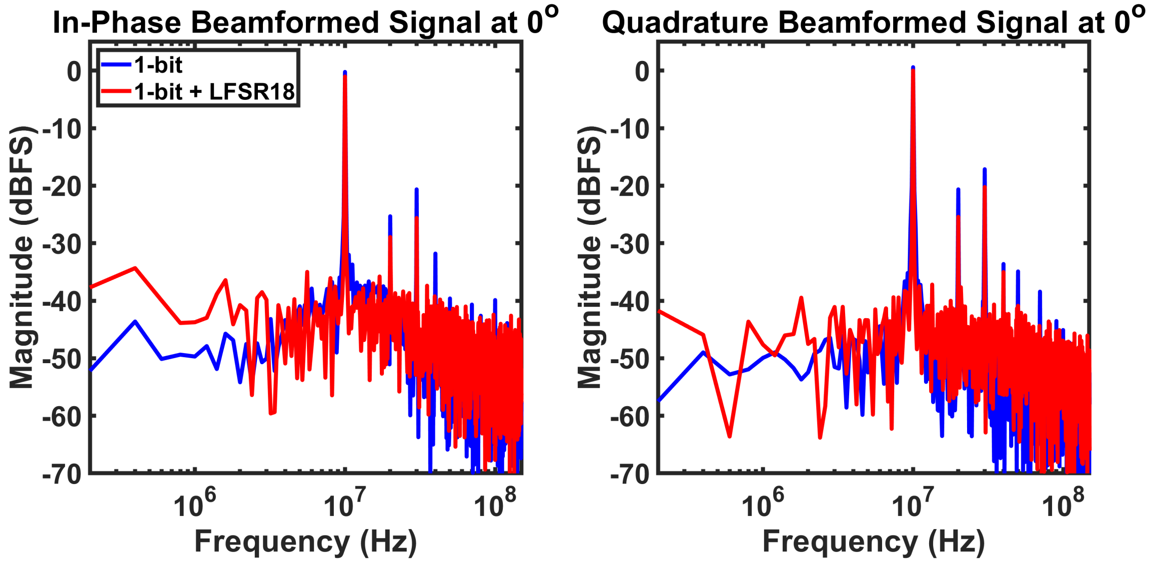

Figure 4.5: Measured spectra of beamformed In-Phase and Quadrature channels applying a 52.51 GHz tone at RF IN and enabling the 1-bit ADC mode

data was fed back and beamformed on-chip. The beamformed output was then demodulated in MATLAB. The constellations using ADCs in 1-bit + LFSR18 and CTDSM modes are shown in Fig. 4.7, with calculated EVMs of -13.4 and -16.2 dB, respectively. Per-element power consumption is illustrated in Fig. 4.8. The measured power consumption of the ADC in 1-bit + LFSR18 (Dither) mode is eclipsed by that of the custom digital timing blocks required to downsample the ADC bits prior to being processed by the DBE. However, the ADC itself is not a power bottleneck in scaling this architecture. The DBE power consumption during beamforming is 373.2 mW with a VDD = 1.2 V and clock rate of 1.2 GHz (divided down from the 4.8 GHz CLK supplied at the FE). Table 4.3 compares this work with prior mm-Wave phased-array RXs.

Table 4.2: SNDR Before and After Beamforming

	1-bit (mean)	1-bit BX	1-bit + LFSR18 (mean)	1-bit + LFSR18 BX
$SNDR_I$ (dB)	9.5	14	7.5	15
$SNDR_Q$ (dB)	10.8	12.8	9.3	16

