Title
Configurable Data Converters for Digitally Adaptive Radio

Author
Whitcombe, Amy

Publication Date
2018

Peer reviewed|Thesis/dissertation

Configurable Data Converters for Digitally Adaptive Radio

by

Amy Whitcombe

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Borivoje Nikolić, Chair
Professor Elad Alon
Professor Aaron Parsons

Fall 2018

Configurable Data Converters for Digitally Adaptive Radio

by

Amy Whitcombe

Abstract

Configurable Data Converters for Digitally Adaptive Radio

by

Amy Whitcombe
Doctor of Philosophy in Engineering - Electrical Engineering and Computer Sciences

University of California, Berkeley
Professor Borivoje Nikolić, Chair

In the past few decades, wireless connectivity has grown from an expensive luxury to an integral component of everyday life, creating a need for wireless systems that can satisfy exponentially growing demand for data in an energy- and cost-efficient manner. New techniques such as digital beamforming can utilize spatial diversity to support multiple users with arrays of radio elements, but the required performance of these elements varies with array size. Additionally, while conventional single-element radios are typically designed to meet stringent performance requirements outlined in wireless standards, peak performance is rarely required. Adaptive wireless receivers that configure performance to suit a system's needs may therefore serve as building blocks for energy-efficient wireless platforms, and enhanced digital processing capabilities afforded by CMOS technology scaling can help realize fully integrated smart wireless systems. However, while computation improvements due to process scaling have driven the adoption of smart devices, complex design rules and lower supply voltages often make it difficult to construct high-performance analog circuits such as wireless receivers in advanced process nodes. To that end, this dissertation discusses the design and implementation of analog-to-digital interface circuits for receivers that are both configurable and well-suited to implementation in scaled nodes optimized for digital performance. It first explores circuit-level techniques for constructing resolution-configurable receivers from scalable elements, and then discusses an alternative receiver design that can integrate scalability at the architecture level.

To explore how circuit-level design techniques can enable configurability, this dissertation first discusses the design of a resolution-scalable successive approximation (SAR) analog-to-digital converter (ADC) for wireless receiver applications. The converter is built with a scalable capacitive digital-to-analog converter (DAC), comparator, and tunable switching algorithm to trade power for resolution. The 80 MS/s prototype converter implemented in a general purpose 65nm CMOS process consumes 0.4-0.8 mW and provides 7.0-9.1 effective bits of resolution in a 10 MHz signal bandwidth. It is integrated with a power-scalable receiver to demonstrate its suitability for wireless systems. By trading noise for power consumption, this receiver can serve as a building block for energy efficient digital beamforming arrays.

Finally, this thesis demonstrates architectural techniques for affording configurability by discussing the design and implementation of a reconfigurable, digital-intensive RF-to-digital converter. The proposed receiver uses a SAR ADC with integrated discrete-time

filtering to provide high linearity and a voltage-controlled oscillator (VCO) based ADC to improve sensitivity. By replacing the high-performance active amplifiers and filters used in conventional receivers with an ADC constructed from digital building blocks, the 16nm CMOS FinFET prototype can leverage the benefits of technology scaling. In low-power mode, the 0.7-1.9 GHz receiver is configured as a VCO-based design drawing 9-16 mW and providing -82 dBm sensitivity in a 10 MHz bandwidth. In blocker-tolerant mode, the SAR ADC can be enabled to obtain 60 dB SNDR and an in-band IIP3 of +16 dBm. The prototype achieves performance comparable to state-of-the-art RF-to-digital converters using an easily configurable digital-intensive design. Overall, both configurability techniques discussed in this dissertation are promising means of leveraging the advantages of CMOS scaling to enable future digitally adaptive wireless systems.

Contents

List of Figures

Acknowledgments

I'm incredibly grateful to the many people who have shaped this dissertation, both through direct technical contributions and by supporting me in my time at Berkeley. First of all, I'd like to thank my advisor Bora Nikolić for giving me the opportunity to come to Berkeley and encouraging me to try a broad range of projects during my PhD. I've learned a lot in the past five years, and really appreciated Bora's support of my ideas, his eager feedback on my work, and the way he encouraged me to work with many people both at Berkeley and in industry. Having the opportunity to collaborate with many people is one of the reasons I came to BWRC, and I definitely had the chance to do that while I was here. In addition to Bora's advising, I'd like to thank Elad Alon for his helpful feedback on this dissertation and, more importantly, for his dedicated efforts leading many of the projects I've worked on. Elad has an incredible way of providing insightful feedback on seemingly any subject, and his hands-on involvement with the CRAFT tapeout made it really feel like a group effort. I'm also very thankful to my outside committee member, Aaron Parsons, for showing interest in my work and taking the time to sit through my qualifying exam and to read this dissertation.

I feel lucky to have been a part of some great teams at Berkeley; none of the chips in this thesis could have been taped out and tested without significant effort from a lot of people. On the eWallpaper/FADER project, I'm especially glad I had the chance to work with Kosta Trotskovsky, whose attention to detail has always been something to aspire to. His patience, good humor, and methodical way of thinking makes him a great person to collaborate and learn with. On top of this, the radio test chip never could have happened without Greg LaCaille and Antonio Puglielli's perseverance and willingness to organize the top-level chip integration. Putting together the untested designs of six fairly new grad students into a testable, working chip is difficult to say the least, and I'm certainly grateful to them for stepping up to the challenge. When it came time to test the chip, Pengpeng Lu and Kosta's work setting up the PCB and test infrastructure made it possible to get the measurement results presented in this dissertation. I'm thankful to all of the students involved in the large eWallpaper project — including Nathan Narevsky, Eric Chang, and Zhongkai Wang — for their friendship and technical advice. In addition to Bora and Elad's help advising, Ali Niknejad provided insightful feedback throughout the project. And no eWallpaper acknowledgements section would be complete without thanking Greg Wright of Nokia, whose extensive knowledge and enthusiasm for the project (and many other projects) has always been appreciated.

The second chip discussed in this dissertation was made possible by the dedicated team of people involved in the CRAFT project, and I'm still fairly amazed that we managed to get the (working!) chip taped out in such a short timeframe. First of all, the tapeout could not have happened without Stevo Bailey's thankless work putting together the top-level digital flow so that a bunch of analog designers could work until the last minute to finish their designs. There wouldn't be a chip to test if he hadn't stepped up to help. Beyond that, Angie Wang's help figuring out top-level DRC errors and Jaeduk Han and Eric Chang's efforts figuring out the PDK details from the analog side made it much easier to put this chip together. I'm also very grateful to Angie, Eric, and Jaeduk for having gone through the pains of taping out the first chip in this process. Contributions of the aforementioned people

also made it possible for me to build a receiver from scratch in three months. Using blocks from the BAG framework, including Nathan's clock receiver, Eric's sampling switches, and Jaeduk's strong-arm latch, helped reduce design time. I especially appreciate Angie's work developing a memory for my ADC, which gave me the chance to focus on other aspects of the design. When it came time to test the chip, I'm really grateful to Stevo and James Dunn for working with me to put together the PCB. Finally, Brian Richards helped with everything from top-level integration to dealing with MOSIS and the CRAFT program, and of course Elad and Bora really helped drive the project.

Beyond Berkeley, I've had great opportunities to collaborate with industry. My first chip was a partnership with Raytheon Vision Systems, and I'm especially thankful to Scott Taylor, Bruce Bozovich, Dave Madajian, Eileen Herrin, and Marty Denham for helping with the implementation and testing of that design. Everyone at Raytheon was supportive of the project and taught me a lot when I was a first year graduate student still figuring out exactly what I was getting myself into by studying IC design. My third main chip, and the second major part of this dissertation, emerged from an internship at Intel. It never would have grown into such a large part of my research without Farhana Sheikh's persistence and encouragement. I'm truly thankful for the time and effort Farhana has invested in me, both in mentoring me throughout my time at Intel and in guiding my project and encouraging me to take it further. During my internships, technical discussions with Erkan Alpman, Jay Wang, and Shreyas Sen helped give me direction; Erkan's background in ADC design was especially useful in defining the project. In my second internship, Ashoke Ravi's advice and guidance was very helpful, particularly as I was still fairly new to the wireless field. I'm also thankful to Chris Hull and Brent Carlton for their useful feedback and eagerness to listen during technical reviews both at Intel and back at Berkeley.

Many people I worked with directly in research provided friendship and helped give me reasons to laugh even through long nights and early mornings at BWRC. Outside of research, I've really appreciated meeting and learning from a diverse range of people in grad school. Having Kristel Deems and Seobin Jung around made transitioning to Berkeley a lot easier. I never expected you guys to actually write formal LaTeX reports about the random YouTube videos I sent when I asked you to, but I'm glad you're the sort of people who would. Mira Videnovic-Misic (and Una), thanks for the all of the snacks and for being so welcoming from the moment you got here. A number of older grad students helped welcome me when I first came to Berkeley, including Katerina Papadopoulou, Amanda Pratt, Brian Zimmer, Matt Weiner, Sharon Xiao, and Matt Spencer. And throughout my time at Berkeley, I've appreciated the helpfulness, ridiculousness, and sense of humor (or at least a subset of these traits) of many people in Bora's group and beyond, including Pi-Feng Chiu (despite all the teasing, I do appreciate you), Ben Keller, Rachel Hochman, Sameet Ramakrishnan, Paul Rigge, John Wright, James Dunn, Luke Calderin, Keertana Settaluri, Krishna Settaluri, Emily Naviasky, and Sidney Buchbinder.

I've mentioned some of them already, but I'd like to reiterate my thanks to the staff at BWRC for all of their help over the past few years, both in research and in making the BWRC a welcoming environment. Everyone in the front office over the past few years — Candy Corpus, Yessica Bravo, Melissa Trevizo, Olivia Nolan, and Sarah Jordan — has somehow managed to keep the BWRC running while being incredibly friendly and putting up with the craziness of a bunch of grad students. Candy's enthusiasm and eagerness to

help especially makes the BWRC a better place. Leslie Nishiyama and Amber Sanchez's help getting purchases through quickly made it possible to get research done. James Dunn has been tremendously helpful over the years with everything from PCB design to software tool support, and somehow manages to maintain a positive attitude and good sense of humor no matter how much he has on his plate. Fred Burghardt's wonderful organization of the BWRC lab made it a great environment to work in, and his prompt help and solid advice made it possible to test my chips throughout the years. Brian Richards has helped me throughout grad school, from working with me to figure out how to set up DRC and LVS in a new technology for my first tapeout to helping take die photos of my last chip, and I'm very grateful for all of his work over the past five years. And in the EECS department, I'd like to thank Shirley Salanio and Susanne Kauer for all of their friendly help in my time at Berkeley.

At a higher level, I wouldn't be writing this dissertation without the people who got me interested in electrical engineering in the first place. Mrs. Stadler's Science Olympiad team helped draw me to engineering, and Mr. Stadler's time coaching me and my awesome partner Kim in Circuit Lab helped me enter college with an interest in circuits. In undergrad, Brad Minch and Kent Lundberg really introduced me to IC design, and Siddhartan Govindasamy introduced me to wireless systems. A combination of their encouragement, interesting classes, and Sherra Kerns's passion for research helped steer me towards graduate school. All of the crazy people at Olin made engineering fun even when we were pulling all-nighters to meet project deadlines. For better or worse, this was good preparation for tapeouts. And finally, thanks to my parents for always supporting and encouraging me while working hard to make sure I had the opportunity to explore what I wanted.

Chapter 1

Introduction

In the past few decades, wireless connectivity has evolved from an expensive luxury to an integral component of everyday life. While the first handheld cellular devices sold commercially in the late 1980s cost the modern-day equivalent of over $9000, by 2017 roughly 95% of US adults were estimated to own a cell phone, and 77% a smartphone [3]. Adoption of wireless devices is only expected to grow with time as connectivity becomes standard among younger generations — the estimated average age of a child receiving a new cell phone is 10 years old [4]. Aligning with the growing adoption of wireless devices, demand for wireless data has grown exponentially. Mobile data traffic has grown 47% annually since 2016, and is projected to continue growing at this rate in the near future [5]. Due to these trends, successive generations of wireless systems are expected to support faster data rates and improve processing capability within a modest power budget at little to no increased cost, posing a significant engineering challenge.

1.1 Research scope

Sustaining continuous improvements in data rates and processing capabilities requires innovations in the design of wireless systems. Specifically, this work focuses on radio receiver design, with a particular emphasis on the analog-to-digital converters (ADCs) that bridge wireless receivers and digital baseband processors. A first major challenge in modern receiver design is compatibiliy with CMOS process technology scaling. Historically, CMOS scaling has played a major role in facilitating the growing capabilities of wireless systems; however, shrinking supply voltages, higher levels of process variation, and increasing design complexity in advanced nodes has made it increasingly challenging to efficiently design highly sensitive analog circuits such as receivers in these nodes. Second, wireless systems must be configurable to enhance energy efficiency. While receivers are often designed for worst-case performance conditions outlined in wireless standards, actual capacity requirements are heavily dependent on a user's environment. As a result, this work explores approaches to developing receivers that are both amenable to implementation in scaled CMOS and reconfigurable to afford power savings.

1.1.1 CMOS scaling and analog design

One of the major innovations that has facilitated the proliferation of handheld wireless platforms in the past decade is semiconductor technology scaling in accordance with Moore's

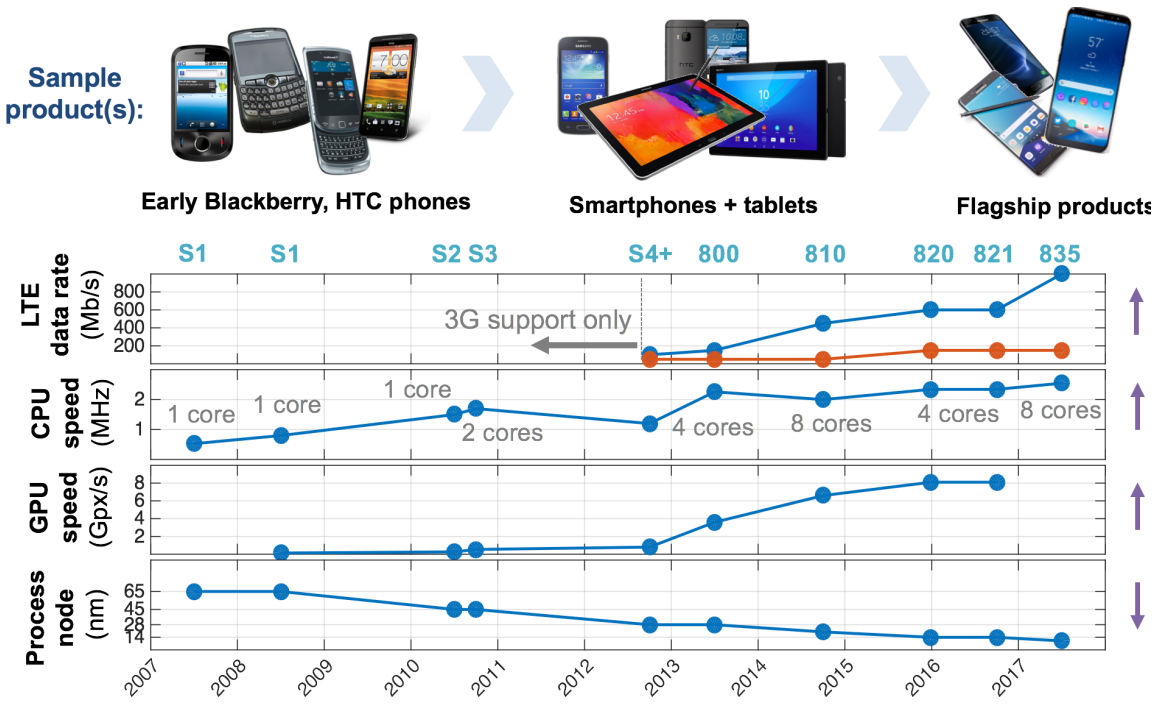

Figure 1.1. Capabilities of Qualcomm Snapdragon SoCs over time. Data from [7].

law. While Moore's law, which predicts that advances in fabrication technologies will allow the number of transistors on an integrated circuit (IC) to double every 2-3 years [6], emerged as a short-term prediction, this trend has been sustained over half a century since it was first proposed. As an example to demonstrate the effects of CMOS scaling on wireless systems, Fig. 1.1 shows how the computing capabilties of the Qualcomm Snapdragon system-on-a-chip (SoC) — contained in many modern wireless devices — have steadily expanded in the past decade as process technology has scaled. Each new version of the processor has come to support steadily growing data rates and improved computational and graphics processing capabilities at an affordable cost, all in accordance with Moore's law scaling. Correspondingly, this SoC has enabled increasingly complex products, from early smartphone models with moderate display capabilities before 2010, to a broad range of smartphones and tablets before 2015, to flagship products such as the Samsung Galaxy 8 in the past few years.

While technology scaling enables significant advances in digital processing power, this processing power must interface with an analog world. To truly benefit from CMOS scaling, integrated SoCs must use analog circuits that are well-suited to implementation in scaled process nodes. Scaling affords some natural advantages for analog circuits — for instance, scaled devices are typically faster, have larger current drive for a fixed leakage current, and can take advantage of the large increase in digital computational power made available to calibrate and tune analog circuits. However, technology scaling comes with many challenges for high-performance analog design. For instance, supply voltages have traditionally scaled faster than threshold voltages as shown in Fig. 1.2(a), leading to reduced voltage headroom that makes it difficult to design high-gain, low-noise amplifiers. Additionally, performance variability for minimum sized transistors in a given process node grows as feature sizes shrink,

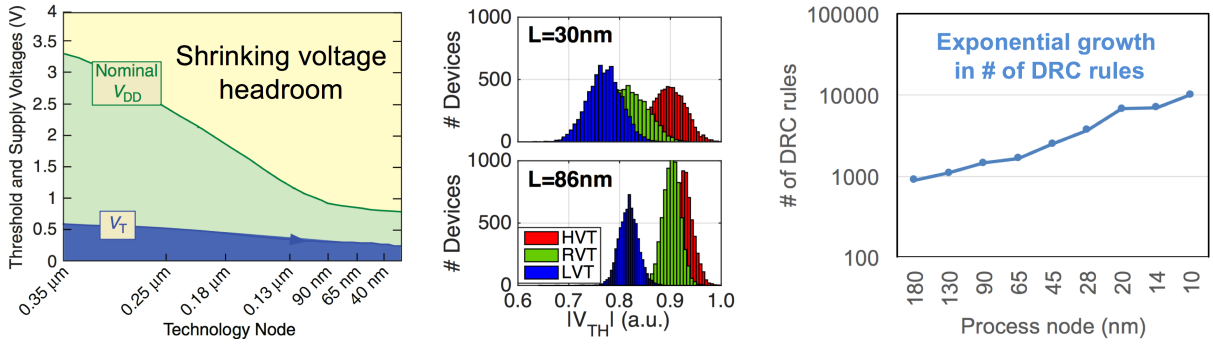

(a) Supply (V_{DD}) and threshold voltage (V_T) vs. process node. From [9].

(b) Variability for two device lengths. From [8].

(c) Estimated # of design rules vs. process node. Data from [10]

Figure 1.2. CMOS technology scaling challenges.

and non-ideal effects such as random telegraph noise become increasingly significant design challenges [8] (Fig. 1.2(b)). Finally, as device sizes have shrunk exponentially, the number of design rule checks (DRCs) has grown exponentially, making it difficult to quickly prototype designs in modern process technologies (Fig. 1.2(c)).

The analog-to-digital converters (ADCs) that enable SoCs to use the increased digital computation power afforded by technology scaling should therefore utilize architectures that can benefit from scaling as well. To that end, the first major focus of this work is to explore the design of process-scalable converter architectures, and receiver topologies that can benefit from CMOS scaling. This means a shift away from architectures that require high-gain static amplifiers to architectures that require predominantly switches, capacitors, and digital blocks.

1.1.2 Motivation for reconfigurable receivers

In addition to being compatible with scaled CMOS process nodes, wireless receivers should be performance-scalable to maintain low energy consumption. Despite the fact that wireless standards outline worst-case operating conditions, actual performance requirements vary greatly for different users. For instance, the 20 MHz Wi-Fi specification requires a receiver to detect signal powers between -82 to -74 dBm in the presence of adjacent channel blockers that are 20-40 dB more powerful, with a maximum signal strength of -30 dBm. The 5 MHz LTE specification is more stringent, requiring a -94 dBm signal to be detectable in the presence of -30 dBm adjacent channel blockers, and out-of-band blockers as high as -15 dBm. To support such a wide dynamic range, either a high-order filter must be present in the receiver or the ADC must provide almost 13 effective bits of resolution.

Despite these requirements, actual operating conditions can vary. Figure 1.3 shows the breakdown of Wi-Fi and cellular signal strength levels measured from a randomly sampled group of users over more than a one month period [11]. Some users see predominantly strong signals, while some see predominantly weak signals, and many see a broad range of signal strengths. While Wi-Fi standards require receivers to detect signals down to -74 dBm, most users observe these signal levels less than 25% of the time. Similarly, the -94 dBm sensitivity level outlined in LTE standards is detected by most users less than 25% of the time. To

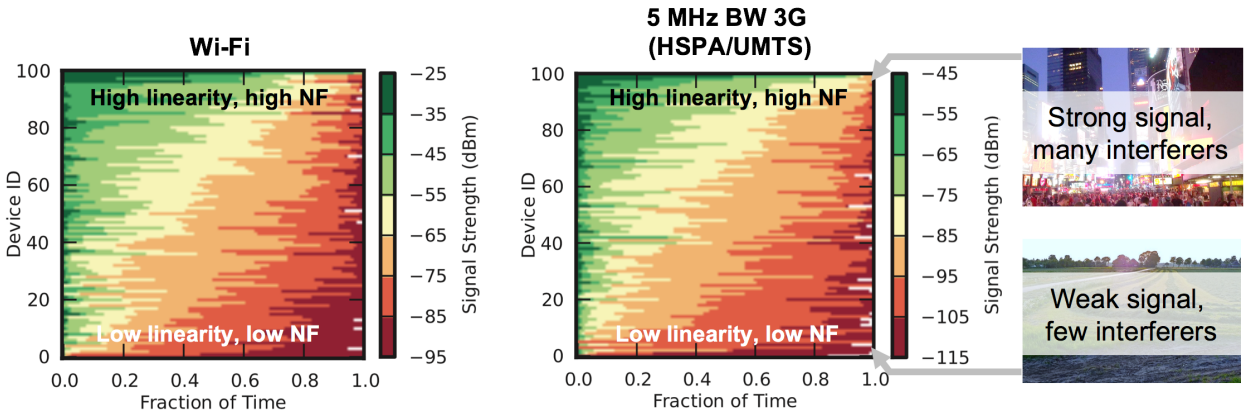

Figure 1.3. Measured wireless signal strength in a variety of handheld devices, from [11].