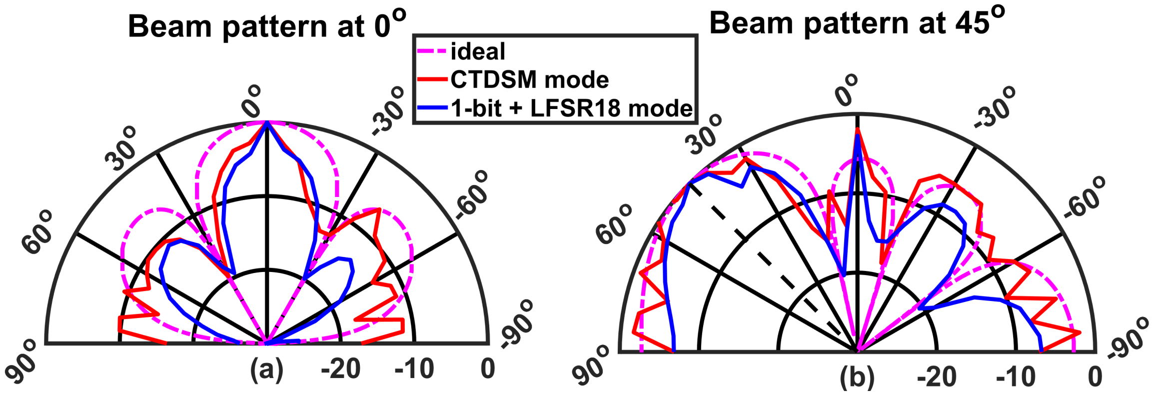

Figure 4.6: (a) Beam pattern at 0° (b) Beam pattern at 45°

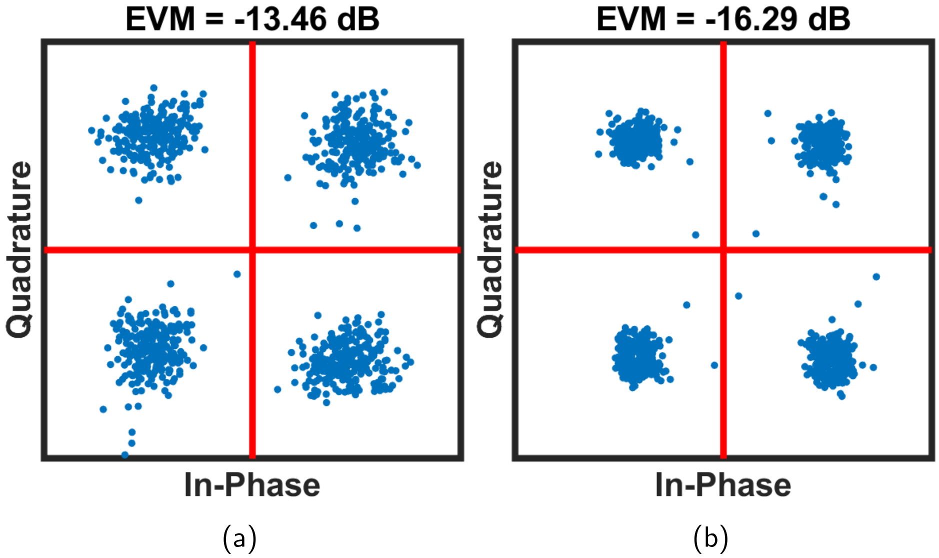

Figure 4.7: (a) QPSK constellation using 1-bit + LFSR18 (b) QPSK constellation using CTDSM

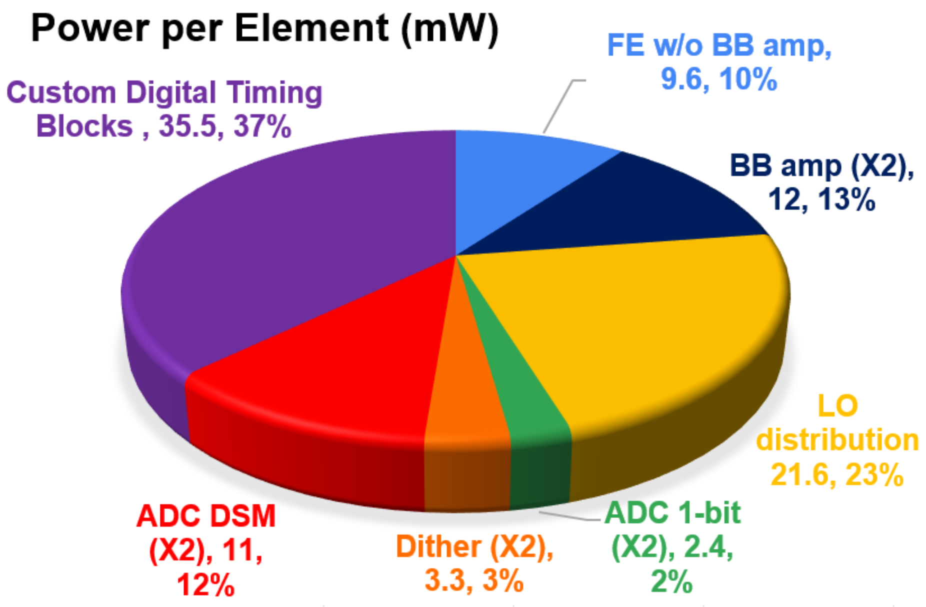

Figure 4.8: RX Phased-Array Receiver Power Breakdown Per Element

Table 4.3: Comparison Table

		This Work	Lu et al. JSSC'21 [6]	Pellerano et al. ISSCC'19 [7]	Park et al. ISSCC'20 [8]	Naviasky et al. ISSCC'21 [9]
Technology		28nm	40nm	22nm FinFET	28nm	28nm
Frequency (GHz)		47-57	28	71-76	39	71-86
Elements per IC		4	16	4	16	16
Beam Former	Type	Digital	Digital	Analog	Analog	Analog
	Phase Res (bits)	10	10	8	4	5
RX FE	Gain (dB)	27.1	21	36.7		35
	NF (dB)	7.8-12	7	6	4.2-4.6	9-11
	BB BW (MHz)	400	100	2000	800 / 100	2000
ADC	Number of ADCs	8	64	-	-	-
	SNDR	21	32	-	-	-
Data Mod.	Constellation	QPSK	4QAM	16QAM	64QAM	16QAM
	Data Rate (MS/s)	400	5	1000	-	250
	EVM (dB)	-16.2	-18	-24	-34.8 / -37.9	-19.2
Power per Element (mW)		96*	177**	42	39	107
Area (mm^2)		0.18 (1 RX)	0.48 (1 RX)	1.26 (1 RX)	1.875 (1 RX)	16

***MFFE+BB Amps+ADC+Digital Timing** ****Front End+ADC+PLL+BSP**

Chapter 5

CONCLUSION

5.1 Thesis Summary

This work describes the motivation for low-resolution ADCs in digital beamforming systems and presents study of open-loop 1-bit ADCs using behavioral models as well as measurement results of a four-element digital beamforming RX using open- and closed-loop 1-bit quantizers. The impact of dithering to recover the array gain is demonstrated. Demodulation of a 400 MS/s QPSK signal with a signal bandwidth 4x greater than measured in [6] was achieved while reducing per-element power and area consumption.

5.2 Future Directions

The following consists of a few suggestions for investigating the design and scalability of wide bandwidth mm-Wave digital beamforming receivers:

1. The 1-bit ADC of Chapter 2 and 4 utilized non-subtractive dither to reduce output distortion and improve the array gain. However, in the actual silicon implementation, the higher order harmonics were not able to be attenuated below the noise floor due to the limited amplitude of the dithering sequence. Preserving a similar architecture, current tuning with parallel devices to the input devices can be investigated to increase the dither amplitude range while optimizing for the minimum pre-amplifier time constant and power consumption. Additionally, given the dither sequence is pseudo-random, a delayed sequence can be subtracted digitally as discussed in [14] to reduce the noise floor and improve the ADC resolution. Finally, the shape of the probability distribution function of the dither sequence can be altered (e.g. from a rectangular to triangular PDF) to not only un-correlate the quantization error with the input signal but also with the injected dither itself, which may improve the array gain by further de-correlating the dither noise powers.

2. The analysis in Section 2.1 only evaluated the digital beamforming system with a single-tone input at 0 degrees. In practice the system will be receiving multiple desired and undesired signals from different angles. Thus, further system-level simulations will need to be designed to determine the bit error rate (BER) with varying modulation order (as tested in [1]) and multiple input incident angles. The minimum ADC resolution with an N-element system for all these cases is yet to be determined. Additionally, given each RX chain in a digital beamforming system is (ideally) identical and sums the phase-shifted signals in the digital domain, there is no spatial rejection that is present in analog beamforming systems, so the specifications of the analog and ADC blocks can not be relaxed. Investigation into front-end circuits that perform spatial rejection can reduce the minimum ADC resolution required.