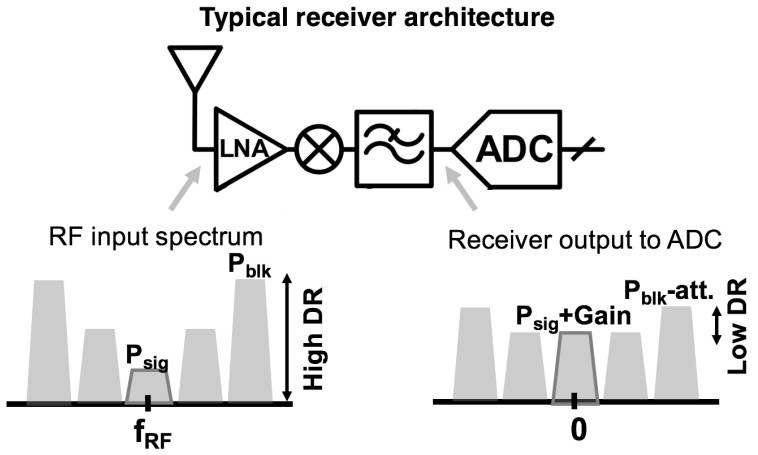

Figure 1.4. Generic receiver topology and spectrum before and after receiver.

truly operate efficiently, radios should not be designed to meet the stringent requirements outlined in the standard, but adapt to the actual needs of each user.

Power scalability can also make new wireless technologies feasible. For instance, massive multiple-input multiple-output (MIMO) beamforming applications require large arrays of radios to support spatial selectivity. By leveraging the noise averaging properties of the array to relax the noise figure requirement of each receiver element, power-scalable radios can allow these arrays to be built in an energy-efficient manner. Scalability therefore makes it possible to improve throughput via spatial directivity with an acceptable power overhead.

1.1.3 Research focus

The goal of this work is to explore analog-to-digital conversion techniques suited to radios that are both easy to integrate with complex digital systems in advanced CMOS process nodes and configurable for improved energy efficiency. Typically, the ADC appears at the end of a receiver chain, following low-noise amplification, down-conversion of the RF signal, and filtering that prevents aliasing and relaxes the dynamic range requirements on the ADC

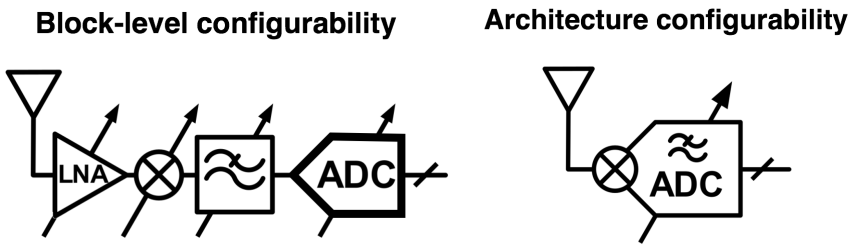

Figure 1.5. Comparison between approaches to designing configurable receivers.

(Fig. 1.4). Considering this, design configurability can be implemented in two ways: by constructing a conventional receiver from independently tunable elements (as shown in Fig. 1.5(a)), or by reworking the receiver architecture to incorporate configurability (as shown in Fig. 1.5(b)).

To investigate the first technique, a resolution-configurable successive approximation (SAR) ADC was designed and implemented in 65nm CMOS. The SAR architecture is inherently well-suited to implementation in scaled CMOS, and the low-power ADC can help enable digital beamforming in radio arrays. Studying this architecture provides insight into the benefits and limitations of designing a power-scalable, variable-resolution ADC for these applications. To explore the second approach to scalable receiver design, an architecture-configurable ADC suitable for direct RF-to-digital operation was designed and implemented in 16nm FinFET. The proposed hybrid SAR and voltage-controlled oscillator (VCO) based receiver presents a means of building software-defined radio that is highly suitable for modern process technologies optimized for digital performance.

1.2 Related work

The contributions of this dissertation build on a foundation of existing work in both reconfigurable and process-scalable receiver design. Efforts to design configurable wireless receivers have been ongoing for decades, especially as a multitude of wireless standards have emerged that require receivers to support a broad range of bandwidths and carrier frequencies. Additionally, desire for receivers compatible with deeply scaled CMOS process nodes has given rise to a range of alternative receiver design approaches, from digitally-assisted performance enhancement to switched-capacitor architectures that benefit from improved switching speeds. This section provides a high-level overview of research in these two areas. More detailed summaries of relevant research are presented in later chapters.

1.2.1 Configurable receivers

Receivers must be capable of detecting a wide range of bandwidths and carrier frequencies to support low-bandwidth legacy wireless standards and new high-bandwidth standards throughout the globe. Frequency and bandwidth configurability is therefore an attractive low-cost alternative to designing customized receivers for each standard. Traditionally, carrier frequency tunability can be difficult to achieve because low-noise amplifiers (LNAs) are often optimized to provide a frequency-dependent impedance match to the antenna. As an

alternative, wideband operation can be achieved using mixer-first receiver architectures or custom LNA designs. In a mixer-first receiver, the RF signal is immediately down-converted to baseband, enabling impedance matching to be performed after the mixer using an amplifier with a lower gain-bandwidth product [12–14]. This approach can be coupled with passive discrete-time filtering techniques to provide both wideband operation and digitally programmable filtering [15]. Alternatively, custom LNA solutions can include multiplexing between different LNAs optimized for different frequency bands [16,17] or simplified low-gain LNA topologies [18].

While carrier frequency can be tuned using mixer-first topologies or alternative LNAs, bandwidth tunability can be implemented with custom filter designs. In the simplest case, the bandwidth of an active RC filter can be tuned by making passive component values digitally programmable [13, 16, 19–21]. Similarly, the sizing and architecture of transconductor-capacitor (G_m-C) filters can be tuned using CMOS enable switches [18, 22, 23]. In a more scaling-friendly approach to configurability, the filter can be constructed using a discrete-time architecture requiring only capacitors and CMOS switches [15, 17, 24–28]. In this case, the filter profile can be configured by modifying capacitor sizes.

Though most configurable receiver research has emphasized bandwidth and tuning frequency configurability, additional effort has been made to implement power-scalable receivers. This is especially necessary in wake-up receiver applications designed to minimize power in wireless sensor nodes. Wake-up receivers that trade sensitivity for power consumption have been proposed using programmable amplifier size [19, 29] or clock frequency [30]. Alternative efforts have studied ways to construct receiver elements that trade power for performance, including LNAs with programmable noise figure and linearity [31] and filters with programmable gain-bandwidth [32]. Comparatively, limited prior research has studied fully power-scalable receivers with integrated ADCs, as presented in this work [1].

1.2.2 Process scalable receivers

In an effort to apply the enhanced digital processing capabilities afforded by CMOS scaling to the challenge of supporting diverse wireless standards, [33] initially proposed a fully software-defined radio, in which an ADC directly digitizes the RF input spectrum, and a digital processor performs channel selection and filtering. While the ADC resolution and sample rate requirements render this vision largely impractical [34], partially software-defined radio can be constructed from digitally-tunable components with performance optimized using a digital processor. Blending digitally-configurable receivers with on-chip processors can enable "cognitive radio" — receivers capable of adapting performance to a new environment (e.g., detecting and utilizing available spectrum). Prior work has demonstrated that receivers with tunable bandwidth, carrier frequency, and gain can be integrated with digital feedback loops to tune performance [16, 18, 21, 35].

At the block level, digital assist techniques have been applied to enhance the performance of individual receiver elements. For instance, digital feedback loops can be used to optimize passive component values and bias voltages for improved LNA linearity and bandwidth [36, 37] or filter linearity [38]. In more complex systems, digital optimization algorithms can be applied to maximize noise cancellation [13]. Overall, digital processing is a powerful tool for improving receiver performance.

1.2.3 Contribution of this work

Relative to previous studies of configurable receivers emphasizing frequency tunability, this work focuses on power scaling to enhance energy efficiency in wideband receivers. These power-scalable receivers are suitable for emerging wireless applications such as low-power beamforming arrays that provide a new approach to improving wireless systems capacity. Existing implementations of massive MIMO systems typically use high-performance, power-hungry radio elements that make large array sizes impractical and require the use of analog signal recombination techniques [39]. Relative to digital beamforming enabled by power-scalable ADCs, analog beamforming does not easily support multiple users and does not benefit from CMOS scaling. Additionally, while most tunable receiver designs use either an off-chip ADC or a fixed-performance ADC, this research provides a more comprehensive look at designing power-scalable ADCs specifically for reconfigurable wireless receivers. Finally, this research studies a new CMOS scaling-friendly SAR+VCO receiver architecture implemented in 16nm FinFET to provide an additional example of how to construct efficient configurable receivers.

1.3 Thesis organization

Overall, this dissertation evaluates the design and implementation of two techniques for constructing ADCs for process-scalable, reconfigurable receivers. In the first technique, receiver building blocks are independently designed to trade performance for power. After Chapter 2 presents an overview of major design considerations for ADCs in wireless systems, Chapter 3 describes a resolution configurable successive approximation (SAR) ADC integrated with a power-scalable receiver in 65nm CMOS to study the effectiveness of this block-level configurability technique. The power savings afforded by this ADC can be used to facilitate digital beamforming in large radio arrays of variable sizes. In Chapter 3, Section 3.1 first reviews prior configurable ADC designs. Next, Section 3.2 provides a general review of SAR ADC design techniques, particularly in the context of performance scalability. Details of the 65nm CMOS prototype are then presented in Section 3.3, and measurement results are given in Section 3.4 prior to concluding.

Next, Chapter 4 discusses the design and implementation of a reconfigurable-architecture RF-to-digital converter to demonstrate the feasibility of using custom architectures for power savings. This converter is constructed with a hybrid SAR and VCO based ADC, two architectures that are inherently digital and therefore well-suited to implementation in scaled CMOS. In Chapter 4, Section 4.1 begins by describing examples of software-defined radio in more detail. Then, Section 4.2 reviews design considerations relevant to the proposed architecture and Section 4.3 outlines the prototype implementation in detail. Measurements and comparisons with prior work are presented in Section 4.4 before Section 4.5 concludes the chapter.

Finally, the conclusion of the thesis in Chapter 5 presents a comparison between the two design approaches to evaluate when each proposed technique is best suited to enhance the capabilities of a wireless system. This chapter also summarizes potential extensions of this work for future research.

Chapter 2

Data Conversion for Wireless Applications

As the interface between the receiver front-end and the digital baseband processor, the analog-to-digital converter (ADC) in a wireless receiver must have sufficient dynamic range to detect a broad range of input signal and blocker strengths and a sampling rate fast enough to detect the target signal bandwidth. Because the ADC requirements depend strongly on the target signal spectrum and receiver capabilities, the two components must be designed together.

This chapter provides an overview of major considerations for designing data converters for wireless applications. First, Section 2.1 provides an overview of the wireless spectrum that receivers for various applications must be designed to detect. This includes the range of blocker powers and signal strengths outlined in wireless standards. Next, Section 2.2 discusses how the receiver and ADC must be designed to handle these signals, describing typical specifications for ADCs used in wireless applications. Finally, Section 2.3 presents some alternative receiver topologies that do not consist of a single receiver front-end and ADC, and explores how these different topologies affect ADC requirements.

2.1 Performance requirements of wireless systems

Wireless communications require standardized protocols to transmit information between two different devices. Among other things, these protocols dictate which frequency bands and bandwidths can be utilized, how data will be modulated within this bandwidth, and what power levels equipment must be capable of transmitting or detecting. A high-level frequency-domain overview of these concepts is summarized in Fig. 2.1. Communications occur using a specific frequency channel of interest, centered around a particular carrier frequency, which must be detectable even as adjacent channels and other frequency bands transmit data. Within a channel, different modulation schemes can be used to encode data in the frequency domain. Figure 2.2 illustrates orthogonal frequency division multiplexing (OFDM), a common encoding scheme. In OFDM, evenly-spaced sinc functions (subcarriers) are used with peaks chosen to align with nulls of the other subcarriers; data is encoded in the peak amplitude of each subcarrier. Figure 2.2 also indicates the presence of sidelobes generated within each carrier that may interfere with signals in adjacent channels.

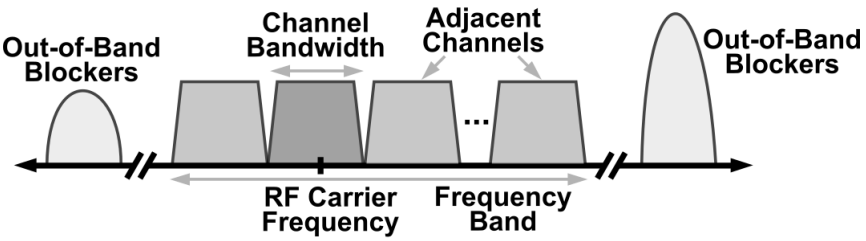

Figure 2.1. Generic wireless spectrum.

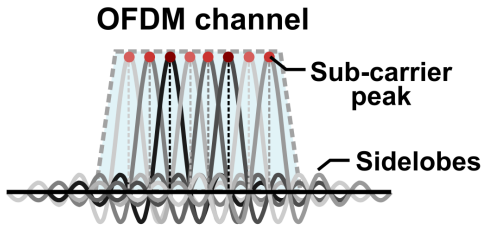

Figure 2.2. Sample channel using orthogonal frequency division multiplexing (OFDM).

While frequency bands for operating wireless equipment are typically allocated by national governments, bandwidths and modulation schemes can be chosen by communication protocols to maximize data throughput. Innovations in encoding schemes and hardware developments supporting higher bandwidth operation have facilitated an exponential growth in mobile data traffic within the past decade [5]. From an implementation perspective, the parameters shown in Fig. 2.1 define the required noise figure, bandwidth, and gain of the receiver, which in turn determines the necessary resolution and sample rate of the ADC. This section will first briefly review different types of wireless devices, and then generally describe LTE and Wi-Fi, two of the most commonly used wireless communications protocols. The next section will discuss the impact of these protocols on receiver and ADC design requirements.

2.1.1 Applications

The number of devices supporting wireless connectivity has grown rapidly in the past few years, particularly with the emergence of interconnected household devices in addition to existing handheld mobile devices and laptop computers. Each application and device type has different system-level requirements. In handheld applications, minimizing power consumption is essential, as any excess power will degrade the battery life of a device. Handheld equipment typically achieves network connectivity using a wireless base station. Figure 2.3 provides a broad overview of classifications of wireless base stations. Small applications with a handful of users (e.g. in-home applications) are suitable for compact femtocells. These stations favor small form factor and low power consumption. Picocells are designed to support tens of users in areas such as commercial buildings or small portions of a city. At the highest level, macrocells are deployed throughout to provide extensive wireless coverage. While power consumption and form factor is not a major concern in macrocells, they must be designed to provide extensive wireless coverage.

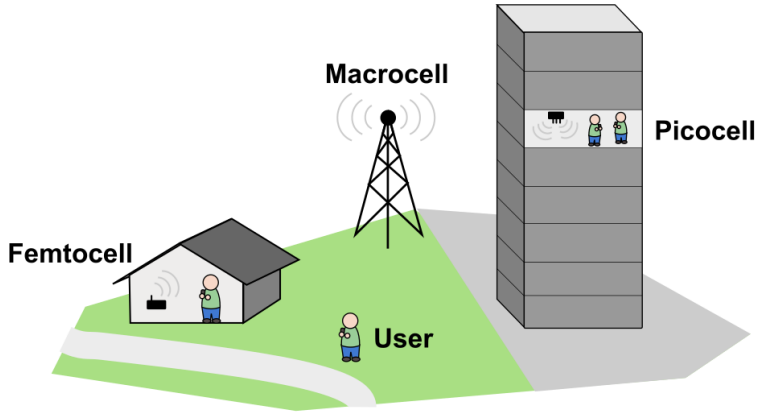

Figure 2.3. Examples of various wireless devices.

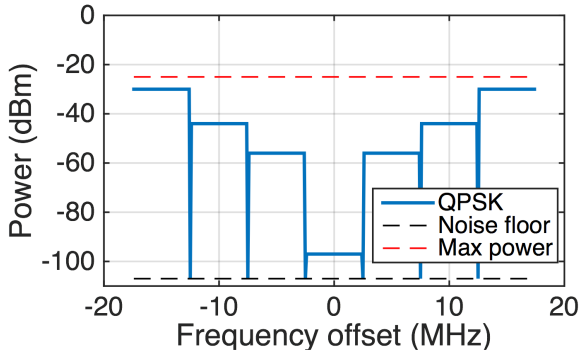

Figure 2.4. Adjacent channel selectivity requirement of 5 MHz LTE. Data from [41].

2.1.2 LTE standards

The Long-Term Evolution (LTE) standard is used globally to connect mobile devices and data terminals. As an extension of previously existing wireless standards designed with a broad range of bandwidths, LTE requires support for 1.4, 5, 10, 15, and 20 MHz bandwidths. This standard also supports carrier aggregation, in which up to five bands can be used for communication simultaneously to enhance bandwidth. Larger bandwidths are utilized to support higher data rates for modern wireless systems, while the smaller bandwidths are required to maintain compatibility with pre-existing standards such as GSM (Global System for Mobile communications) and CDMA (Code Division Multiple Access). The LTE standard operates over a broad range of frequency bands, from 700 MHz up to a few bands in the 3.7 GHz space [40].

The bandwidth support, carrier frequency range, and signal strength levels specified by the wireless protocol ultimately determine the sample rate and dynamic range of the ADC. A blocker mask profile as shown in Fig. 2.5 for 5 MHz LTE [41] is a useful tool for determining these parameters. This plot shows the minimum signal strength the receiver must detect while maintaining <5% symbol error rate, even in the presence of adjacent channel blockers (positioned at integer multiples of the signal bandwidth from the carrier frequency) with the power levels shown. Overall, Fig. 2.4 illustrates that LTE requires narrow signal bandwidths

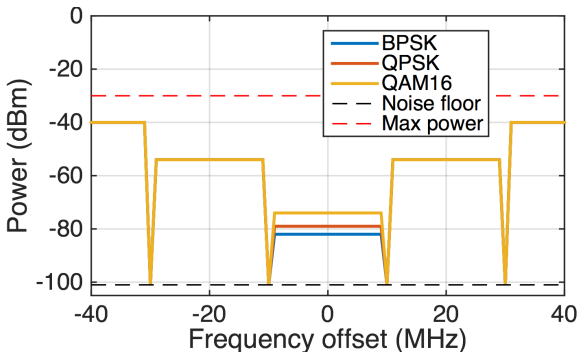

Figure 2.5. Adjacent channel selectivity requirement of 20 MHz Wi-Fi. Data from [43].

but also larger blockers. Specifically, the receiver must be able to detect signals down to -94 dBm with blockers up to 38 dB higher in the adjacent 5 MHz channel, and 64 dB higher only 15 MHz away from the desired channel.

2.1.3 Wi-Fi standards

The wireless local area network (WLAN) protocol, also branded as Wi-Fi, is the most common protocol providing internet access to laptops and smartphones. First introduced in 1997 as the IEEE 802.11 standard, Wi-Fi initially offered data rates of up to 1 Mbps using exclusively the 2.4 GHz frequency band. Since then, updates to the standard have been implemented to support higher data rates and a broader range of frequency bands to address increasing demand for wireless access. Most devices today utilize either the 2.4 GHz 802.11n standard, providing data rates up to 300 Mbps, or the 5 GHz 802.11ac standard, providing data rates up to 4 Gbps. Both 802.11n and 802.11ac utilize 20 and 40 MHz frequency bands; the 802.11ac standard can also support 80 and 160 MHz bandwidths, and emerging applications require bandwidths up to 320 MHz. A growing number of devices have come to support 60 GHz frequencies using the 802.11ad standard [42].

Relative to LTE, Wi-Fi requires wider signal bandwidths that have grown as the standard has evolved to provide higher data rates. Figure 2.5 illustrates the blocker profile that must be supported by 20 MHz Wi-Fi. For 20 MHz bandwidth operation, the receiver must be able to detect signals down to -74 dBm using the highest-throughput modulation scheme (16-level quadrature amplitude modulation, QAM) and -82 dBm using the lowest-throughput modulation scheme (binary phase-shift keying, BPSK). It must maintain a symbol error rate under 5% at these signal strengths even in the presence of -54 dBm adjacent channel blockers. These requirements are somewhat relaxed relative to LTE, as Wi-Fi is typically intended for picocells or femtocells instead of communication directly to a macrocell base station from long distances.

2.2 Receiver & ADC specifications

The requirements outlined by a particular wireless standard specify the minimum signal levels a receiver must detect in a given bandwidth, even in the presence of strong blockers. Because the ADC is typically optimized to handle large-amplitude signals but a receiver must detect signals on the order of microvolts, the front-end circuitry must amplify the received signal to

match the full-scale range of the ADC. Additional filtering in the receiver attenuates large blockers prior to the ADC, relaxing its required dynamic range and preventing large signals at frequencies higher than the ADC sample rate from folding onto the signal bandwidth. After providing a brief overview of major receiver design parameters, this section describes how the ADC sample rate and dynamic range are shaped by the desired wireless standard and receiver characteristics.

2.2.1 Receiver design

As illustrated in Fig. 1.4, a typical receiver chain consists of low-noise amplification, down-conversion, and filtering prior to the ADC. These stages typically occur in that order, though alternative architectures (e.g. mixer-first receivers [12]) have been proposed. The main parameters required to characterize a receiver, in addition to its bandwidth and frequency range, are its noise contribution and linearity. These requirements are dictated predominantly by the standard; the gain and filter order of the receiver can be co-optimized with the ADC to simplify its overall design.