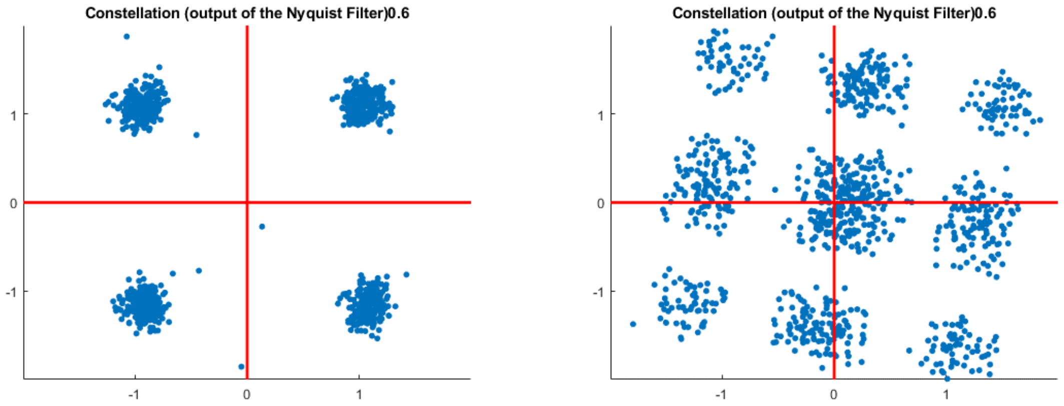

(a)

(b)

Figure 5.1: (a) 400 MS/s QPSK constellation using CTDSM with a 0.6 clock period delay correction (b) Second collected sample using settings from (a)

3. The LO and CLK for this chip were supplied off-chip for the 4-element RX. While this simplified the design of the RX, the lack of clock synchronization between the signal sources, the chip analog front end and digital back end impeded the demodulation of the input QPSK signal on-chip. All utilized signal sources were synchronized to the

same 10-MHz reference clock, but the same beamforming weights and delay settings were unable to produce consistent results as both the LO and CLK sampled at different phase offsets between each collected samples. As a visual example Fig. 5.1a shows an acquired beamformed and demodulated QPSK signal post-beamforming. The delay of the Nyquist filter used prior to demodulation is adjusted to $0.6 \cdot T_{\mathrm{CLK}}$, where $\mathrm{T_{CLK}}$ is the sampling clock period. A second sample is collected using identical settings (Fig. 5.1b) and the output constellation appears rotated (due to LO phase mismatch) and altered (due to CLK phase mismatch). While a tunable delay line was implemented on-chip to adjust the timing between the ADC clock and digital input sequential logic, improper routing rendered it unable to adjust the delay. Thus, more robust clock synchronization methods between the utilized test signal sources and different analog and digital domains on-chip must be studied and implemented for future mm-Wave beamforming receivers. Furthermore, an additional fixed phase delay was observed between each RX element that needed to be manually compensated via the programmable delay line to ensure the spacing between each wavelength was fixed at half the carrier wavelength. This could have been due to differences in the LO distribution between each element, but further investigation is required.

4. The higher-resolution ADC used in this work is a 3rd-order CTDSM. However, this architecture comes with a number of constraints. First, the pole-zero locations of the loop filter are dependent on the clock frequency, which limits the clock tuning to a finite range to preserve the shaped noise transfer function. Second, the mismatch of the resistors and capacitors used in the loop filter can result in different optimal tuning settings of the multiple delta-sigma converters used in the digital beamforming system, as shown in Fig. 5.2. The optimum settings for the resistors and capacitors of the third integrator differ for the ADCs of elements 1 and 4. Additionally, the search space is non-convex, which increases the complexity of a potential search algorithm that can be used. Alternative high-speed ADC architectures should be studied to determine the SNDR frequency-dependence and rapid tuning feasibility.