2.2.1.1 *Noise*

Receiver noise must be minimized to detect low-power signals. This is generally accomplished using a carefully designed LNA, which provides impedance matching and a modest amount of gain to relax the noise requirements of successive stages of the receiver. While LNA design details are outside the scope of this work, a receiver's noise performance is typically summarized in terms of its noise figure (NF), which indicates the contribution of the receiver noise relative to the noise generated by the signal source. Because this source is typically a 50Ω antenna, the noise figure can be calculated by dividing the total effective noise at the receiver input by the noise of just the source. Assuming this is thermal noise specified by Boltzmann's constant (k) and temperature (T), the noise figure is calculated as:

$$NF = 10\log_{10}\left(\frac{4kT(50\Omega) + Noise_{receiver}}{4kT(50\Omega)}\right)$$

While the above expression implies that the noise figure of an impedance-matched receiver is limited to 3 dB, various design techniques such as noise cancellation can afford sub-3 dB noise figure. The minimum detectable input power (P_{min}) can be calculated as a function of the signal bandwidth (f_{BW}) and noise figure (NF) as follows:

$$P_{min} = -174\text{ dBm} + 10\log_{10}(f_{BW}) + NF$$

The -174 dBm limit is set by the thermal noise floor of a 50Ω antenna at room temperature. In the context of ADC design, the receiver NF predominantly dictates the amount of additional amplification required to align with the smallest detectable signal of the converter.

2.2.1.2 *Linearity*

To detect both large-amplitude blockers and in-band signals that may be close to the noise floor, the transfer function of a receiver must be highly linear, as shown in Fig. 2.6. Nonlinearities in the receiver transfer function will generate harmonic distortion products. Without sufficient linearity, large blockers may generate spurious tones that fall in the desired frequency bandwidth. Practically speaking, linearity is limited by both the finite supply voltage

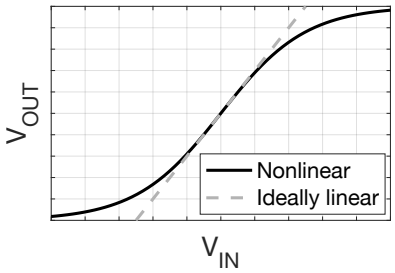

Figure 2.6. Ideally linear (dashed) vs. nonlinear (solid) voltage transfer curve.

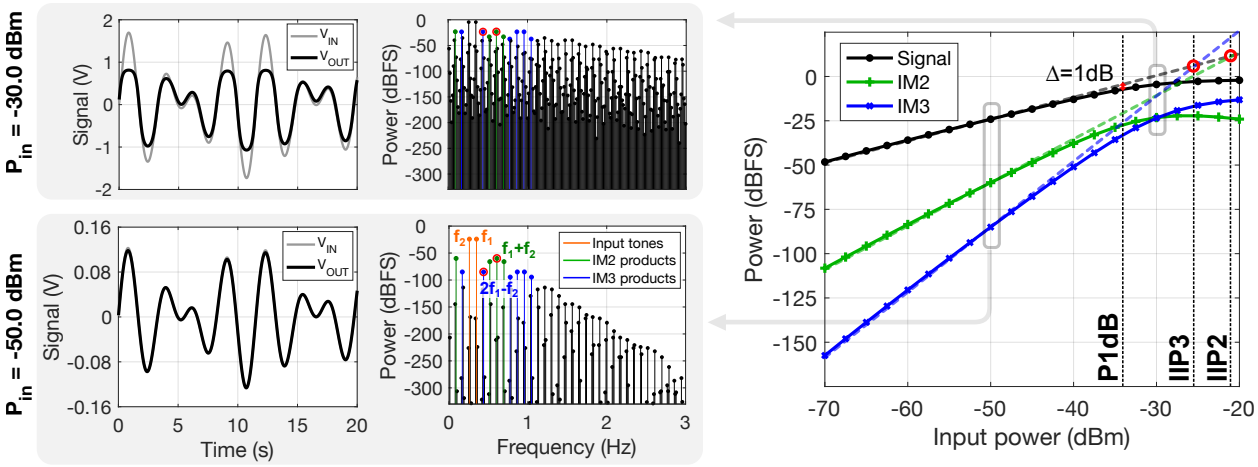

Figure 2.7. Linearity characterization metrics.

of the circuit (e.g., maximum and minimum output bounds in Fig. 2.6) and the nonlinear voltage-to-current transfer characteristic of the transistors that comprise the receiver. A broad range of architectural and circuit-level techniques can be applied to improve linearity, but this section focuses on how receiver linearity is characterized and its implications on ADC design.

Figure 2.7 illustrates how receiver nonlinearity is characterized using a two-tone test. The left sub-plots compare the distorted signal to the ideal signal in the time domain, and the center plot shows the receiver output in the frequency domain. For two tones at frequencies of f_1 and f_2, second order intermodulation products (IM2) can be seen at $f_1 \pm f_2$, $2f_1$ and $2f_2$. Third order intermodulation products (IM3) exist at $2f_1 \pm f_2$, $2f_2 \pm f_1$, $3f_1$, and $3f_2$. As the right plot shows, these intermodulation products grow exponentially with input power until the receiver output begins to saturate. These curves are used to measure the 2^{nd} and 3^{rd}-order intercept point input power level (IIP2 and IIP3) and 1 dB compression point (P_{1dB}) of the receiver. The IIP3 is an especially relevant metric because the spurious tones at $2f_1 - f_2$ and $2f_2 - f_1$ generated by large blockers can fall in the desired frequency bandwidth.

While the ADC design requirements are largely separate from the linearity of the receiver, distortion from the receiver will contribute to the overall signal to noise and distortion ratio (SDNR) of any signal received by the ADC. If the receiver linearity is poor, the noise

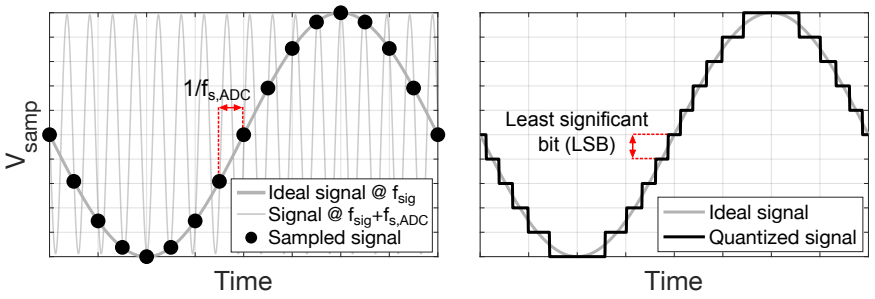

Figure 2.8. ADC performance limitations — ADC sample rate ($f_{s,ADC}$) and aliasing effects, left, and dynamic range (minimum detectable signal amplitude), right.

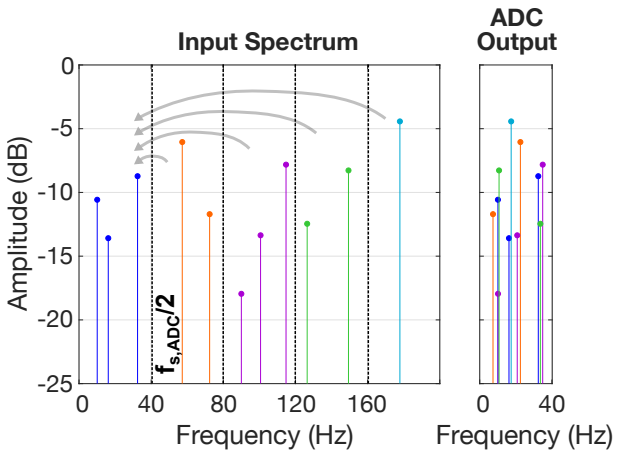

Figure 2.9. Aliasing effects in the frequency domain.

and distortion of the ADC should be lower than the distortion generated by the receiver to avoid further degrading performance. Moreover, the interface between the ADC and the last stage of the receiver may influence the overall receiver linearity.

2.2.2 ADC design

ADC design requirements will depend on both the wireless standard(s) supported and the properties of the receiver preceding the ADC. As illustrated in Fig. 2.11, the converter discretizes a signal in both time and amplitude; two key corresponding metrics required to characterize the ADC are its sample rate and dynamic range. The figure illustrates how discrete sampling causes signals at frequencies higher than the ADC sample rate to be indistinguishable from lower-frequency signals. This frequency folding property is known as aliasing. As a result, the ADC sample rate will depend on both the degree of receiver filtering and target blocker profile. The dynamic range of the ADC will be affected by both these two factors and the signal-to-noise ratio (SNR) required to support a chosen modulation scheme. The next sections provide a discussion of how these parameters are selected.

2.2.2.1 *Sample rate*

The sampling frequency of an ADC determines the maximum bandwidth that the receiver can detect reliably. An ADC's sample rate ($f_{s,ADC}$) can be chosen either for Nyquist-

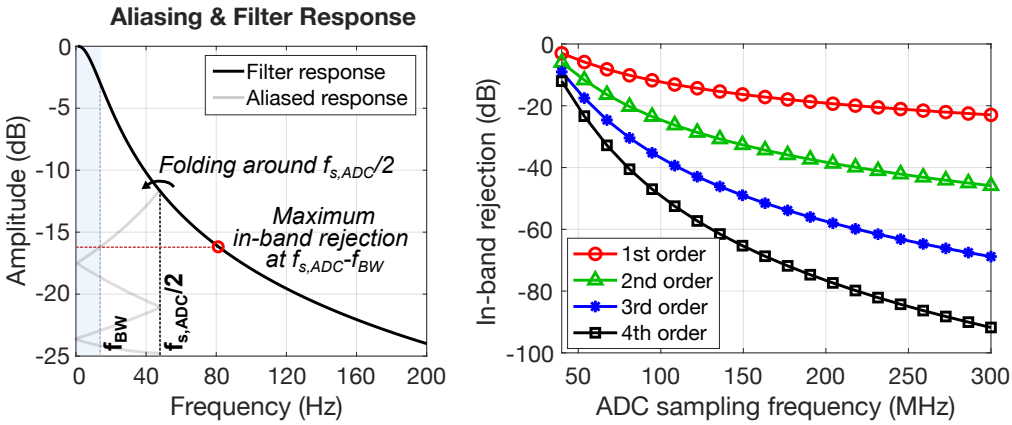

(a) Determining maximum achievable attenuation from bandwidth f_{BW} and ADC sample rate ($f_{s,ADC}$). (b) Maximum attenuation vs. filter order for $f_{BW} = 10$ MHz.

Figure 2.10. Impact of ADC sampling rate on anti-alias filtering.

rate operation ($f_{s,ADC}$ is twice the signal bandwidth, f_{BW}) or for oversampled operation ($f_{s,ADC} > 2f_{BW}$). Oversampling is used in most receiver applications because it relaxes the amount of active filtering that must precede the ADC. Additionally, certain oversampled ADC architectures allow for separate signal and noise transfer functions, thereby improving the in-band SNR through oversampling (see Section 3.1.2).

Active receiver filtering is typically required to mitigate the aliasing effect shown in Fig. 2.11. Figures 2.9 and 2.10 illustrate the effects of aliasing in the frequency domain. As shown in Fig. 2.9, any frequencies over $f_{samp}/2$ will fold into the sampled ADC spectrum, making higher frequency signals indistinguishable from lower-frequency ones. In order to sustain acceptable in-band SNR, any blockers that alias into the signal bandwidth must be attenuated by the receiver.

Due to aliasing, a tradeoff exists between the ADC sampling rate and the order of the filter prior to the ADC. Figure 2.10(b) shows the maximum in-band anti-alias filtering as a function of ADC sample rate and receiver filter order. For reference, Fig. 2.10(a) illustrates how these parameters are calculated — the maximum rejection is evaluated as the filter rejection at the lowest frequency that will fold into the signal bandwidth. To be tolerable, any blockers that fold into the signal bandwidth must not significantly degrade the in-band signal to noise ratio of the receiver. Therefore, the blockers must be attenuated to a fixed level below the minimum signal power. Considering the 20 MHz Wi-Fi blocker requirements illustrated in Fig. 2.5, at least 40 dB blocker rejection is required. Using the example shown in Fig. 2.10(b), this is achievable with sampling frequencies of 80-110 MHz using a 3^{rd} or 4^{th} order filter, but the sampling frequency must exceed 200 MHz to support a 2^{nd} order receiver filter. Operating the ADC at a higher sample rate increases the ADC power but simplifies the design of the filter, leading to a tradeoff that must be considered when the full wireless system is designed.

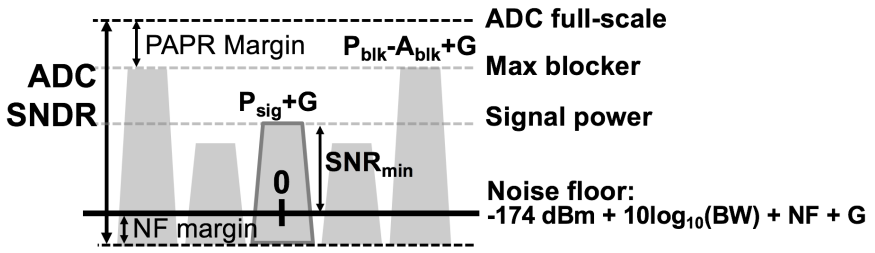

Figure 2.11. Receiver design parameters + ADC design requirements.

2.2.2.2 *Dynamic range*

The dynamic range of an ADC indicates the maximum and minimum signal levels that the converter can detect simultaneously. In a receiver, an ADC's required dynamic range is set by the amplitude of signals the converter may encounter after the receiver filter. By attenuating large blockers, the receiver filter relaxes the dynamic range requirements of the ADC in addition to suppressing high-frequency tones that may alias into the signal bandwidth. The receiver gain is then chosen to ensure that the maximum signal from the receiver aligns with the full-scale input range of the ADC.

Figure 2.11 presents a more detailed look at how the receiver gain, blocker attenuation, and required channel signal to noise ratio (SNR) dictate the ADC dynamic range. This dynamic range is measured as the required signal to noise and distortion ratio (SNDR) of the ADC. As shown in Fig. 2.11, the receiver gain should ensure that the full-scale amplitude of the ADC aligns with the peak amplitude of the receiver output. This peak amplitude can be calculated by adding margin to the average blocker power (P_{blk}) that has been amplified by the receiver gain G but attenuated by a factor of A_{blk} from the filter. The amount of margin depends on the peak-to-average power ratio (PAPR) of a particular modulation scheme; for OFDM signals, the typical PAPR is 12 dB. The noise floor of the ADC will be determined by adding margin onto the noise floor of the receiver. To ensure that the noise contribution of the ADC degrades the receiver's noise figure (NF) by only ΔNF dB, the margin M can be calculated as follows:

$$M = 10\log_{10}\left(10^{\frac{\Delta NF}{10}} - 1\right)$$

Assuming that the receiver gain G is sufficient to boost the peak signal to the full-scale range of the ADC, the required SNDR illustrated in Fig. 2.11 is independent of G and can be calculated as:

$$SNDR = [P_{blk} - A_{blk} + PAPR] - [-174\text{ dBm} + 10\log_{10}(f_{BW}) + NF - M]$$

The dynamic range requirements can be relaxed by enhancing the order of the receiver filter to increase A_{blk}. Most other parameters will be dictated by the requirements outlined in a standard, resulting in a tradeoff between the filter complexity and ADC complexity that must be carefully considered when designing a full wireless system.

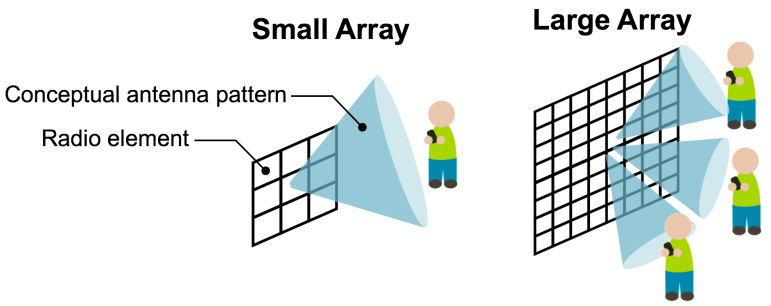

Figure 2.12. Concept of spatial multiplexing via beamforming. Increasing the number of elements in the array improves directivity to reliably support more users.

2.3 Alternative receiver architectures

The discussion in the previous section pertains to the standard receiver architecture shown in Fig. 1.4, in which the RF signal from a single antenna is fed to a low-noise amplifier, mixer, and additional filtering and gain stages prior to the ADC. However, alternative architectures exist that can improve throughput or implementation feasibility. For instance, emerging wireless systems such as massive multiple input, multiple output (MIMO) beamforming systems utilize an array antennas to support many users through spatial selectivity. Alternatively, direct RF-to-digital receivers constructed from ADCs optimized for wireless systems can enable baseband filtering to be implemented in the digital domain. This section describes these two applications in more detail, as they pertain to the two major prototypes presented in this dissertation.

2.3.1 MIMO radio arrays

Previous wireless standards that have evolved to meet the exponentially growing demands for spectral capacity incorporate a variety of access and encoding schemes to utilize existing resources in frequency and time. Spatial multiplexing (Fig. 2.12) provides an additional dimension for growth, using directed beams that can help mitigate interference from other devices that traditionally limits the reliability of wireless systems. Beamforming is accomplished using arrays of radio elements with unique phase shifts applied to direct transmitted signals and receive signals coming from a particular direction. While beamforming coefficients can be applied in the receive direction in the analog domain, digital beamforming provides a much higher degree of flexibility to support multiple users with customizable antenna patterns and array sizes. To facilitate digital beamforming with low power consumption, however, moderate-resolution and low-power ADCs are required. This section first describes how the array size places unique requirements on the resolution of ADCs in radios, and then discusses how to translate array-level performance requirements to the specific design requirements of a fixed-performance ADC. The final section describes what this means in the context of designing a configurable ADC suitable for this application.

2.3.1.1 *Array scaling properties*

In addition to enabling spatial directivity, averaging signals from multiple antennas reduces uncorrelated noise. Because the signal is correlated between N_{ant} array elements, the total

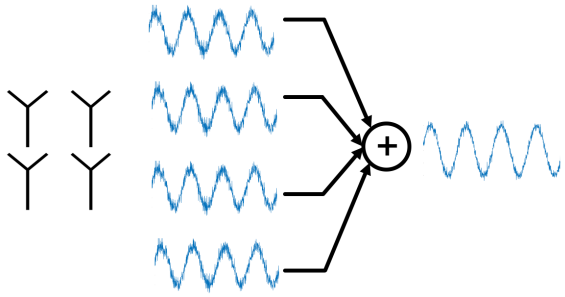

Figure 2.13. Example of 4-element array averaging.

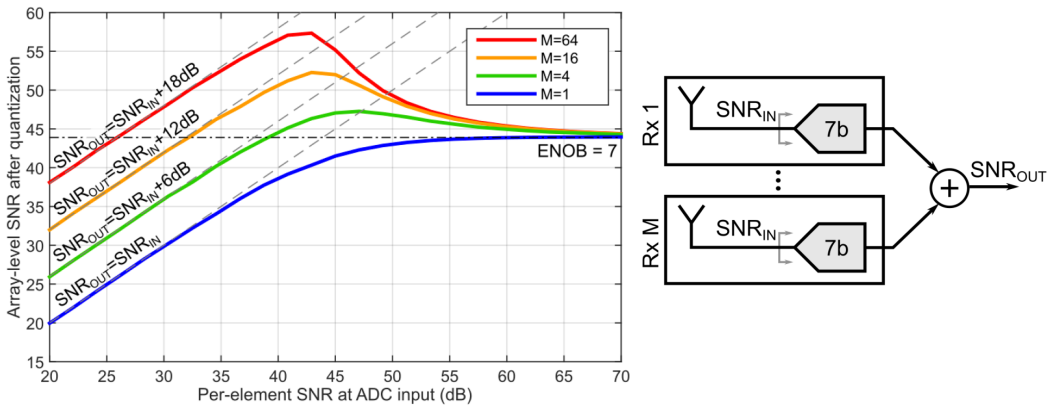

Figure 2.14. Theoretical array-level signal-to-noise ratio (SNR) achievable as a function of input SNR and number of receiver elements M, as shown in [44].

array-level signal power (P_{sig}) will scale with N_{ant}^2, while the uncorrelated noise power (P_{noise}) will scale with only N_{ant}. This means that the array-level signal-to-noise ratio (SNR) should improve by 3 dB for every factor of 2 increase in N_{ant}:

$$SNR_{dB} \quad \propto \quad 10\log_{10}\left(\frac{P_{sig}}{P_{noise}}\right) = 10\log_{10}\left(N_{ant}\right)$$

This means that the radio array provides an effective gain that increases the radio's SNR by $10\log_{10}(N_{ant})$ dB if the noise of each radio is uncorrelated between array elements. As a result, radio arrays can provide the same overall noise performance of a single-radio system using radio elements with a much higher noise floor than conventional high-performance designs, as discussed in [44] and illustrated in Fig. 2.13. Relaxing the noise requirements on each element allows arrays to be constructed from low-power elements. While building full beamforming arrays from conventional high-performance radio elements would scale the array power consumption by a factor of N_{ant}, using low-power elements can keep the radio power consumption relatively constant as N_{ant} grows.

2.3.1.2 *Translating array requirements to ADC requirements*

Selecting an appropriate ADC topology requires understanding how the system-level design requirements for the full radio array dictate the speed and resolution requirements of the

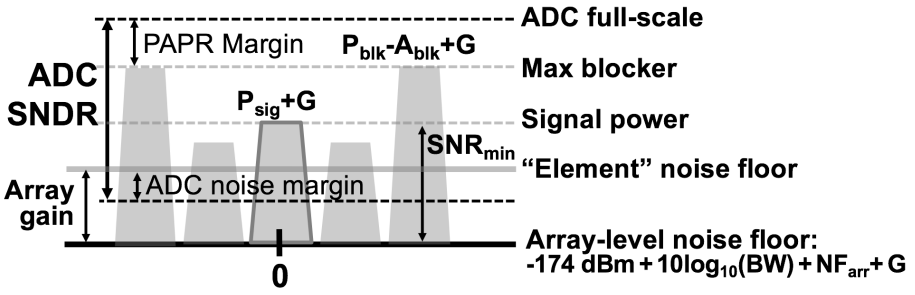

Figure 2.15. Receiver and array design parameters + ADC design requirements.

per-element ADC. As discussed, the array must be able to reliably detect signals in a target bandwidth and power range even in the presence of adjacent channel blockers. For a radio array, while the receiver can tolerate high noise figure (NF), it cannot tolerate a high amount of correlated distortion that will not be eliminated via array averaging. As an example of how uncorrelated noise impacts the array averaging benefits of a receiver, Fig. 2.14 shows the array-level SNR as a function of the SNR at the input of an ideal 7-bit ADC for different array sizes. An ideal 7-bit ADC has an SNR of 44 dB, as marked by the horizontal dashed line in the figure. When the SNR at the ADC input is significantly lower than 44 dB — indicating that uncorrelated thermal noise between elements is the dominant noise source — the array-level SNR improves by the expected $10\log_{10}(N_{ant})$ dB. However, the array-level SNR approaches the 44 dB resolution limit as the SNR at the ADC input increases, leaving correlated quantization noise to dominate the noise of the receiver. While the quantization noise can be randomized with dithering techniques, Fig. 2.14 highlights the importance of minimizing uncorrelated noise in receivers for MIMO applications.

Assuming that uncorrelated noise dominates the effective SNR of the receiver, Fig. 2.15 summarizes how the dynamic range requirement of the ADC is affected by the array-level SNR boost. Relative to the conventional ADC SNDR requirement shown in Fig. 2.11, the per-element noise floor can be shifted upwards by the array gain of $10\log_{10}(N_{ant})$ dB. As a result, the SNDR can be calculated as:

$$SNDR = [P_{blk} - A_{blk} + PAPR] - [-174\text{ dBm} + 10\log_{10}(f_{BW}) + NF - M + 10\log_{10}(N_{ant})]$$

While the same peak-to-average power ratio (PAPR) margin and noise figure degradation margin M discussed in Section 2.2.2.2 must be accounted for, the overall SNDR requirements can be relaxed by $10\log_{10}(N_{ant})$ dB, or 1 effective bit for every 4x increase in the number of array elements.

The array properties also relax the sample rate requirements of the ADC. As shown in Fig. 2.16, any blockers that fold into the signal bandwidth will be attenuated by not only the rejection from individual receiver elements, but also by spatial filtering from the array. This assumes that blockers will originate from a different direction than the desired signal. Because spatial filtering is performed after the array averaging, the ADC must still have a dynamic range high enough to detect these large blockers, but within each single radio element they may fold into the signal bandwidth if there is additional spatial filtering after the array-level averaging in the digital domain. Therefore, a modest oversampling rate may be acceptable for converters in MIMO arrays.

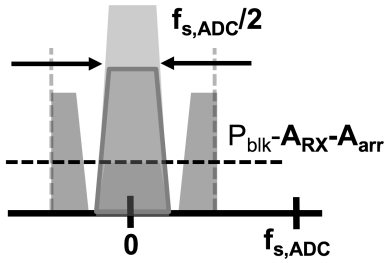

Figure 2.16. Aliased blockers with power P_{blk} are attenuated by spatial filtering from receiver (A_{arr}) in addition to filtering from the receiver (A_{RX}).

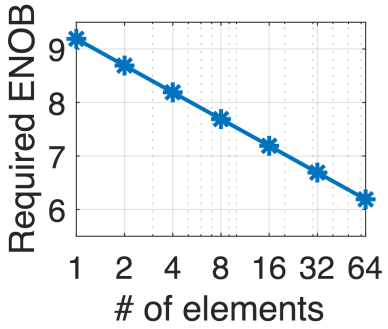

Figure 2.17. Required ADC resolution vs. array size. Assumes target array-level noise figure of 6 dB, maximum blocker strength of -20 dBm, and 3rd order filtering from the receiver.

2.3.1.3 *Considerations for scalable ADCs in MIMO systems*

As discussed, array averaging enables the use of low-resolution ADCs with a moderate oversampling ratio. For target baseband bandwidths of up to 20 MHz suitable for Wi-Fi and LTE applications, an ADC sampling rate of 80 MS/s provides a 4x oversampling ratio that can be used with a third-order filter in the receiver. To support an array-level noise figure of 6 dB with a receiver gain of 40 dB, the ADC can have a resolution of 9-6 effective bits without significantly degrading the resolution of the receiver, as shown in Fig. 2.17. In conventional stand-alone receivers, stringent noise figure requirements encourage the use of highly oversampled converters that can relax the receivers anti-aliasing filter requirements and allow for most filtering to be implemented in the digital domain. The relaxed resolution requirements in a massive MIMO system enable the use of an ADC architecture with lower resolution and speed.

To support per-element power scalability, any fixed overhead power within the ADC must be minimized. This means that any peripheral ADC components, such as external reference voltages or always-on calibration circuitry, should be eliminated if possible. To maintain constant power consumption as the array size increases, the power consumption of the ADC should ideally be cut in half each time the converter's resolution requirements scale by 3 dB. The power scaling range of the converter depends on the selected topology and design of the ADC.

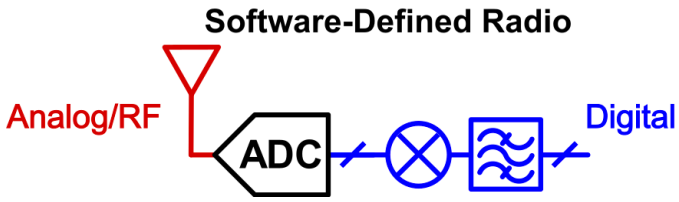

Figure 2.18. Software-defined radio concept.

2.3.2 Software-defined radio

A fully software-defined radio, as proposed in the 1990's [33] and shown in Fig. 2.18, would use a high-resolution ADC to directly digitize the RF spectrum. Down-conversion and filtering would be performed in the digital domain, fully leveraging the benefits of CMOS technology scaling. Realistically, this proposal places such stringent demands on the dynamic range and sampling frequency of the ADC that only a handful of published high-power converters are suitable for these applications [45–47]. As an alternative, both oversampled RF-to-digital receivers and conventional analog receivers with digitally-tunable carrier frequency and bandwidth have been proposed to realize partially software-defined radio. After providing a brief overview of some major approaches to software-defined radio, this section will highlight the major differences between designing ADCs for conventional receivers and for direct RF-to-digital conversion.

2.3.2.1 *Software-defined radio architectures*

A more detailed review of previously published software-defined radio receivers is presented in Section 4.1. Generally speaking, the three main approaches to implementing software-defined radio include the following:

- *High-speed, high-resolution ADCs*: In fully software-defined radio, a high-speed ADC will directly digitize the RF input, allowing down-conversion and filtering to be performed digitally. While recent time-interleaved pipelined ADCs [45–47] have been able to achieve over 60 dB SNDR at sample rates of over 4 GS/s, this performance comes with a substantial power overhead. Converters targeting direct RF operation consume nearly 500 mW, relative to optimized stand-alone receivers which typically consume under 50 mW while offering improved noise figure and blocker tolerance.

- *Direct RF-to-digital receivers*: Relative to high-speed Nyquist-rate ADCs intended to convert the full RF spectrum from DC to the carrier frequency, RF-to-digital receivers typically incorporate a mixer and blocker filtering directly in the converter. This typically repurposes a filter within the ADC, such as the loop filter in an oversampled $\Delta\Sigma$ converter [48], to support blocker rejection. These designs are typically oversampled to enable digital filtering.

- *Digitally-tunable receivers*: Enhancements to existing receiver designs can provide compatibility with software-defined radio by adding components with digitally-configurable performance. This can range from small levels of configurability, such as digitally-tunable passive component sizes (e.g. to enable bandwidth tuning [1]), to fully redesigned architectures (e.g. mixer-first architectures that offer wideband operation

[12]). While this approach does not support a fully software-defined receiver baseband, it can leverage some of the improved digital processing capability offered by CMOS scaling to boost receiver performance.

Many approaches can help support the implementation of software-defined radio. Direct RF-to-digital converters provide a tradeoff between the high power consumption of GS/s Nyquist-rate ADCs and the limited digital processing afforded by making conventional receivers digitally tunable.

2.3.2.2 *Considerations for software-defined radio ADCs*

In software-defined radio, channel filtering is typically performed in the digital domain. As discussed in the previous section, a conventional receiver contains an active analog baseband filter that provides anti-alias filtering and reduces the required dynamic range of the ADC. Without this filter, the ADC needs both a fast sample rate to prevent aliasing and a high dynamic range to tolerate strong out-of-band blockers. This combination of requirements favors oversampled converters with noise shaping, in which the signal and quantization noise have different transfer functions. This allows the ADC to have a high in-band SNDR while requiring only a low-resolution quantizer that is easier to implement with high dynamic range. A further discussion of these design considerations is presented in Chapter 4.

Chapter 3

A resolution-configurable ADC

Performance-tunable components, such variable-gain amplifiers and reconfigurable filters, can serve as building blocks for scalable receivers. To that end, this chapter discusses the design of an ADC that trades resolution for power consumption for use in a power-scalable receiver with tunable noise figure. In addition to being suitable for wireless systems that adapt to their environment, the configurable ADC makes this receiver suitable for digital beamforming, which enables spatial selectivity using an array of radios. By digitally applying appropriate phase shifts prior to summation, these systems can support multiple users with arbitrary antenna patterns. The number of users that can be reliably supported by the array grows with the number of array elements, but incorporating many radio elements can significantly increase power consumption. However, averaging uncorrelated noise between each receiver element in the array allows individual elements to have higher noise figure while maintaining constant array-level performance. As a result, power and resolution scalability is required to sustain energy efficient operation. Scalable ADCs that can trade resolution for power consumption have the potential to keep array-level power consumption comparable to that of a single element with the same single-user performance that cannot support multiple users, provided a sufficient number of elements are used in the array.

This chapter discusses the design and implementation of a resolution- and power-scalable, fixed-speed successive approximation register (SAR) ADC that can be used to build these energy-efficient massive MIMO radio arrays. Previously, Section 2.3.1 provided context for this work by describing design requirements unique to ADCs in massive MIMO systems. In this chapter, Section 3.1 first reviews prior configurable ADC designs to illustrate why existing scalable ADCs are not well-suited to this application, and examine why the successive approximation register (SAR) ADC topology in particular is a logical choice for such designs. Section 3.2 then discusses SAR ADC design in more detail to provide insight into the challenges of developing an efficient resolution- and power-scalable ADC, and Section 3.3 elaborates on the design specifics of the 65nm CMOS prototype. Finally, Section 3.4 reviews measurements of this test chip to emphasize some of the challenges and limitations of designing scalable SAR ADCs for this particular application, which are summarized in Section 3.5.

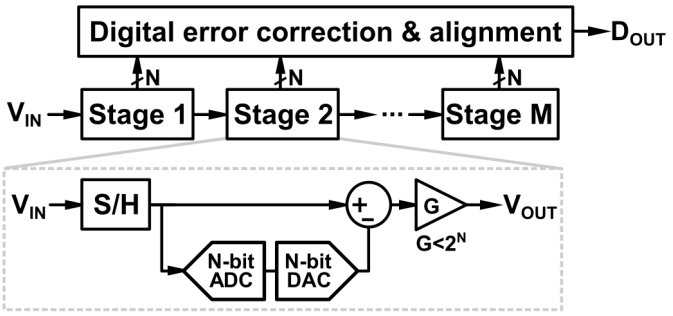

Figure 3.1. General block diagram of a pipelined ADC.

3.1 Prior scalable ADC designs

The architecture and design of an ADC is typically chosen carefully to minimize power consumption for specific speed and resolution requirements. However, many applications require different capabilities under different operating conditions. For instance, a multi-mode wireless platform must be able to detect signal bandwidths ranging from 2.5 MHz in LTE systems to above 160 MHz (320 MHz in energy) for 5 GHz Wi-Fi platforms. In biomedical sensing applications where energy efficiency is crucial, monitoring different vital signs may require different sample resolutions and frequencies, so enabling an ADC to trade resolution for power can optimize the efficiency of these nodes. As a result, many prior efforts have studied the feasibility of developing ADC topologies that can trade performance for power consumption. This section describes some of the major prior work in this area, grouped by converter topology.

3.1.1 Pipelined ADCs

In a pipelined ADC, analog-to-digital conversion is performed by a sequence of low-resolution stages that sample the input signal, perform coarse analog-to-digital conversion, and then amplify the residue for conversion in the next stage, as illustrated in Fig. 3.1. While pipelined topologies are often well-suited for high-speed applications, they require amplification in each stage that increases the overall power consumption of the ADC. These amplifiers are often constructed from active op-amps that require a large gain-bandwidth product to achieve high resolution at fast speeds. Early research into scalable ADC design largely focused on pipelined ADCs with a variable number of stages or variable-performance amplifiers to trade speed for conversion rate. Most designs explore means of tuning conversion rate while keeping power consumption low for applications in multi-standard wireless systems with variable baseband bandwidths.

To support multiple standards, the work in [49] uses a variable number of time interleaved stages to support different conversion rates and includes an optional flash ADC based final stage to tune resolution. In 802.11b Wi-Fi mode, the work in [49] provides 60 dB of SNDR at 44 MS/s while drawing 20.2 mW, while in Bluetooth mode it provides 64 dB of SNDR at 11 MS/s and draws 14.8 mW. The work in [50] also bypasses stages to tune resolution from 10-12 bits. To achieve additional power scaling, the work in [50] also scales the current provided to op-amps used as inter-stage amplifiers. Overall, the combination of amplifier and architecture scaling results in a Walden figure of merit (FoM) of 0.35-0.5

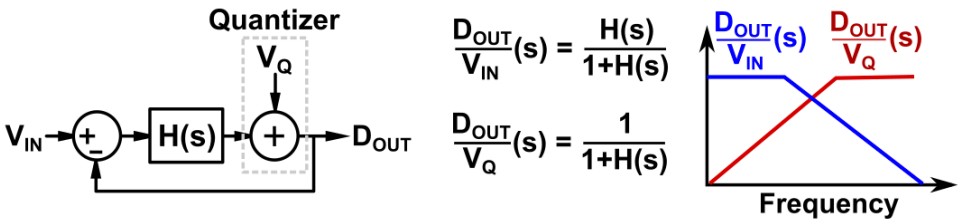

Figure 3.2. General block diagram of $\Delta\Sigma$ converter (left) and corresponding signal (V_{IN}) and quantization noise (V_Q) transfer functions. With an integrating loop filter $H(s) = 1/s$, the transfer functions take the form illustrated at right.

pJ/step maintained across a 0.4-44 MS/s speed range.

In [51] and [52], a power-scalable pipelined ADC for sub-sampling wireless receivers is presented that achieves a broad sampling speed tuning range using fast-settling, current modulated op-amps that only draw power during conversion. This allows the op-amp power to scale directly with frequency in a manner similar to digital circuits. In [51], the power ranges from 0.7 mW to 35 mW as the sampling frequency scales from 0.17 MS/s to 20.94 MS/s while the SNDR is kept above 55 dB, while the work in [52] lowers the peak power consumption to 27 mW by relaxing the requirements on the first stage sample and hold of the pipelined converter.

Most prior scalable pipelined ADCs emphasize speed configurability to support the range of bandwidths required by different wireless standards. This typically involves using relaxed settling time requirements to scale the op-amp power consumption [50–52]. However, this analog-intensive approach requires high-gain amplifiers that can be difficult to construct in scaled CMOS process technologies. Moreover, the large static power (>10 mW) consumed by the amplifiers results in a significant power overhead. As a result, modern ADCs for wireless applications use either an oversampling-based architecture that can relax the receiver's design requirements, or successive approximation ADCs that are inherently process-scalable.

3.1.2 Oversampled ADCs

In oversampled ADCs, the converter samples at a frequency higher than the target bandwidth to allow out-of-band signals to be filtered digitally. This is typically implemented as a $\Delta\Sigma$ converter, in which the signal and quantization noise have different transfer functions to the final digital output, as illustrated in Fig. 3.2. A first-order $\Delta\Sigma$ converter integrates the difference between the analog input signal and its digital approximation, pushing quantization noise to higher frequencies while maintaining a low-pass response for the signal. For many high-performance wireless systems, the oversampling nature of $\Delta\Sigma$ converters can greatly relax the required anti-alias filter order. Similar to pipelined converters, configurable $\Delta\Sigma$ converters have typically focused on means of tuning bandwidth to support various wireless standards.

The oversampling ratio provides one means of configuring $\Delta\Sigma$ performance. One of the early tunable $\Delta\Sigma$ converters presented in [53] suggests tuning the converter architecture to operate as either a $\Delta\Sigma$ converter or an incremental ADC, which averages a number of $\Delta\Sigma$ ADC samples to generate a filtered output. Similarly, [54] describes a resolution-configurable ADC for touch sensing applications that uses a variable number of averaged cycles in an

incremental $\Delta\Sigma$ ADC to trade resolution for frame rate.

Alternatively, the performance of the converter can be configured using the loop filter. The work in [55] uses a hybrid LC/active RC loop filter with tunable components to configure the loop filter characteristics optimized for baseband (0 MHz) and IF (450 MHz, 1 GHz) center frequencies. More recently, [56] uses a reconfigurable loop filter architecture to implement tunable blocker rejection for a fixed bandwidth and overall SNDR. For general reference, [57] presents a methodology for optimizing $\Delta\Sigma$ converters for wireless radios by selecting an optimal number of conversion bits, oversampling ratio, and loop filter topology. A broader overview of $\Delta\Sigma$ converters for wireless applications is presented in [58].

An alternative $\Delta\Sigma$ topology uses a voltage-controlled oscillator (VCO) and digital counter to implement the integrator and quantizer, operating as a feed-forward $\Delta\Sigma$ converter with differentiation performed in the digital domain. This architecture will be discussed in further detail in Chapter 4, but as an example, the work in [59] presents a configurable VCO-based ADC design with background calibration. It maintains an SNDR above 70 dB for sample rates ranging from 1.3-2.4 GHz, with power that scales from 11.5-39 mW across this range. The digital-intensive implementation of this design causes the power consumption to scale directly with sample rate.

Overall, a broad range of $\Delta\Sigma$ architectures can be used to provide higher resolution through oversampling and noise shaping. While some converter architectures are amenable to implementation in advanced CMOS process nodes (e.g. [59]), most require active analog integrators to implement the loop filter. The op-amps required to construct such integrators typically consume static power that often exceeds the power consumption of fully dynamic designs that can benefit from CMOS scaling.

3.1.3 Successive approximation ADCs

In a successive approximation register (SAR) ADC, a binary search algorithm is used to convert the sampled analog input signal to a digital output. A SAR ADC typically uses a capacitive digital-to-analog converter (DAC) to subtract a reference amount from the sampled input signal, a clocked comparator, and digital logic to execute the algorithm, resulting in an architecture highly compatible with deeply scaled CMOS technologies optimized for digital performance. High-performance ADCs have been demonstrated down to 14nm FinFET. The next section discusses the design of SAR ADCs in further detail. These ADCs also do not require precise amplification, resulting in a low-power and energy-efficient design. However, the resolution and speed of these converters is limited by the performance of the comparator and speed of the digital logic. As a result, many previous implementations of scalable SAR ADCs focus on low-speed, low-power, moderate-resolution wireless sensor node applications [60–64], though recent work has emerged using SAR ADCs for higher-performance wireless receivers [65, 66].

Because most of the total power consumption in a SAR ADC is digital, it scales according to CV_{DD}^2, where C is the total switched capacitance and V_{DD} is the supply voltage of the chip. As a result, substantial power savings can be obtained by lowering V_{DD} when lower speed is tolerable. Previous research has used supply voltage control as a means of scaling speed or power consumption. The work in [60, 61] uses supply voltage scaling and tunable component sizes to tune power consumption from 116-206 nW while scaling the resolution from 36.6-55 dB and keeping the sampling rate fixed at 20 kS/s. Similarly, [67] uses supply

voltage scaling to tune power consumption from 0.5-11μW while scaling conversion speed from 0.5-4 MS/s and maintaining an SNDR of 55 dB, while [68] scales power from 0.4-85 μW while tuning sample rate from 0.5-30 MS/s with a fixed SNDR of 44 dB. Additional description of the impact of V_{DD} on SAR ADC operation is presented in [62]. While supply scaling can help optimize performance, it also increases system complexity and requires additional power overhead in the form of tunable voltage references.

Alternatively, ADC resolution can be configured by using tunable components. The work in [63] and [64] presents a resolution and speed configurable ADC using a configurable DAC size and performance-tunable comparator. This allows the design to scale resolution from 6.9-9.3 effective bits and scaling power from 1.6-3.6 μW while sampling at 2 MS/s. The work in [65] integrates a SAR ADC within a triple-mode transceiver, using a variable number of ADC cycles to tune the resolution and speed of the converter for GSM and LTE operation. Alternatively, tunable receiver gain can be integrated within the ADC; [66] incorporates gain into the ADC sampling buffer by using a variable-size G_m cell. This converts the input voltage to a current that is integrated onto the DAC capacitance for a time set by the desired ADC sample rate.

Overall, SAR ADCs can provide highly efficient, digital-intensive operation suitable for moderate conversion speeds and resolution levels. While previous work has demonstrated that highly efficient scalable SAR ADCs can be built for low-power sensor applications, these designs do not support the moderate conversion rates needed for wireless applications. Moreover, prior designs achieve low figure of merit by reducing V_{DD} to significantly reduce power when slow ADC operation is tolerable, but this technique cannot be applied to fixed-bandwidth operation. Advances in CMOS scaling have improved ADC sample rates to levels suitable for wireless baseband applications, though previous work developing bandwidth-configurable ADCs to support various wireless standards similarly does not provide resolution scalability.

3.2 SAR ADC design considerations

To illustrate how to develop a tunable successive approximation register (SAR) ADC, this section reviews major aspects of the design of a SAR ADC with fixed speed and resolution and then summarizes what this means in the context of developing tunable converters. It begins with a general overview of SAR ADC design, and then discusses major sources of noise and distortion found in three main building blocks of a conventional SAR ADC: the capacitive DAC, conversion logic, and the comparator.