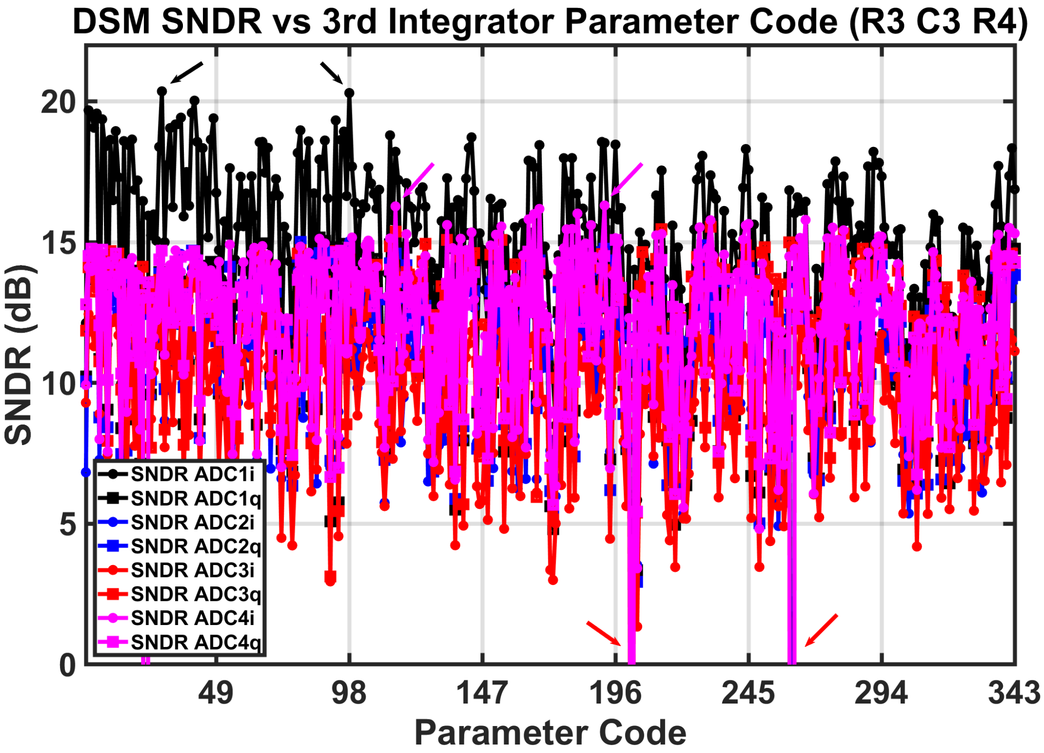

Figure 5.2: DSM Phased-Array Tuning Example. The black and pink errors point to resistor and capacitor settings of the third integrator that yield the highest SNDR for the ADCs in elements 1 and 4, respectively. The red arrows point to corrupted data sets, most likely due to scan chain failure.

BIBLIOGRAPHY

[1] K.D. Chu. Area and Power Reduction Techniques for Millimeter-Wave Phased-Array Transceiver Front-ends, PhD Dissertation, University of Washington, 2021.

[2] A. Ramkaj. In the Pursuit of the Optimal Accuracy–Speed–Power Analog-to-Digital Converter Architecture. *IEEE Solid-State Circuits Magazine*, 14(1):45–53, 2022.

[3] B. Murmann. ADC Performance Survey 1997-2021, 2021. Accessed: May 2022. [Online]. Available: `https://web.stanford.edu/~murmann/adcsurvey.html`.

[4] Y. Duan and E. Alon. A 12.8 GS/s Time-Interleaved ADC With 25 GHz Effective Resolution Bandwidth and 4.6 ENOB. *IEEE Journal of Solid-State Circuits*, 49(8):1725–1738, 2014.

[5] ETSI. "ETSI TS 138 101-2 v17.5.0". Technical report, 3GPP, 2022.

[6] R. Lu, C. Weston, D. Weyer, F. Buhler, D. Lambalot, and M. P. Flynn. A 16-Element Fully Integrated 28-GHz Digital RX Beamforming Receiver. *IEEE Journal of Solid-State Circuits*, 56(5):1374–1386, 2021.

[7] S. Pellerano et al. A Scalable 71-to-76GHz 64-Element Phased-Array Transceiver Module with 2×2 Direct-Conversion IC in 22nm FinFET CMOS Technology. In *IEEE ISSCC*, pages 176–178, February 2019.

[8] H. C. Park et al. A 39GHz-Band CMOS 16-Channel Phased-Array Transceiver IC with a Companion Dual-Stream IF Transceiver IC for 5G NR Base-Station Applications. In *IEEE ISSCC*, pages 76–78, February 2020.

[9] E. Naviasky, L. Iotti, G. LaCaille, B. Nikolić, E. Alon, and A. M. Niknejad. A 71-to-86-GHz 16-Element by 16-Beam Multi-User Beamforming Integrated Receiver Sub-Array for Massive MIMO. *IEEE Journal of Solid-State Circuits*, 56(12):3811–3826, 2021.

[10] V. Bhagavatula, T. Zhang, A. R. Suvarna, and J. C. Rudell. An Ultra-Wideband IF Millimeter-Wave Receiver With a 20 GHz Channel Bandwidth Using Gain-Equalized Transformers. *IEEE Journal of Solid-State Circuits*, 51(2):323–331, 2016.