3.2.1 General topology

Functionally, a SAR ADC uses a binary search algorithm to obtain an increasingly accurate digital approximation of a sampled analog input signal. At each conversion step, the input signal is compared to a reference voltage that is updated based on the result of the previous comparison using a digital-to-analog converter (DAC), as shown in Fig. 3.3.

Specifically, a capacitive DAC-based SAR ADC typically samples an input voltage onto a binary-weighted bank of capacitors, as shown in Fig. 3.4. At each iteration of the conversion, a comparator determines whether this voltage is larger or smaller than a fixed reference voltage, and one terminal of a capacitor in the bank is switched between

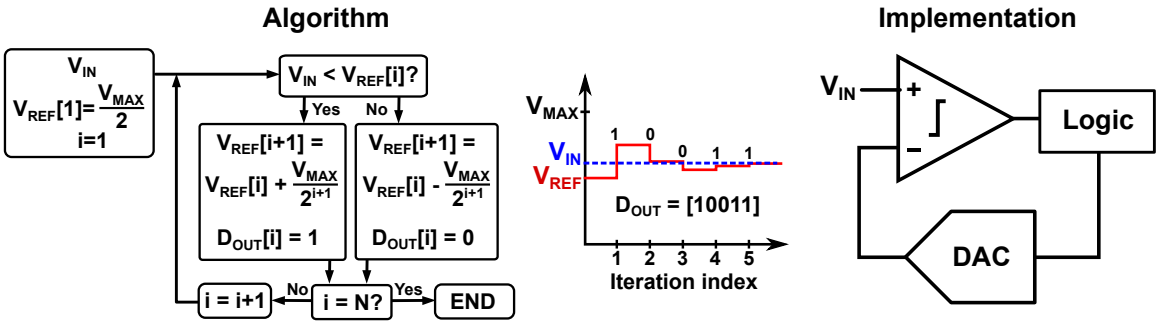

Figure 3.3. General overview of successive approximation algorithm (left) and main implementation blocks (right).

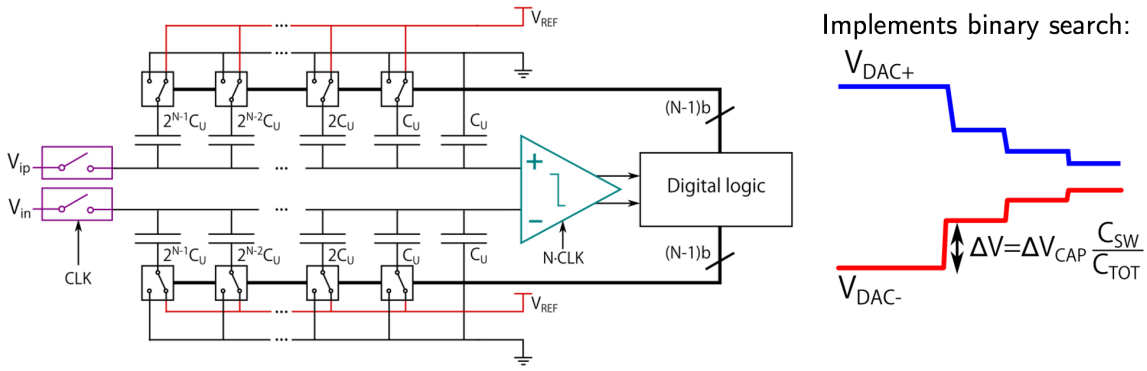

Figure 3.4. Diagram of differential capacitive DAC-based SAR ADC.

reference voltages according to the comparator result. This switching procedure subtracts off a new fraction of the reference voltage from the sampled input signal via charge sharing. This procedure typically progresses for N ADC cycles to obtain an N bit approximation of the input signal, though additional cycles can be used to improve the accuracy of the approximation through redundancy.

The main objective of SAR ADC design is typically to meet the sampling frequency and signal to noise and distortion ratio (SNDR) requirements set by a target application while minimizing power consumption and area. Each component of a SAR ADC introduces non-idealities that limit the converter's SNDR, while design choices that help improve SNDR typically increase the power consumption of each component. Developing an efficient power- and resolution-scalable ADC requires identifying how best to allocate the converter's power budget between different noise and distortion sources for a particular target resolution and operating speed.

In addition to quantization noise, sources of noise and distortion in a SAR ADC include the input-referred noise and offset of the comparator, thermal noise sampled onto the capacitive DAC, mismatch between DAC weights, settling accuracy of the DAC, and the nonlinearity of the input sampling switch. The next sections will describe these effects in more detail, but generally speaking a high-resolution SAR ADC will require a DAC with sampling capacitance large enough to meet thermal noise and component matching constraints, and a comparator with low input-referred noise.

Most power consumption in a capacitive SAR ADC is dynamic, unlike pipelined or $\Delta\Sigma$ converters that typically require active amplification stages that draw static power. Switching the capacitive load of the comparator, logic, and DAC quickly enough to meet the target sampling frequency requirements results in power proportional to CV^2f, where C is the total capacitance, V is the supply voltage, and f is the switching frequency. More specifically, the total ADC power (P_{ADC}) includes a component proportional to the total DAC capacitance (C_{DAC}), supply voltage (V_{DD}), and reference voltage (V_{ref}), a component set by the number of bits (N_{bits}) and the comparator switching energy (E_{comp}), and a component set by the per-bit logic switching energy (E_{comp}):

$$P_{ADC} = f_s \left(\alpha C_{DAC} V_{ref} V_{DD} + N_{bits} E_{comp} + N_{bits} E_{logic}\right)$$

The α factor in the above equation depends upon the DAC switching methodology. As seen above, if C_{DAC} is increased to reduce DAC noise, the power consumed by the ADC will increase accordingly. Similarly, as the comparator size is increased to reduce noise, E_{comp} will also increase. To build a resolution-scalable ADC that can maintain efficient low-resolution operation, both the DAC capacitance and comparator sizing must be tunable, and any fixed power overhead in the ADC (e.g. conversion logic power consumption) must be minimized.

3.2.2 Capacitive DAC and Sampling

The main role of the capacitive DAC is to both sample the input voltage and to subtract a fixed fraction of the reference voltage from this sample at each ADC conversion step. The accuracy of this sample depends on the noise and bandwidth of the sampling network, while the accuracy of the scaled reference voltage is determined by the degree of capacitor mismatch, settling errors, and supply noise. Overall, the DAC power consumption is dominated by the capacitance size, reference voltage, and choice of switching procedure, though additional sources of fixed overhead power exist.

3.2.2.1 *Sources of noise and distortion*

The accuracy of the feedback DAC is a major limitation of ADC performance. The capacitive DAC in a SAR ADC introduces thermal noise due to sampling, distortion due to component mismatch, and additional error due to DAC settling limitations.

Sampling noise The finite resistance of the sampling switch contributes thermal noise each time the input voltage is sampled onto the DAC capacitor. Because the sampling switch resistance (R_{sw}) is directly proportional to the thermal noise variance but inversely proportional to the bandwidth of the noise transfer function, the variance of the total sampled noise is independent of R_{sw}. In contrast, the sampling capacitance C is inversely proportional to the noise bandwidth but has no effect on the thermal noise variance, so the total integrated sampling noise variance is given by kT/C, where k is Boltzmann's constant and T is temperature in K. As a result, a differential DAC design with total capacitance C_{DAC} will see a total sampling noise variance of $4kT/C_{DAC}$ when the noise variances of the two capacitor banks of size $C_{DAC}/2$ add together:

$$\begin{aligned} v_{n,tot}^2 &= v_{n,DACP}^2 + v_{n,DACN}^2 \\ &= \frac{kT}{\frac{1}{2}C_{DAC}} + \frac{kT}{\frac{1}{2}C_{DAC}} = \frac{4kT}{C_{DAC}} \end{aligned}$$

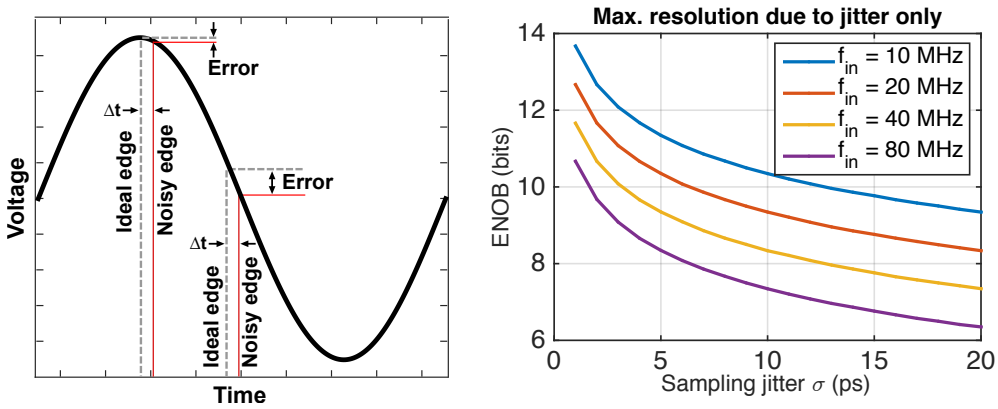

(a) Sampling error due to clock jitter. (b) Simulated resolution dependence on jitter.

Figure 3.5. Sampling jitter effects.

Because the noise voltage variance is inversely proportional to C_{DAC}, relaxing the allowable thermal sampling noise by a factor of 6 dB (1 effective bit) can enable a factor of four reduction in C_{DAC} (and therefore DAC switching power) if the converter's resolution is limited by sampling noise.

Sampling clock jitter Because ADC samples are assumed to occur at fixed time intervals, any variation in the sampling clock will introduce additional variation, as shown in Fig. 3.5(a). This error depends on the particular sampling instant and frequency of the input signal, as sampling clock variation during fast input sample transitions will translate to a larger error than jitter that occurs during slow signal transitions. Assuming a normally distributed clock jitter variance of σ_j, the corresponding maximum achievable ADC resolution is shown in Fig. 3.5(b).

While moderate speed and resolution applications with low-frequency input signals will not be heavily impacted by sampling clock jitter, this can be a significant source of performance degradation for high-speed ADCs. If noise due to sampling jitter becomes significant, the converter's clock source will draw additional power. However, wireless baseband applications targeting signal bandwidths on the order of 20 MHz and below will not experience significant degradation due to sampling clock jitter.

Sampling distortion In addition to thermal noise, the finite resistance of the sampling switch contributes distortion that further degrades the SNDR of the ADC, particularly for high-speed input signals. The low-pass RC filter formed by the finite sampling switch resistance and the sampling capacitance limits the input signal bandwidth, and the signal-dependent sampling switch resistance introduces distortion. Because the switch resistance is set largely by the gate-to-source voltage (V_{GS}) of the switch transistor, it will vary with the input voltage if the source or drain terminal of the sampling switch is connected directly to the input signal and the maximum gate voltage is fixed. The effects of this nonlinearity can be mitigated by lowering the nominal sampling switch resistance to keep the worst-case sampling bandwidth higher than the bandwidth of the input signal. Additionally, a bootstrap circuit can be used to maintain a constant V_{GS} across the switch as the input voltage varies. This approach is particularly necessary when the input signal swing is large,

causing a large fluctuation in the switch V_{GS}. Alternatively, the input signal can be sampled by opening switches at the bottom plate of the sampling capacitors, which will operate with a constant V_{GS}. A full quantitative analysis of sampling switch distortion effects is presented in [69]. Using a long-channel device model, the amplitude of the second and third harmonic signals relative to the input signal (HD_2 and HD_3) are given by:

$$\begin{aligned} HD_2 &= \frac{A}{2}\left(\frac{\omega C_{samp}}{K(V_G - V_{TH})^2}\right) \\ HD_3 &= \frac{A}{4}\left(\frac{\omega C_{samp}}{K(V_G - V_{TH})^3}\right) \end{aligned}$$

In the above expressions, A is the peak-to-peak amplitude of the input signal, ω is the frequency of the input signal in radians/second, C_{samp} is the total sampling capacitance, V_G is the maximum clock input voltage, V_{TH} is the threshold voltage of the sampling switch, and K is a proportionality factor that determines the switch resistance. In a long-channel device model, $K = \mu C_{ox}\frac{W}{L}$, where μ and C_{ox} are process-dependent mobility and gate oxide capacitance parameters and W and L are the switch width and length. The model above assumes that the fall time of the sampling switch is much smaller than the frequency of the input signal, and that the gate voltage applied by the clock is constant. Overall, this emphasizes that distortion can be reduced by lowering the input frequency, C_{samp}, A, or V_{TH}, or by increasing V_G or W/L. Because A is fixed by the output swing of the receiver, V_{TH} is dictated by the process technology, and V_G is limited by the maximum supply of the ADC, distortion can be minimized by keeping C_{samp} as small as possible and using a large W/L for the switch transistor.

Relaxing the resolution requirements on the ADC by a factor of 6 dB allows the spurious free dynamic range (SFDR) of the ADC to increase by this amount. If the thermal noise-limited DAC sampling capacitance scales by a factor of 4 for each bit reduction in required resolution, to first order the sampling switch size can ideally be reduced by a factor of 16, using the HD_2 and HD_3 expressions above. Therefore, the power required to drive the sampling switch will drop accordingly until a minimum-sized sampling switch transistor is used.

Unit capacitor mismatch The accuracy of capacitor weights is another important consideration in designing a high-resolution SAR ADC, as small capacitors are difficult to fabricate accurately due to process variability. Mismatch in capacitor sizes causes nonlinearities in the ADC's transfer characteristic, which maps the analog input voltage of the ADC to a digital output code. Ideally, this would occur in uniform steps according to the relationship $D_{out} = (2^N - 1)\left\lfloor \frac{V_{in} - V_{min}}{V_{max} - V_{min}} \right\rfloor$, where D_{out} is the digital output code, N is the number of ADC bits, V_{in} is the input voltage and V_{max} and V_{min} are the maximum and minimum ADC input voltages. However, mismatch creates systematic nonlinearities in this transfer function, as shown in Fig. 3.7. These nonlinearities can be split into two categories: differential nonlinearity (DNL), which measures code width mismatch, and integral nonlinearity (INL), measures deviation of the ADC's transfer characteristic from an ideal line.

To quantify the effect of DAC unit element mismatch, the unit DAC capacitance can be modeled as a Gaussian random variable with mean C_U and standard deviation σ_u. Because

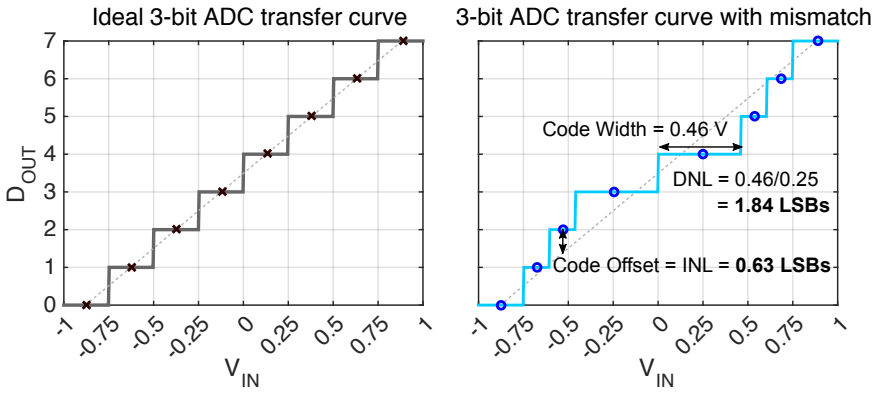

Figure 3.6. Illustration of differential nonlinearity (DNL) and integral nonlinearity (INL) using a 3-bit DAC with capacitor weights of 1.9, 0.46, and 0.8 instead of the ideal weighting factors of 2, 1, and 1.

DNL is determined by mismatch between switched DAC elements, the maximum DNL variance of an N-bit binary-encoded DAC will occur during the $2^{N-1}-1$ to 2^{N-1} code transition. During this transition, $2^{N-1}+2^{N-1}-1=2^N-1$ unit DAC elements will switch, leading to a DNL variance (σ^2_{DNL}) of $2^{N-1}\sigma^2_u$. As derived in [70], the INL variance (σ^2_{INL}) can be approximated as $\sigma^2_U 2^{N-2}$.

Assuming $\sigma_u \propto 1/\sqrt{C_u}$, this model can be used to compute how the DAC capacitance can be scaled while keeping σ_{DNL} fixed if the required ADC resolution is relaxed by a single bit. Because eliminating 1 conversion bit cuts the number of switched elements in half, σ_{DNL} will be constant if σ_u is also cut in half. This translates to a total $N-1$ bit DAC capacitance of $2^{N-1}\frac{C_u}{2}$ relative to the N bit DAC capacitance of $2^N C_u$, facilitating a 4x reduction in capacitance to maintain mismatch-limited scaling.

Note that while increasing DAC capacitance by 4x can guarantee a 4x reduction in thermal noise variance for any specific ADC, the random nature of DAC mismatch will not guarantee that the DNL of a particular scalable ADC will remain fixed when its DAC capacitance scales by 4x. Only the standard deviation of DNL can be compared, which relies upon large-scale statistics. Mismatch calibration techniques such as redundant conversion bits and non-binary DAC weights are required to guarantee that high resolution operation can be achieved, but this translates to a fixed amount of overhead power. To reduce the likelihood of DAC mismatch degrading ADC performance, a thermometer-encoded DAC can be used to limit the number of switched elements. While this does not improve the INL variance, it can be used to limit DNL. Thermometer encoding may reduce σ_{DNL} such that additional mismatch calibration is not required, but the power consumption of a binary-to-thermometer encoder will similarly increase the fixed overhead power of the ADC and limit its efficient scaling range.

Settling error The dynamic settling error caused by the finite resistance of switches toggling between the DAC reference voltage and V_{SS} must also be kept sufficiently small. To prevent systematic comparator decision errors, the bottom plate of the DAC must settle to within an acceptable fraction of the final value (smaller than the target LSB size) before the next comparator decision is initiated. The switch resistance, R_{sw}, required to settle the

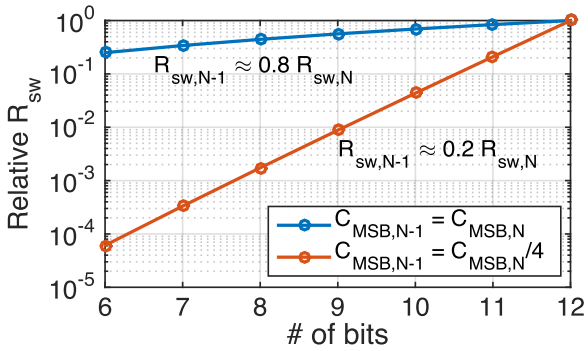

Figure 3.7. Scalability of required DAC switch resistance (R_{sw}) as the target ADC resolution scales. The relative R_{sw} values are presented relative to R_{sw} required by a 12-bit ADC.

MSB capacitance (C_{MSB}) within half an LSB in a settling time of t_{settle} is given by:

$$R_{sw} = \frac{t_{settle}}{C_{MSB}(N+1)\ln(2)}$$

In a synchronous ADC design, $t_{settle} \approx ((N+1)f_{s,ADC})^{-1}$. Using the above equation, reducing the required ADC resolution by a single bit allows scaling R_{sw} by a factor of $\frac{C_{MSB,N-1}(N)^2}{C_{MSB,N}(N+1)^2}$, where $C_{MSB,N}$ represents the MSB capacitance of an N-bit ADC. If the DAC capacitance does not scale with resolution such that $C_{MSB,N-1} = C_{MSB,N}$, Fig. 3.7 shows that the resolution scaling leads to a modest reduction in R_{sw} when N is large that improves as N scales. If C_{MSB} is reduced by a factor of 4 per bit to keep the relative contribution of DAC thermal noise constant as the ADC resolution scales, R_{sw} can be reduced by roughly a factor of 5 per bit.

Assuming the switch transistors have minimum length and threshold voltage and the gate voltage of the switches is maximized, R_{sw} can only be lowered by increasing the effective width (W) of the sampling switch. Because the gate capacitance of the switch scales with W, increasing R_{sw} proportionally lowers the dynamic power required to drive these switches. As illustrated in Fig. 3.7, the minimum required R_{sw} scales sharply with the number of conversion bits. Practically, the switch transistor cannot be made smaller than a minimum-sized device, so at low resolutions the dynamic power used to drive the switch will be larger than necessary. Moreover, the leakage current of wide, low-threshold switch devices used to minimize R_{sw} for high resolution operation will draw a fixed amount of overhead power.

3.2.2.2 *Power consumption*

The total power consumption of a capacitive DAC consists of two main sources: dynamic power consumption due to DAC capacitance switching, and a fixed amount of power overhead due to switch drivers and any peripheral control logic in the DAC.

DAC switching procedure The energy used to switch the DAC array between reference voltages while executing the conversion algorithm is proportional to $C_{DAC}V_{ref}^2$, where C_{DAC} is the capacitance being switched and V_{ref} is the DAC's reference voltage. Because the DAC power consumption is directly proportional to the total DAC capacitance, increasing C_{DAC} to reduce sampled thermal noise or mismatch comes at the expense of power. Specifically, increasing the total thermal noise-limited resolution by a single bit (the equivalent of 6 dB

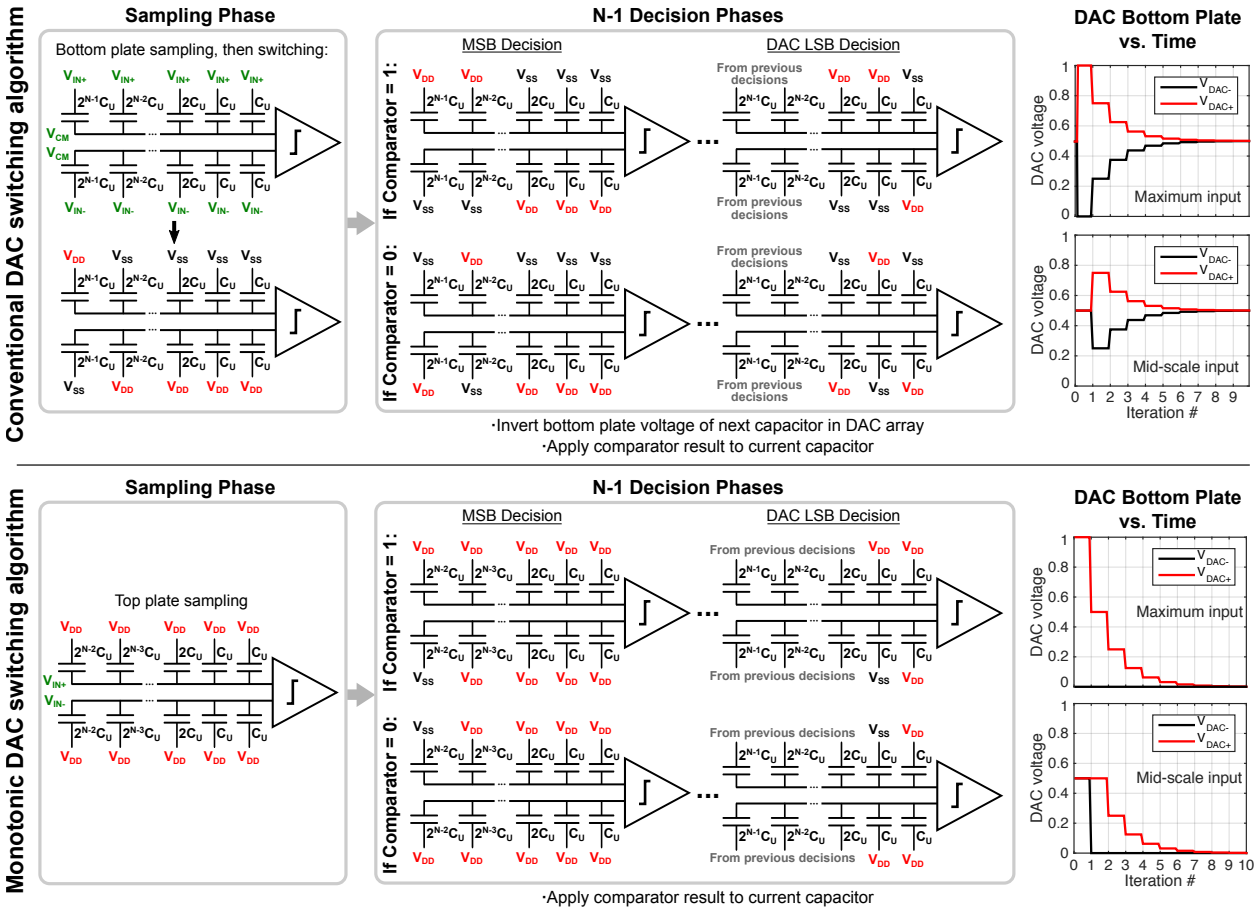

Figure 3.8. Comparison of "conventional" and "monotonic" DAC switching procedures. The conventional scheme maintains a constant DAC common-mode vs. iteration at the expense of higher power and switching complexity.

increase in signal-to-noise ratio) will increase the required DAC capacitance (and therefore DAC switching energy) by a factor of four. In addition, the precise relationship between $C_{DAC}V_{REF}^2$ and the actual DAC switching energy is strongly dependent upon both the input voltage and the switching procedure used to perform the conversion.

Implementing the successive approximation algorithm requires iteratively subtracting fixed fractions of a reference voltage from the sampled input signal in response to a comparator decision. Various input sampling techniques, capacitor weights, reference voltages, and switching procedures can be chosen to implement a valid conversion algorithm. Each method has varying levels of implementation complexity, performance impact, and power consumption. Figure 3.8 compares the "conventional" switching procedure to the monotonic procedure presented in [71]. The conventional procedure samples the input signal via bottom-plate sampling, and then requires both sides of the differential DAC to be switched in each conversion iteration, maintaining a constant DAC common-mode voltage. In the simpler monotonic procedure, the input signal is sampled onto the top plate of the DAC capacitors, and only a single side of the DAC must switch at each conversion step. As detailed in [71], this monotonic switching procedure can reduce the DAC switching energy by 80%.

While the monotonic switching scheme can significantly reduce the DAC switching energy, it causes the common-mode voltage input to the comparator to vary after each conversion cycle. Because various aspects of the comparator's performance — namely input-referred offset and noise — often depend upon this common-mode bias point, a monotonic switching procedure can degrade ADC performance if this effect is not considered. Many alternative low-power switching schemes have been proposed, including an energy-saving approach that maintains constant V_{CM} while reducing power relative to the conventional procedure by 56% [72], a charge-recovery scheme using V_{CM} as a reference that reduces switching energy by 87% [73], and a single-ended switching procedure using V_{CM} as a reference with the potential to reduce switching energy by 95% [74]. Each of these techniques presents tradeoffs between performance impact, system-level complexity (e.g. reference voltage generation), and layout simplicity.

Fixed overhead Some aspects of DAC power consumption are not scalable. A fixed amount of overhead power will be drawn by any peripheral logic blocks, such as binary-to-thermometer encoders used to mitigate DNL degradation due to component mismatch. Digital logic blocks are typically constructed from minimum-sized standard cells, so the power consumption of these components can only be reduced if the entire cell is disabled.

A final non-scalable aspect of DAC overhead power is leakage current. While the leakage of peripheral logic blocks will be small, the large, low-threshold transistors used to minimize the resistance of the MSB capacitor drivers will draw a large amount of non-scalable leakage current. The precise impact of leakage power will depend on the conversion speed of the ADC. In very low speed (kS/s) ADCs with low dynamic power, leakage may contribute a substantial fraction of the total ADC power. For moderate-speed converters suitable for wireless applications, the leakage power of the DAC drivers can consume over 30% of the total DAC power in the low-resolution ADC configuration (as discussed later in Fig. 3.18).

Scalable overhead The dynamic power needed to drive the gate capacitance of the DAC sampling switch and bottom-plate switches also contributes to the total DAC power, but scales with the target resolution of the ADC. Because the minimum size of these switches is a function of the DAC capacitance and required conversion speed, their size can be scaled with the DAC capacitance (with some additional overhead power for control gates to modify the component sizes).

3.2.3 Logic

The digital logic in a SAR ADC generates the control voltages needed to implement the chosen DAC switching procedure. It typically includes a shift register to track the current conversion bit, an array of registers to retain each DAC control setting, and an additional set of registers to store the final ADC decision. The clock can be generated either synchronously or asynchronously. In synchronous operation, an external control clock operating at Nf_{samp} is used to drive the ADC logic, where N is the number of conversion bits and f_{samp} is the ADC sample rate. This allocates a fixed amount of time to the comparator decision and DAC settling for each bit. Alternatively, in an asynchronous SAR ADC, the logic is driven by a "done" signal generated by the comparator until the N decision cycles are complete. While asynchronous operation requires additional control logic, it can support faster sample rates because the ADC speed is not limited by the worst-case comparator delay and DAC

settling time. While the SAR logic will not introduce anything beyond quantization noise that degrades the resolution of the ADC, it consumes a fixed amount of overhead power that fundamentally limits the power and resolution scalability of the converter.

3.2.3.1 *Sources of noise and distortion*

Because an ADC converts a continuous signal into discrete levels, any ideal ADC will contribute a fixed amount of error (quantization noise) when approximating an analog signal. This noise is purely a function of the number of digital bits used to convert an analog signal to the digital domain, and results in a signal to quantization noise ratio (SQNR) that can be calculated as follows (assuming a full-scale input signal):

$$SQNR_{dB} = 6.02N_{bits} + 1.76$$

This indicates that adding a bit of resolution to the digital logic will improve SQNR by roughly 6 dB. Because the logic is implemented in the digital domain, it will not introduce more noise or distortion to the converter. However, the logic does introduce a fixed delay that limits the achievable conversion speed. For moderate sampling rate wireless baseband applications, however, this is not a major limitation of the ADC's performance.

3.2.3.2 *Power consumption*

If the digital logic switching energy per iteration of the SAR algorithm is fixed, the total logic power will scale linearly with N_{bits}, which determines the number of algorithm iterations. Since the digital logic switching energy may depend linearly on the number of bits as well (e.g. a typical shift register contains N_{bits} elements), the logic power can scale with N_{bits}^2. As a result, if the resolution requirements on the ADC are relaxed by a single bit and the converter is quantization noise limited, the power consumption of the digital logic will shrink by at most a factor of $\frac{N_{bits}^2}{(N_{bits}+1)^2}$.

The Walden figure of merit (FoM) for an ADC is defined as the total ADC power (P) divided by the product of the converter's Nyquist sampling rate (f_s) and number of effective conversion levels (2^{ENOB}, where $ENOB = (SNDR - 1.76)/6.02$). To maintain constant FoM as the converter size scales, the ADC power must be cut in half for each bit reduction in resolution. However, this is impossible to achieve when the converter is quantization noise-limited. To illustrate, the equation below expresses the ratio between the FoM of an N bit ADC ($FoM_{w,N}$) and an $N-1$ bit ADC ($FoM_{w,N-1}$) if both are quantization limited and logic power dominates, which is possible if N is small. E_{logic} represents the energy consumed each time the logic switches:

$$\frac{FoM_{w,N-1}}{FoM_{w,N}} = \frac{E_{logic}(N-1)^2 f_s}{f_s \times 2^{N-1}} \times \frac{f_s \times 2^N}{E_{logic} N^2 f_s} = 2\frac{(N-1)^2}{N^2}$$

Fig. 3.9 summarizes how quantization noise-limited FoM degrades (increases) when N_{bits} is reduced. If logic power varies quadratically with N_{bits}, FoM increases by up to a factor of 4.5 relative to that of an 8-bit converter, while it increases by up to a factor of 16 if logic power scales linearly with N_{bits}. For $N_{bits} \geq 8$, ADC resolution is typically limited by thermal noise and distortion, rather than quantization. At low resolutions, however, the digital logic power can greatly limit the efficiency of a scalable converter, indicating that the total digital logic power must be minimized in a resolution-scalable design.

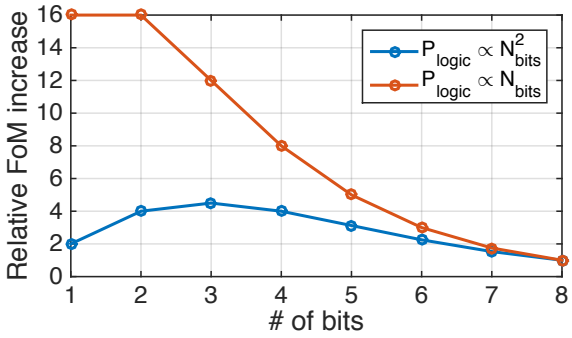

Figure 3.9. Degradation in Walden FoM (fJ/step) relative to 8-bit ADC due to logic power scaling. The blue and orange curves assume that logic power varies quadratically and linearly (respectively) with the number of ADC bits.

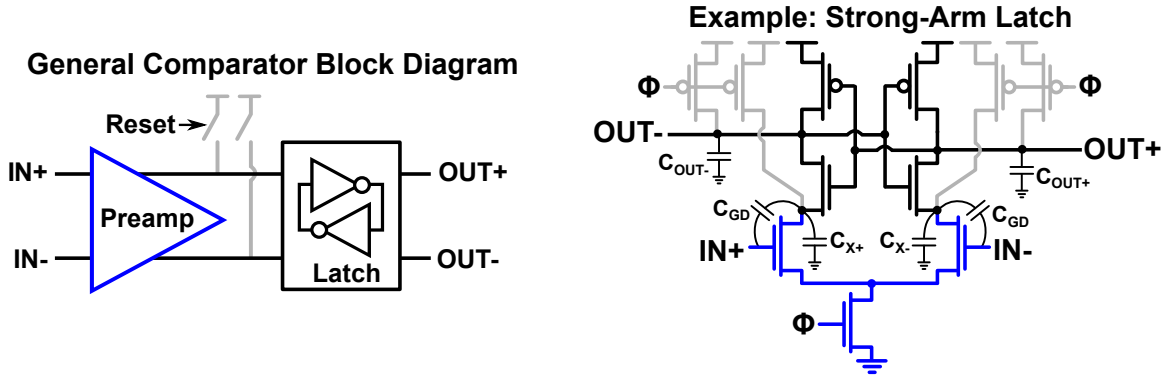

Figure 3.10. General comparator block diagram (left) and diagram of strong-arm latch comparator topology with some major parasitic capacitances marked (right).

3.2.4 Comparator

The comparator must accurately determine the sign of the DAC voltage. While specific implementation details may vary, comparators for SAR ADCs are typically implemented as a clocked circuit that feeds current from a preamplification stage into a regenerative latch built from cross-coupled inverters. Positive feedback in the latch amplifies the input signal to generate a digital output voltage that is usually reset during one phase of the sampling clock. Figure 3.10 presents a schematic diagram of a common comparator topology, the strong-arm latch, along with a corresponding functional block diagram.

3.2.4.1 *Sources of noise and distortion*

Many comparator non-idealities restrict the resolution of a SAR ADC, including input-referred thermal noise and offset, kickback noise, and metastability. In deeply scaled CMOS process nodes with low supply voltages, the input-referred noise of the comparator can significantly limit the ADC's achievable resolution. If the comparator decisions are inaccurate, the DAC will subtract an incorrect voltage from the input sample.

Input-referred thermal noise The input-referred thermal noise of the comparator appears in series with the ADC input voltage, introducing a fixed amount of noise that limits

the converter's SNDR, particularly if the ADC's input swing is limited. The nonlinear, time-varying nature of comparator circuits complicates noise analysis, but the comparator operation can be separated into linearized phases as described in [75] to quantitatively analyze its noise. The work in [75] derives a model for the input-referred noise of the strong arm latch shown in Fig. 3.10 that can be generalized to clocked comparator topologies consisting of a preamplifier and latch. Overall, the analysis indicates that the main design parameters contributing to comparator noise are the transconductance (g_m) of the preamplifier input devices, g_m of the NMOS transistors in the latch, parasitic node capacitances (C_X and C_O in Fig. 3.10), and duration of the preamplification and latching initiation phases. Generally speaking, increasing g_m, parasitic capacitances, and decision times will all improve input-referred noise. Specifically, the result from [75] expresses input-referred noise using a square-law model for device operation as:

$$\begin{aligned}
\mathcal{F} &= \frac{V_{TH,NLAT}}{V_{ov,IN}} = \frac{g_{m,IN} t_{preamp}}{2C_X} \\
\mathcal{H} &= \frac{V_{DD} - V_{CM}}{V_{ov,NLAT}} \frac{I_{D,NLAT}}{I_{D,IN} - I_{D,NLAT}} = \frac{g_{m,NLAT} t_{pre_regen}}{2C_X} \\
v_{n,in}^2 &= \frac{kT}{C_X} \left(\frac{2\gamma}{\mathcal{F}} + \frac{1 + C_X/C_O}{2\mathcal{F}^2} + \frac{4(1+\gamma)}{8\mathcal{F}^2\mathcal{H}} + \frac{(1+\gamma)C_O/C_X}{8\mathcal{F}^2\mathcal{H}^2} \right)
\end{aligned}$$

In the above expression, $V_{TH,NLAT}$, $g_{m,NLAT}$, $I_{D,NLAT}$, and $V_{ov,NLAT}$ are parameters for the NMOS devices in the latch; $g_{m,IN}$, $I_{D,IN}$, and $V_{ov,IN}$ are parameters for the input devices in the differential pair; t_{preamp} is the duration of the preamplification phase and t_{pre_regen} is the amount of time the NMOS transistors in the latch are active before the PMOS devices have turned on. Because the input-referred voltage noise of the differential pair is inversely proportional to both the g_m of the input devices and the integration time set by C_X and the differential pair drain current (I_D), comparator noise can be lowered using wide devices with a high g_m/I_D. Similarly, designing the NMOS latch devices to have large g_m/I_D will improve the total integrated noise, though by a smaller amount than the input devices. Because t_{preamp} and t_{pre_regen} set the effective noise bandwidth of the comparator, a direct tradeoff exists between the comparator's input-referred noise and decision time. Reducing the overdrive voltage ($2I_D/g_m$) of the comparator by lowering its common-mode input voltage can also improve noise, though at the expense of decision speed.

To scale the input-referred noise of the comparator without degrading the comparator's decision speed, the g_m of the input devices could be scaled by tuning the effective width of the input devices, assuming the tail transistor in the differential pair is small enough to fix I_D. While tuning C_X as in [60] and [64] will scale comparator noise well for low-speed designs that are not restricted by comparator speed, this approach may not scale to wireless baseband applications requiring signal bandwidths of up to 80 MHz. Because t_{preamp} and t_{pre_regen} scale linearly with C_X, the effect of C_X on $\mathcal{F}$ and $\mathcal{H}$ cancel such that $v_{n,in}^2 \propto kT/C_X$. As a result, increasing the ADC's required SNDR by 6 dB requires C_X to quadruple. If C_X is increased to reduce noise, I_D must increase correspondingly to maintain the same speed. If $\mathcal{F}$ and $\mathcal{H}$ are large, $v_{n,in}^2 \approx \frac{2\gamma kT}{C_X \mathcal{F}}$. While $\mathcal{F}$ may not be significantly larger than 1 (especially in modern processes with low V_{TH}), it will be the dominant factor in setting comparator noise. Therefore, reducing the required ADC resolution by 6 dB requires the g_m of the input devices

to quadruple. Physically, this can be accomplished using parallel input devices that can be individually enabled, though the parasitic capacitance of the enable switches will increase C_X and degrade the comparator's decision speed. To minimize noise, the comparator's decision time should be made as slow as the target ADC sampling speed will allow.

Input-referred offset Mismatch between the threshold voltage of input devices in the differential pair and regenerative latch, as well as mismatch between the capacitive loading at the output of the comparator, will create a systematic input-referred offset to the comparator. In a SAR ADC using one comparator, this restricts the full-scale range of the ADC to $V_{max} - 2V_{offset}$, where V_{max} is the nominal maximum input swing of the ADC. For wireless applications where information is typically encoded in the frequency domain, this fixed offset can be removed digitally and does not degrade the ADC's performance other than reducing its input swing. Comparator offset is also a strong function of the common-mode input voltage to the comparator (V_{CM}). If a DAC switching procedure is used that does not maintain constant V_{CM}, offset mitigation and calibration techniques will be required to avoid degrading the DNL of the ADC.

Similar to DAC mismatch, comparator offset can be modeled as a zero-mean Gaussian random variable. Threshold voltage variance will be inversely proportional to device area, just as capacitor variance will be inversely proportional to capacitance. While increasing device area decreases the likelihood that offset due to the presence of random device mismatch will be problematic, explicit offset calibration techniques can eliminate offset without requiring large devices that typically consume more power. From a scalability perspective, these calibration techniques will require a fixed amount of power overhead. Careful layout techniques will help minimize the amount of systematic offset in a comparator, but the relatively small size of devices used in comparators can introduce a large degree of random mismatch that must be calibrated.

Kickback noise Kickback noise is introduced when the large output swing of the comparator capacitively couples onto the comparator's input through the gate to drain capacitance (C_{GD}) of the input devices in the differential pair. For any comparator topology in which one side of the differential pair output is pulled high and the other is pulled low, kickback noise will be signal dependent and should be minimized for high-resolution ADC operation. This is true of the strong-arm latch shown in Fig. 3.10. As a result, the magnitude of kickback noise will depend on both the swing at internal nodes of the comparator, dictated by its topology, and the ratio between the C_{GD} of the comparator's input devices, which are typically made large to reduce thermal noise, and the capacitive load of the DAC. Generally speaking, kickback noise can be modeled as:

$$\Delta V_{kickback} = \frac{C_{GD}}{C_{GD} + C_{DAC}} \Delta V_{swing}$$

From a noise scaling perspective, this would indicate that to reduce the relative contribution of kickback noise, the fraction $\frac{C_{GD}}{C_{GD}+C_{DAC}}$ must scale by a factor of 2 for every 6 dB change in ADC resolution. However, both C_{DAC} and C_{GD} (directly proportional to the size of the comparator's input devices) can be scaled by a factor of 4 per 6 dB change in ADC resolution. This means if kickback noise is significant, its relative contribution to the total noise of the ADC cannot be made scalable unless C_{GD} and C_{DAC} scale differently. Therefore,

it is important to select a comparator topology that minimizes ΔV_{swing}.

If thermal noise from the comparator is significant, a preamplifier consuming static power can limit ΔV_{swing}. Alternatively, comparator topologies using a dynamic preamplifier with a consistent rail-to-rail output swing may be added to reduce the impact of kickback noise. As discussed in [76], additional kickback reduction techniques include sampling the input voltage to the comparator through a set of switches that provide an additional degree of isolation from the large output swing of the comparator, or architectural modifications designed to cancel the kickback effect.

Metastability Because the comparator's gain and decision time are finite, there exists a nonzero probability that the input signal will be too small for the comparator decision to resolve fully within the target clock period. To achieve high gain, the comparator relies on positive feedback to rapidly amplify its differential input to full scale. The minimum tolerable comparator input voltage is a function of the preamplifier gain, A_{preamp}; the regeneration time constant of the latch, dictated by the effective $-1/g_m$ resistance of the cross-coupled inverters and the capacitive load it must drive (C_O in Fig. 3.10); and the maximum allowable comparator decision time, t_{comp}. The minimum non-metastable input voltage V_{min} is simply the desired full-scale output of V_{DD} divided by the overall comparator gain:

$$V_{min} = \frac{V_{DD}}{A_{preamp} \exp\left(\frac{t_{comp} g_m}{C_O}\right)}$$

The likelihood of comparator metastability is then given by $P\left(V_{in,comp} < V_{min}\right)$. Degradation in ADC performance due to metastability can be characterized as a signal to metastability ratio (SMR) that relates the signal power to the size of the metastability window set by V_{min} [77,78]. Assuming a fixed N-bit ADC conversion time of t_{conv}, the SMR of a synchronous ADC can be expressed as [78]:

$$SMR = \frac{1}{\ln(10)} \frac{t_{conv} g_m}{C_O} + 10 \log_{10}\left(\frac{2}{3} \frac{1}{1 - 2^{-N}}\right)$$

The above expression assumes the comparator input voltage ($V_{in,comp}$) is uniformly distributed across the full-scale input range of the ADC, and assumes that most of the comparator gain is obtained through the latch. As discussed in [78], Gaussian or Laplacian distributions are more consistent with multi-tone sinusoidal inputs similar to received wireless signals. Assuming this distribution of input signals can degrade the SMR 18 dB further.