[11] K.D. Chu et al. A Mm-Wave Gm-Assisted Transformer-Based Matching Network 2x2 Phased-Array Receiver for 5G Communication and Radar Systems. In *2021 IEEE Radio and Wireless Symposium (RWS)*, pages 26–28, 2021.

[12] H.T. Friis. A Note on a Simple Transmission Formula. *Proceedings of the IRE*, 34(5):254–256, 1946.

[13] B.H. Ku, P. Schmalenberg, O. Inac, O. D. Gurbuz, J. S. Lee, K. Shiozaki, and G. M. Rebeiz. A 77–81-GHz 16-Element Phased-Array Receiver With ±50° Beam Scanning for Advanced Automotive Radars. *IEEE Transactions on Microwave Theory and Techniques*, 62(11):2823–2832, 2014.

[14] R. A. Wannamaker. Subtractive and Nonsubtractive Dithering: A Mathematical Comparison. *Journal of the Audio Engineering Society*, 52, 2004.

[15] S. Golomb. *Shift-Register Sequences, Revised Edition.* Aegean Park Press, 1981.

[16] W. R. Bennet. Spectra of Quantized Signals. *Bell System Technical Journal*, 27:446–471, 1948.

[17] R. A. Shafik et al. On the Extended Relationships Among EVM, BER and SNR as Performance Metrics. In *ICECE*, pages 408–411, December 2006.

[18] D.V. Ess. Signals From Noise: Calculating Delta-Sigma SNRs.

[19] I. Zhao. Custom High-Speed ADC for mmWave Digital Beamformers, Masters Thesis, University of Washington, 2021.

[20] D. Dosluoglu et al. A Reconfigurable Digital Beamforming V-Band Phased-Array Receiver. In *ESSCIRC 2022 - IEEE 48th European Solid State Circuits Conference (ESSCIRC)*, to appear in 2022.

[21] D. Pena-Colaiocco et al. An Optimal Digital Beamformer for mm-Wave Phased Arrays with 660MHz Instantaneous Bandwidth in 28nm CMOS. In *IEEE ISSCC*, pages 1–3, February 2022.

[22] P. Song and H. Hashemi. mm-Wave Mixer-First Receiver With Selective Passive Wideband Low-Pass Filtering. *IEEE Journal of Solid-State Circuits*, 56(5):1454–1463, 2021.

[23] D. Murphy, H. Darabi, A. Abidi, A. A. Hafez, A. Mirzaei, M. Mikhemar, and M. Frank Chang. A Blocker-Tolerant, Noise-Cancelling Receiver Suitable for Wideband Wireless Applications. *IEEE Journal of Solid-State Circuits*, 47(12):2943–2963, 2012.

[24] J. Mo and R.W. Heath. Capacity Analysis of One-Bit Quantized MIMO Systems With Transmitter Channel State Information. *IEEE Transactions On Signal Processing*, 63(20):5496–5512, 2015.

[25] O. Mattia and B. Murmann. An 80 GS/s 5.5 ENOB Time-Interleaved Inverter-Based CMOS Track-and-Hold. *Electronics Letters*, 56(7):328–329, 2020.

[26] A.T. Ramkaj, M.S.J. Steyaert, and F. Tavernier. A 13.5-Gb/s 5-mV-Sensitivity 26.8-ps-CLK–OUT Delay Triple-Latch Feedforward Dynamic Comparator in 28-nm CMOS. *IEEE Solid-State Circuits Letters*, 2(9):167–170, 2019.

[27] M. van Elzakker, A. J. M. van Tuijl, P.F.J. Geraedts, D. Schinkel, E. A.M. Klumperink, and B. Nauta. A 10-bit Charge-Redistribution ADC Consuming 1.9µW at 1MS/s. *IEEE Journal of Solid-State Circuits*, 45(5):1007–1015, 2010.

[28] B. Razavi. The Design of a Comparator [The Analog Mind]. *IEEE Solid-State Circuits Magazine*, 12(4):8–14, 2020.

[29] D. Schinkel, E. Mensink, E. Klumperink, E. van Tuijl, and B. Nauta. A Double-Tail Latch-Type Voltage Sense Amplifier with 18ps Setup+Hold Time. In *2007 IEEE International Solid-State Circuits Conference. Digest of Technical Papers*, pages 314–605, 2007.