From a scalability perspective, relaxing the required comparator SMR by 6 dB allows g_m/C_O to decrease by a factor of 4 if the ADC conversion time is fixed. To lower power consumption, g_m could be reduced using smaller devices in the latch. Practically speaking, if the regenerative latch is already built from minimum-sized transistors, this result mostly indicates that the likelihood of metastability reduces exponentially as the required ADC resolution scales.

For moderate speed ADCs designed for wireless applications, packets may be re-transmitted if metastability triggers an incorrect comparator decision that significantly degrades the resolution of the received packet. If packet retransmission occurs infrequently, this will not

substantially increase the overall radio power. The ADC sample rate required for Nyquist-rate (or near-Nyquist) wireless baseband applications is typically low enough that metastability is not a large concern. However, high-speed (GS/s) ADC designs that require high reliability and resolution may use metastability detection circuitry to avoid using high-power comparators with large g_m to drive a particular output load capacitance, C_O. Nevertheless, this detection circuitry would contribute a fixed amount of power overhead and increase C_O.

3.2.4.2 *Power consumption*

The overall power consumption of a comparator is largely implementation-specific. If no active preamplifier is used, the comparator power will be predominantly dynamic, dictated by the parasitic capacitance of devices within the comparator C_{comp}, the supply voltage V_{DD}, the number of comparisons N and the ADC sampling frequency $f_{s,ADC}$:

$$P_{comp} = C_{comp} V_{DD}^2 N f_{s,ADC}$$

Given that the number of comparator decisions N is typically equivalent to the number of ADC bits (N_{bits}), P_{comp} will reduce linearly with N_{bits} if C_{comp} is fixed. If C_{comp} is made tunable to scale comparator noise with the required ADC resolution, P_{comp} can be reduced by roughly a factor of 4 per bit as the resolution requirements of the ADC are relaxed. Relaxing the comparator's noise requirement would allow the preamplifier's input devices and the latch devices to have a lower g_m, and therefore proportionally lower comparator power. However, the overhead of switches to enable comparator configurability can introduce additional parasitics that fundamentally limit its minimum power consumption, as will be discussed in Section 3.3.4. Moreover, some comparator non-idealities such as offset may require additional calibration circuitry. In a resolution-scalable ADC targeting energy efficient MIMO systems, the overhead power of any calibration circuitry restricts the minimum power of the ADC, which will limit its overall power efficiency in large arrays.

3.2.5 Considerations for scalable SAR design

Major considerations in designing the capacitive DAC, digital logic, and comparator for a resolution-scalable SAR ADC are summarized as follows:

1. **Capacitive DAC:** In high-resolution SAR ADCs, a large DAC capacitance (C_{DAC}) is required to minimize sampling noise and nonlinearity induced by DAC component mismatch. However, C_{DAC} is directly proportional to the converter's power consumption, resulting in a direct tradeoff between DAC resolution and power. Overall, C_{DAC} can be reduced by a factor of four when the ADC resolution requirements are relaxed by a single bit. To keep sampling switch distortion and settling time errors low when C_{DAC} is large, however, the DAC must use small switch resistances (large switch sizes) that may increase the overhead power of the DAC (e.g. via leakage current). This may degrade the converter's efficiency for low-resolution operation.

2. **Digital logic:** The number of conversion bits in the digital logic can be reduced when a higher degree of quantization noise is tolerable. However, the power consumed by the digital logic will not scale exponentially with the number of ADC bits, which

fundamentally limits the low-resolution FoM of a scalable ADC. The logic power consumption must therefore be minimized to improve the low-resolution efficiency of the converter.

3. **Comparator:** High-resolution ADCs require low input-referred comparator noise, which typically requires large input devices or large parasitic capacitances to reduce the comparator's noise bandwidth. In turn, low-noise comparators often consume more power. To trade switching energy for noise, the size of the input devices can be tuned or the internal capacitances within the comparator can be scaled. Additional comparator non-idealities include offset, kickback noise, and metastability. While metastability may scale with device sizing, offset must typically be improved via calibration and kickback noise should be reduced using the comparator topology. Alternatively, the comparator can be biased to reduce the overdrive voltage of the input pair, which trades reduced noise for slower conversion speed, but does not significantly impact dynamic power consumption.

The next sections describe a 65nm CMOS implementation of a resolution-scalable SAR ADC. This implementation includes a scalable DAC capacitance of 250 fF to 1 pF, 6-11 conversion bits, 8 unit comparators that can be individually enabled and disabled to tune noise and offset, and low-power single-sided DAC switching procedure that facilitates tuning the comparator common-mode input voltage.

3.3 Prototype implementation

A prototype converter was designed to examine the feasibility of building a resolution- and power-scalable ADC suited to massive MIMO applications. Due to the noise averaging properties of the array, the maximum resolution of the converter could be restricted to 9 effective bits. Additionally, a moderate sampling rate of 80 MS/s is sufficient to support baseband signal bandwidths of 10-20 MHz with a slight degree of oversampling to relax the anti-aliasing filter requirements of the receiver. Figure 3.11 illustrates the full scalable receiver system [1], with the configurable mixer-first receiver and the configurable ADC from this work. The next chapter will discuss mixer-first receiver topologies in further detail. Resolution configurability is implemented using a capacitive SAR ADC with scalable DAC capacitance, comparator size, and number of conversion bits as shown in Fig. 3.12. A programmable single-sided DAC switching procedure is used to control the bias voltage of the comparator to regulate its offset and noise. This section reviews the design of four main components of the ADC — the capacitive DAC, digital logic, comparator, and DAC switching procedure.

3.3.1 Comparator bias tuning and DAC switching

Because the input-referred comparator noise and offset depends on the comparator's common-mode input voltage (V_{CM}), tuning this bias point provides one mechanism for configuring comparator performance. To adjust V_{CM}, a single-sided DAC switching procedure is used in which the digital logic applies the differential comparator result to a new bit of the DAC during each SAR conversion step. This enables V_{CM} tuning by ensuring that only one side of the differential DAC array switches at a time, as shown in Fig. 3.13. Using this approach,

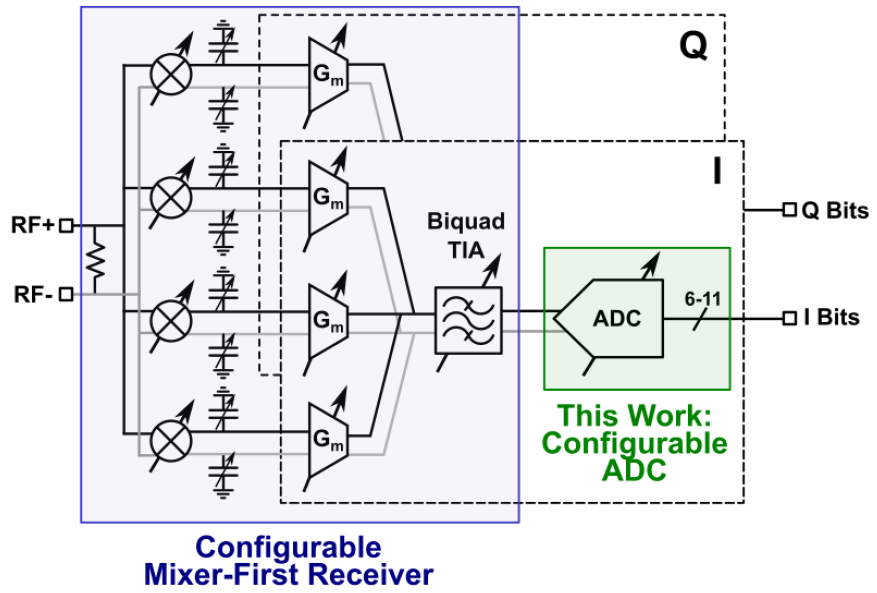

Figure 3.11. Diagram of prototype receiver + ADC, as presented in [1].

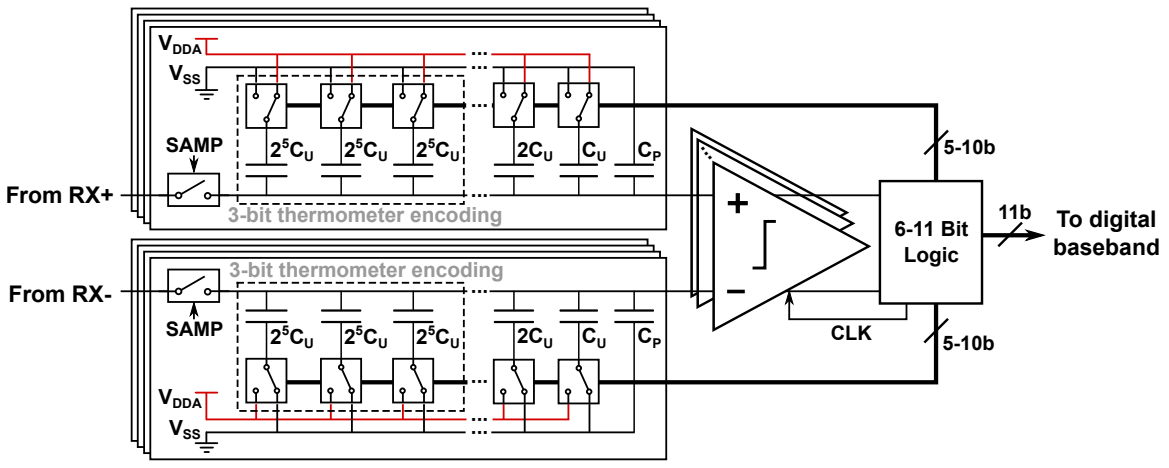

Figure 3.12. Scalable ADC design topology.

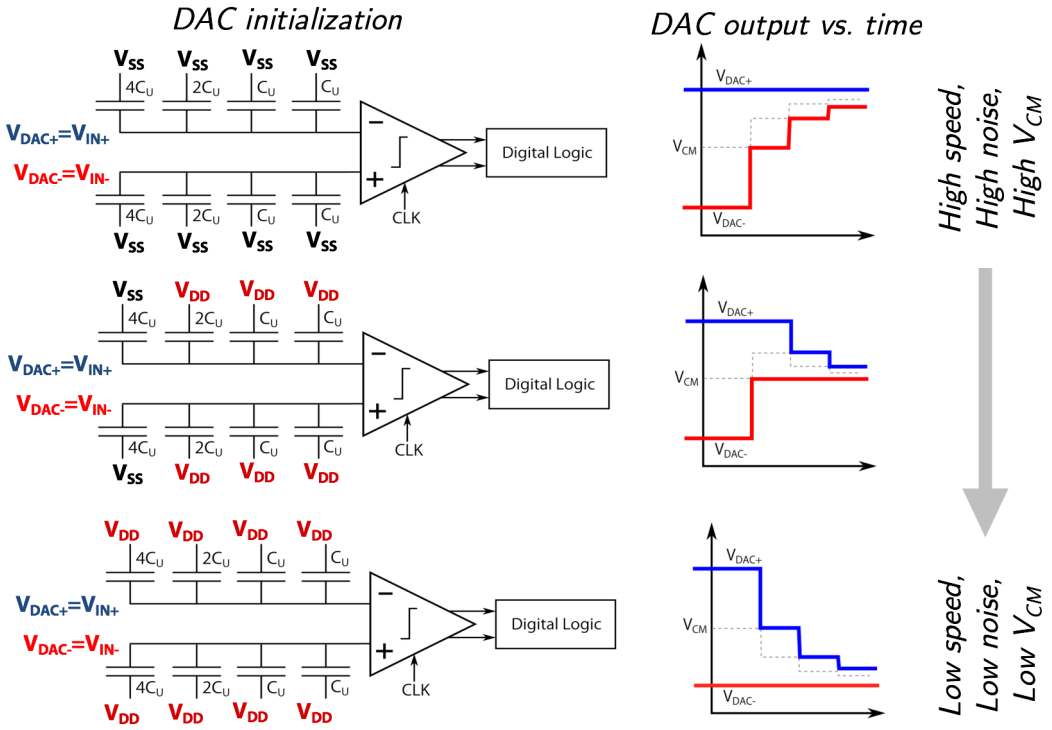

Figure 3.13. Illustration of programmable DAC switching procedure designed to set the comparator bias point.

the initial voltage applied to the bottom plates of capacitors in the DAC array dictates the final V_{CM}. For instance, if the bottom plates are all initialized to V_{SS}, one of the two bottom-plate voltages will rise to V_{DD} as the successive approximation algorithm proceeds, creating a positive voltage jump that shifts V_{CM} upwards. The high V_{CM} is suitable for high-speed yet high-noise ADC operation. To minimize noise, the DAC bottom plates can be initialized to V_{DD} in order to lower the final V_{CM}. This is only an option when timing margin is available, so V_{CM} should be set as low as possible to minimize noise while meeting the overall conversion speed requirements for a given technology. The initial DAC bottom-plate voltages can be modified as needed to generate a mid-range V_{CM}.

Because only one side of the differential DAC switches during each conversion step, this switching scheme is also highly energy-efficient. The work in [71] introduced the monotonically decreasing V_{CM} switching procedure as a means of reducing DAC switching energy. Figure 3.14 compares the switching energy of the DAC using a conventional fixed-V_{CM} technique to the switching energy of the tunable-V_{CM} technique. The switching energy is output code-dependent because the DAC bottom-plate switching sequence determines the total charge consumed by the DAC. Similarly, as the DAC reset code is tuned to vary V_{CM}, the average switching energy varies by 16% but enables more than a 3x reduction in DAC switching energy overall. This is lower than the 5x improvement reported in [71] because the DAC prototype incorporates a large parasitic capacitance in parallel with the switching DAC capacitance to restrict the full-scale input range of the ADC. Because the radio receiver driving the ADC input does not have a rail-to-rail output swing, the parasitic capacitance attenuates the DAC voltage change when the bottom-plate voltage swings by V_{DD}. This

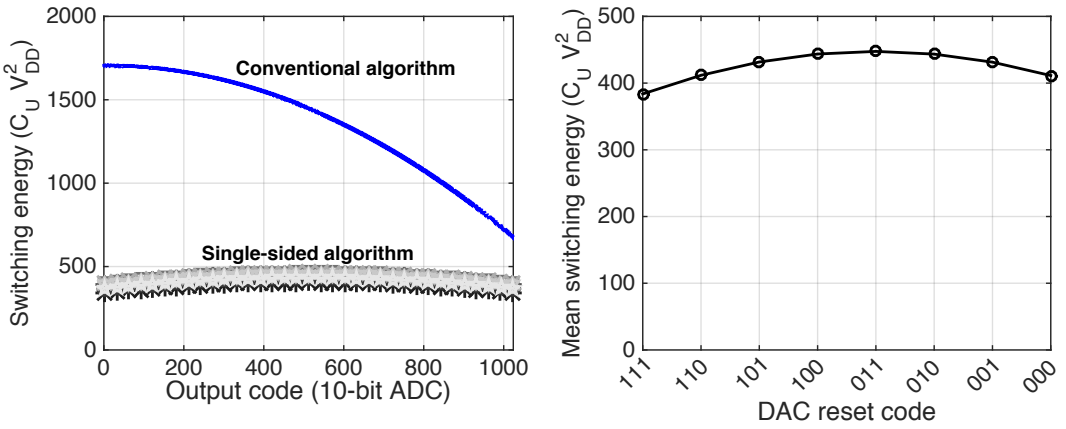

Figure 3.14. Simulated DAC switching energy for programmable switching procedure.

allows the full-scale input range of the ADC to align with the full-scale output swing of the receiver while obviating the need for an additional mid-rail reference voltage.

One key disadvantage of the tunable V_{CM} switching procedure is the dependence of comparator offset on V_{CM}. When the common-mode input to the comparator changes at each ADC conversion step, fluctuation in comparator offset can translate to missing codes within the ADC, particularly near MSB transitions. As a result, the comparator offset must be designed to have a low sensitivity to V_{CM}. In this prototype, redundancy in the comparator design is used to greatly reduce V_{CM} sensitivity; more details are provided in Section 3.3.4.

Previous implementations of single-sided DAC switching schemes use a monotonically decreasing configuration in which V_{CM} approaches ground at each conversion step [71, 79]. Because many stand-alone ADCs are designed to support a rail-to-rail input swing, this lowers V_{CM} to V_{SS} in the final iterations of the SAR algorithm. To prevent biasing the comparator's input devices in the subthreshold regime as the algorithm progresses, a PMOS-input comparator is typically used. However, this causes the overdrive voltage of the input devices to increase as the algorithm progresses, which raises the input-referred comparator noise and offset when the comparator input is small. Because the V_{CM} is programmable in this implementation, an NMOS-input comparator can be used without sacrificing speed. Moreover, the large parasitic capacitance used to restrict the input range of the ADC similarly limits the V_{CM} swing of the DAC, making it possible to use an NMOS-input comparator with the monotonically decreasing switching procedure.

3.3.2 Capacitive DAC

3.3.2.1 *Design*

To trade sampling noise and distortion for power consumption, the total DAC capacitance can be scaled with the number of conversion bits. As the analysis in Section 3.2.2 demonstrates, relaxing the resolution requirements on the ADC by a single bit allows the DAC capacitance to be cut by a factor of 4. This increases the variance of sampled thermal noise by the allowable 6 dB and keeps the worst-case DNL standard deviation (σ_{DNL}) constant. Reducing the DAC capacitance by a factor of 4 for a single bit reduction in ADC resolution

would require eliminating the two MSB capacitors and introducing a new least significant bit (LSB) capacitor with half the size of the previous one. This increases the unit element mismatch, but also requires fewer switched elements, keeping σ_{DNL} constant. However, introducing a new LSB capacitance each time the ADC resolution changes presents practical layout challenges because the DAC should be constructed from uniform unit capacitors to maintain a high degree of matching. To simplify the physical implementation of the DAC, it is constructed as four identical 9-bit, 250 fF (single-ended) sub-DACs, as shown in Fig. 3.15. The sub-DAC technique sacrifices the power savings in favor of implementation simplicity.

While using four sub-DACs enables only two bits of DAC capacitance scaling, the low-power DAC switching procedure ensures that capacitive DAC switching is not the dominant source of power consumption in this ADC. While a wider scaling range could have been implemented by disabling the most significant bit (MSB) capacitances to reduce the total DAC capacitance as in [60] and [64], the sub-DAC approach was used to minimize overhead power by eliminating the need for changes to the SAR conversion logic and digital baseband circuitry.

The 250 fF sub-DAC capacitance is approximately equally divided between actively switched capacitance and parasitic capacitance used to reduce the full-scale input range of the ADC. As discussed previously, including a large parasitic capacitance allows the DAC to use a reference voltage of V_{DD}, and limits the full-scale input range of the ADC to the linear output range of the RF receiver. Each DAC contains a fixed-size sampling switch, allowing the total switch resistance to scale with the number of unit DAC cells. The supply switches are implemented as inverters sized to drive the corresponding DAC capacitances. To reduce mismatch effects, the three MSBs of the capacitive DAC are thermometer-encoded, as illustrated in Fig. 3.15. This limits the worst-case σ_{DNL} by preventing the DAC from switching between a large number of unique unit cells as the SAR switching procedure progresses. The power overhead of the binary-to-thermometer encoder is kept low by using only three thermometer-encoded MSBs. Due to the sub-DAC capacitance scaling method, the three MSBs are thermometer-encoded even in the low-power 9-bit DAC setting.

Physically, each sub-DAC is constructed as an array of identical 250 aF (125 aF switched and 125 aF parasitic to AC ground) unit capacitors built in the top metal layers. Control signals are routed along the bottom of each unit cell, with ground shields surrounding the capacitor plates to minimize any possible coupling from the control traces to the DAC (as shown in Fig. 3.16). A common centroid layout is used to minimize the effects of systematic mismatch in the DAC. Four unit cells per sub-DAC can be controlled independently, allowing the three LSBs that are not contained within a single sub-DAC to be placed in a common centroid fashion as shown in Fig. 3.17.

3.3.2.2 *Expected performance*

The major sources of resolution degradation caused by the DAC are DNL and INL due to unit capacitor mismatch, sampled thermal noise, and sampling switch nonlinearity. The unit capacitance chosen for this DAC is 125 aF (250 aF total, incorporating parasitics to ground). In this process technology, the standard deviation of capacitor mismatch (estimated via Monte Carlo simulations of the provided metal-oxide-metal (MOM) capacitor cells), σ_u, can be estimated as $\frac{1}{230\sqrt{C}}$, where C is the unit cell capacitance in femtofarads. For this design, $\sigma_u = 0.0126$. Because the DNL variance scales directly with the number of

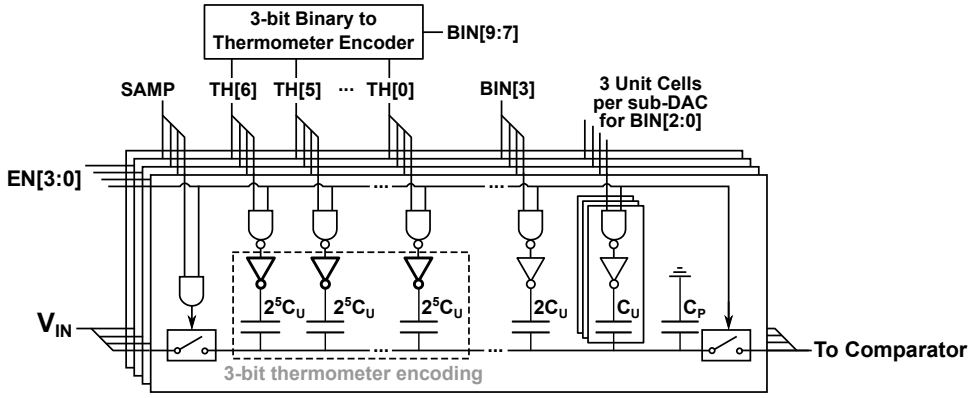

Figure 3.15. Schematic of DAC diagram.

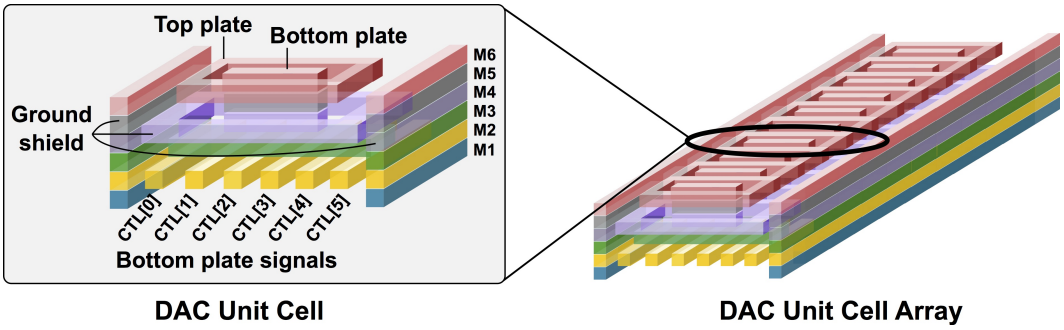

Figure 3.16. Diagram of DAC unit cell layout.

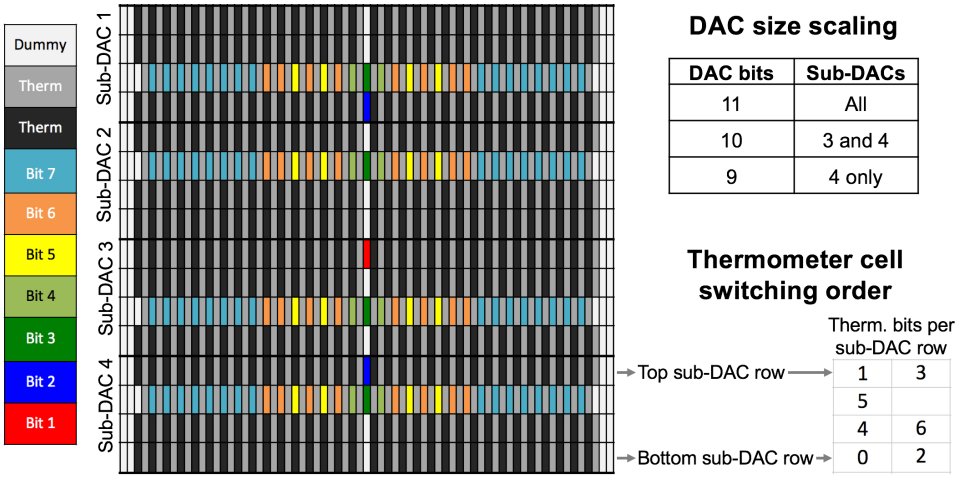

Figure 3.17. Diagram of common centroid DAC unit cell layout. The three unit cells in each sub-DAC are connected to bits 2-0 or dummy cells to maintain common centroid placement. The 32 thermometer cells for each thermometer bit are grouped in each row, and then switched in an order that also maintains a common centroid layout.

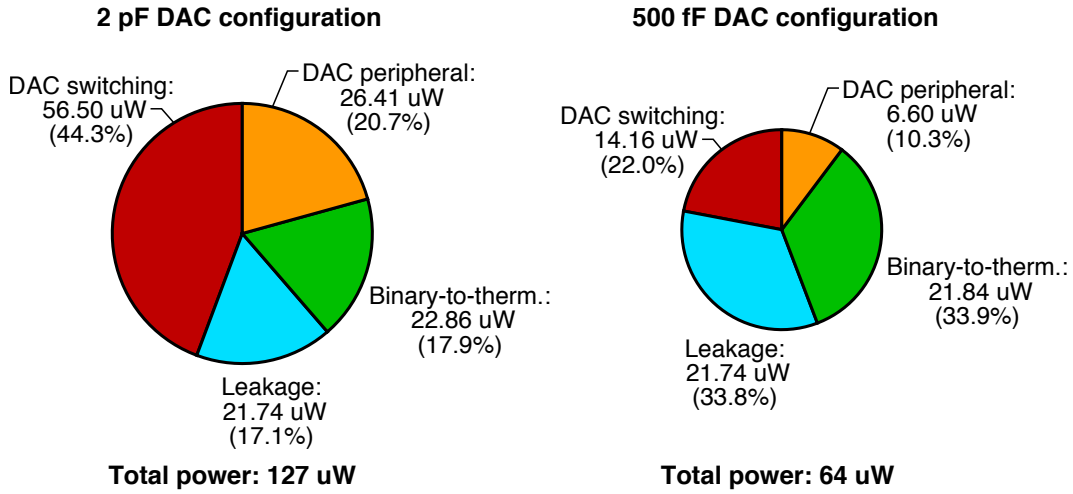

Figure 3.18. Simulated DAC power consumption breakdown in maximum-sized 2 pF (left) and minimum-sized 500 fF (right) settings.

switched unit elements, the expected σ_{DNL} of this design with 7 binary bits (N_{bin}) and 3 thermometer-encoded bits is given by $\sigma_u\sqrt{2^{N_{bin}+1}-1} = 0.20$ LSBs. The thermometer encoding does not improve the converter's integral nonlinearity (INL); its standard deviation (σ_{INL}) is given by $\sigma_u\sqrt{2^{N_{bits}-2}} = 0.20$ LSBs (N_{bits} is the number of total ADC bits). Overall, the mismatch performance should be sufficient to support the target 9 effective bits of resolution in the highest-resolution configuration. With a 1.2 V differential input amplitude, the thermal sampling noise-limited SNR of the converter is restricted to 73 dB, guaranteeing that sampling noise is not a significant concern for this DAC implementation.

While the resolution of the DAC is largely dictated by the total capacitance, which sets the unit element mismatch and sampling noise, the power consumption of the DAC is not entirely scalable. In addition to DAC switching energy, the total DAC power consumption will include the power consumption of the binary-to-thermometer encoder, enable logic, and the energy needed to drive the DAC switches. Because the on resistance of the switches must be sufficiently low to support the conversion speed of the ADC, they will also introduce leakage current that contributes to the static power consumption of the converter. Figure 3.18 summarizes the simulated breakdown of DAC power consumption. Overall, the expected power consumption of the DAC is 64 μW in the 9-bit setting and 127 μW in the 11-bit setting, but only 22% and 44% of this power (respectively) can be attributed to DAC switching energy. This indicates that the DAC consumes a fixed amount of overhead power that will restrict the efficiency of the ADC in low resolution configurations.

3.3.2.3 *Scalability impact*

The sub-DAC structure fundamentally limits the scalability of the DAC switching energy in low resolution configurations, which reduces the efficiency of the ADC by contributing a fixed amount of overhead power. However, as seen in Fig. 3.18, DAC components with scalable power consumption account for less than half of the total DAC power in the minimum-resolution (9-bit) DAC configuration. Most power is consumed by the binary-to-thermometer encoder and leakage current. Figure 3.18 demonstrates that in this low resolution setting,

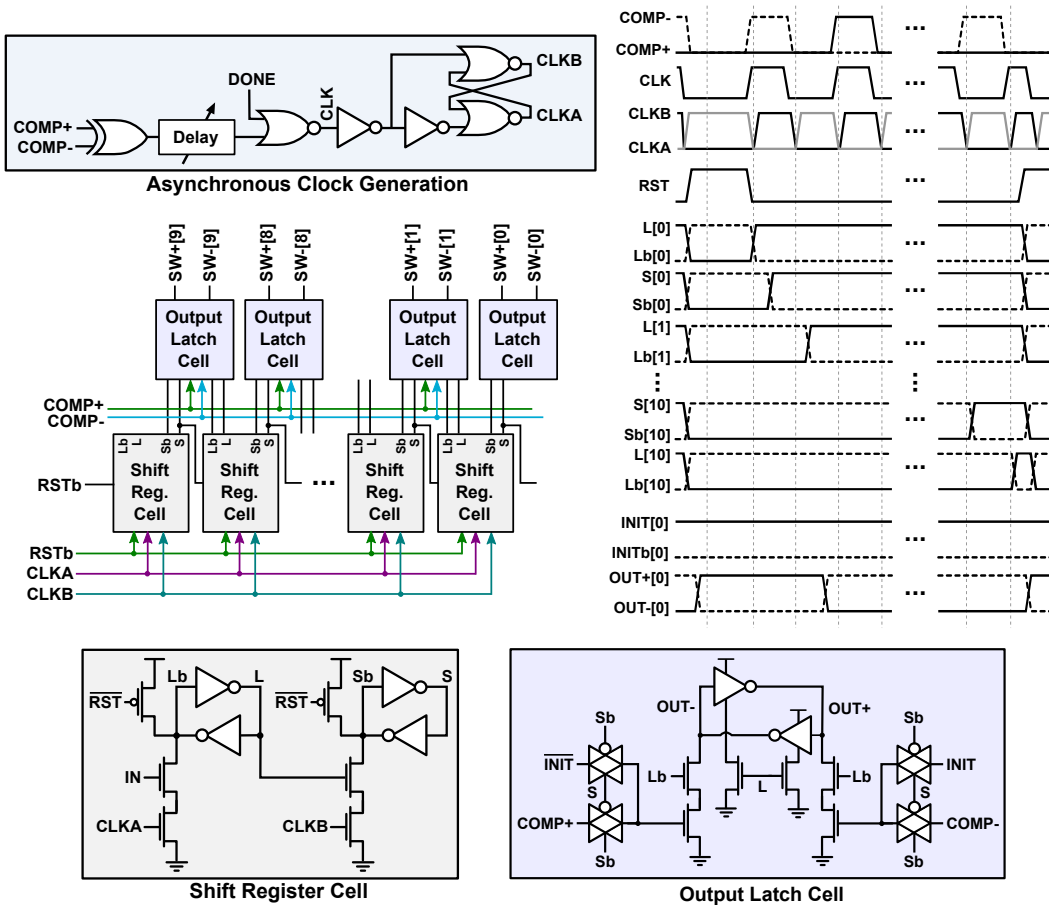

Figure 3.19. Schematic diagram of ADC logic and timing diagram of 11-bit logic operation.

the scalable DAC topology consumes 244% more power than a stand-alone 9-bit DAC that only consumes DAC switching power, peripheral power within the sub-DAC, and leakage scaled to the size of the DAC. Realistically, relaxed timing margin in a lower resolution ADC would enable using smaller DAC drivers with even lower leakage power. This overhead power penalty has a more significant impact on ADC efficiency in the lowest-resolution (6-bit) operating mode, considering that this particular DAC implementation scales down to only 9 physical bits from the maximum 11-bit setting.

3.3.3 Logic

3.3.3.1 *Design*

The logic must support a programmable DAC switching procedure while tuning the converter's resolution from 6 to 11 bits. Functionally, implementing the programmable DAC switching procedure simply requires applying the result of successive comparator decisions to the appropriate DAC control bit. At the end of each conversion, the DAC control bits are reset to a programmable code that configures the DAC switching procedure to set the V_{CM} of the comparator. To do this, the logic uses a shift register to track the current conversion bit and a set of custom multiplexers to latch the appropriate comparator result. The num-

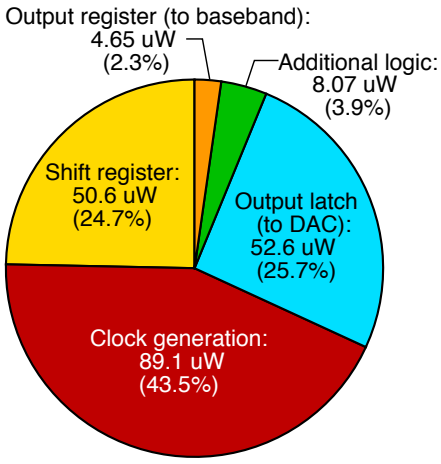

Figure 3.20. Power breakdown of total 11-bit logic power consumption.

ber of bits is adjusted by selecting which bit of the shift register marks the final conversion cycle. In this final cycle, the DAC bottom-plate voltages are reset and the sampling switch is closed until the next rising edge of the sampling clock. Figure 3.19 shows a general block diagram of the logic implementation that includes the shift register and custom level-shifting multiplexers that latch the comparator result and apply the required reset code. A custom logic implementation was used to lower power consumption and area, given that logic power scales only linearly with the number of conversion bits. Reducing logic power consumption helps reduce the amount of overhead power that limits the efficiency of the ADC at low resolutions.

The comparator clock is generated asynchronously by the digital logic using the circuit shown in Fig. 3.19. Since both outputs of the comparator are initially reset low when the comparator clock is low, the decision is complete when one of the two outputs is pulled high. Therefore, a DONE signal can be generated by taking the XOR of both comparator outputs. This can then be inverted to generate the comparator clock. An adjustable delay is included in the clock generation path to incorporate margin for the logic delay and DAC settling time.

Because the comparator and DAC operate on the analog supply while the digital baseband circuitry operates on the lower digital supply, the logic also requires level shifters to translate between the two voltage levels. Figure 3.19 shows the schematic-level implementation of the output latch cell that incorporates level shifting. The four control signals from the custom shift register cells are used to select between the programmable initial voltage (INIT), comparator output (COMP), or to latch the previous decision. This 18-transistor cell replaces a full flip-flop with programmable reset capabilities and a level shifter, which would typically require 31 and 17 transistors, respectively. Overall, in the 65nm process the custom block consumes a 6.8 μm^2 area relative to the 22.9 μm^2 that a standard cell flip-flop and level shifter would require. The shift register cell also shown in Fig. 3.19 contains 14 transistors and consumes an area of 5.31 μm^2, relative to the 27-transistor standard cell flip-flop with reset that requires 7.92 μm^2 area.

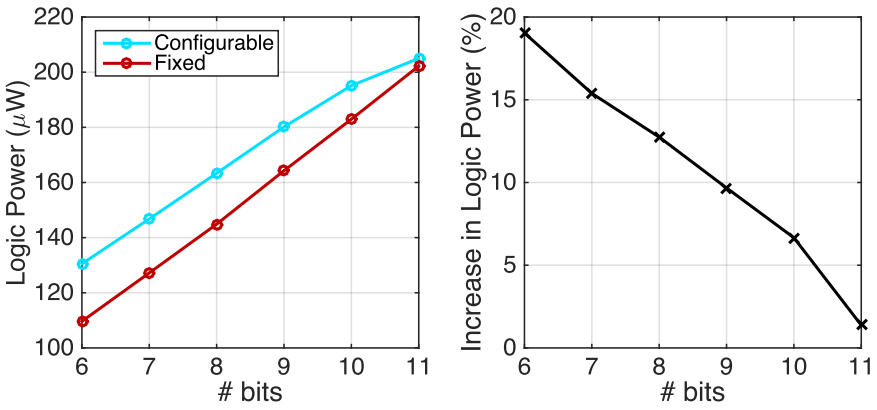

Figure 3.21. Comparison of power consumption for configurable (blue) and fixed (red) digital logic implementations.

3.3.3.2 *Expected performance*

Figure 3.19 shows the timing diagram of main control signals in the digital logic. The ADC conversion speed is limited by the 146 ps delay (characterized via post-layout simulation) from the rising edge of the comparator output to the transition of the appropriate DAC control signal. Overall, the logic consumes 205 μW of power when operating with an 80 MS/s sampling clock under typical operating conditions. The simulated power breakdown at full-scale operation is shown in Fig. 3.20. Almost half (43.5%) of the digital logic power consumption is attributed to comparator clock generation and driving the requisite logic blocks with this clock. Because the switching frequency of the comparator clock for an N bit, 80 MS/s ADC is $N \cdot 80$ MHz, any node driven by the clock or comparator draws a large amount of dynamic power. For instance, most power consumed by the output latch cells generating the DAC control signals is due to the comparator output buffer driving the input to all of the latch cells.

3.3.3.3 *Scalability impact*

Figure 3.21 compares the simulated power drawn by a fixed-resolution logic scheme to the proposed scalable-resolution logic. Relative to a fixed-resolution logic implementation containing only the required number of output latch and shift register cells, the configurable topology slightly increases the parasitic capacitive loading, which translates to higher power consumption. The results in Fig.3.21 are obtained by scaling the simulated schematic-level power consumption by a factor of 1.5 to match the total power consumption obtained via post-layout extraction simulation. As expected, the power overhead of the configurable logic has the largest impact when the number of conversion bits is small. In the highest-resolution setting, the power overhead is negligible. Because the power consumption breakdown in Fig. 3.20 indicates that a significant fraction of the total power is consumed by clock generation, any unnecessary loading on the comparator clock introduced by unused cells should be minimized. Nevertheless, the power penalty introduced by implementing resolution-scalable logic using this topology is consistently under 20%.

3.3.4 Comparator

3.3.4.1 *Design*

The comparator must be able to trade power consumption for input-referred noise and offset. Moreover, to support a low-power DAC switching procedure that varies the common-mode input voltage (V_{CM}) to the comparator at each conversion step, its offset must remain constant across V_{CM} to avoid significantly degrading the resolution of the ADC. In this implementation, eight identical comparators that can be individually enabled or disabled are placed in parallel. Noise can be reduced by increasing the number of parallel comparators, and offset can be calibrated by selecting combinations of comparator unit cells whose V_{CM} dependencies cancel. The power consumption of the comparator will be dictated by the number of comparator unit cells and the switching energy of each cell, with some additional overhead power due to the configuration logic and switches.

Typically, comparator offset is eliminated at a fixed V_{CM} by tuning output capacitance or injecting a fixed offset cancellation current at the input of the comparator, but these techniques will not mitigate the comparator's sensitivity to V_{CM}. Prior SAR ADCs with monotonic switching procedures have introduced an analog bias voltage to the comparator's tail current to limit the differential input pair's sensitivity to V_{CM} [71], or utilized separate comparators with independently-tuned offsets at each comparison [79]. To minimize system-level design complexity and reduce overhead power consumption, however, any calibration techniques should be simple to implement and avoid consuming static power. The V_{CM} sensitivity reduction approach in [71] requires a bias voltage generator and limits the comparator's decision speed by reducing the bias current. Moreover, it requires the tail current source transistor to have a high output resistance which is difficult to realize in deeply scaled process technologies. While the latter tuned-comparator approach enables good performance and greatly simplifies the converter's logic implementation, it requires extensive calibration overhead. To minimize V_{CM} sensitivity using the unit cell technique incorporated in this design, individual comparator V_{CM} characteristics can be measured a single time using the variable V_{CM} switching procedure. These eight measurements can be used to determine which combination of unit comparators minimizes V_{CM} dependence. Figure 3.25 illustrates how individual unit comparators could be chosen in a hypothetical example with four parallel comparators. Further analysis of the eight comparator case is presented in Section 3.4.3.3. While the V_{CM} dependence of the lowest-noise comparator setting with all unit cells enabled cannot be tuned, any comparator setting with k enabled cells of n total will have $\binom{n}{k}$ possible V_{CM} dependencies to choose from.

Each sub-comparator is a double-tail comparator based off the topology in [80], shown in Fig. 3.22(a), which contains a preamplifier to reduce noise and prevent signal-dependent kickback onto the DAC capacitor. To allow unit cells to be individually enabled, switches are placed at the output of both the preamplifier and the final comparator stage, and a NAND gate is used to disable the clock input. Because the output of the first stage controls the second-stage reset transistors, only a single clock phase is required. As shown in Fig. 3.22(b), when the clock (Φ) is low, the bias current in the input differential pair is disabled and nodes V_A and V_B are pulled high by the two PMOS reset transistors in the first stage. This closes the NMOS reset switches in the second stage so that both output nodes are pulled low. When Φ goes high, the reset transistors in the first stage turn off and the differential input pair in the

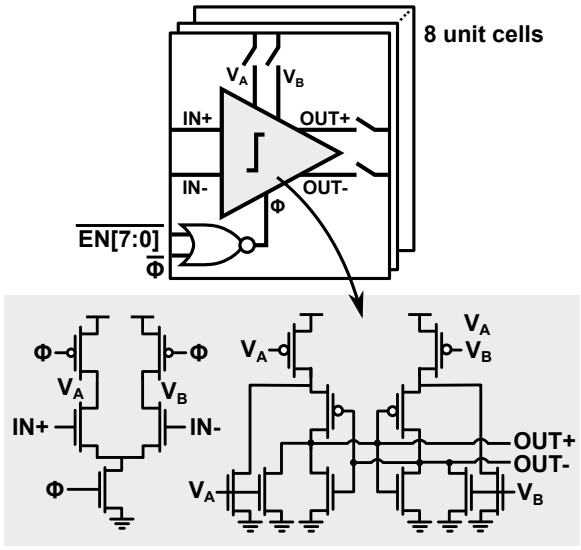

(a) Schematic diagram of comparator unit cell.

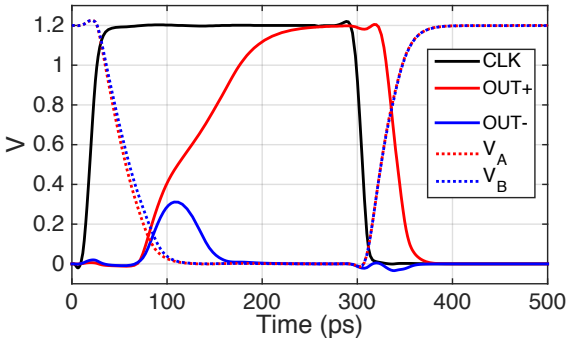

(b) Transient operation of comparator.

Figure 3.22. Operation of scalable double-tail comparator cell.

preamplifier pulls charge from the parasitic capacitance at V_A and V_B. The voltage at these nodes falls linearly according to the relationship $V = \frac{I}{C}t$, with a differential output voltage $(V_A - V_B)$ proportional to the differential input voltage (v_{id}) if the transconductance (g_m) of the input devices are approximately constant in this preamplification phase. As V_A and V_B fall, the NMOS reset transistors in the second stage turn off while the pseudo-differential PMOS input devices turn on, generating a new differential current set by V_A and V_B. When this differential current is sufficiently large, positive feedback in the cross-coupled inverter latch of the second stage will amplify this current and store the comparator decision.

To ensure uniform layout and minimize systematic mismatch, all devices within the comparator are constructed as multiples of a uniform unit finger of fixed length and width. To reduce noise as described in [75], the comparator was sized to minimize the overdrive voltage of the preamplifier's differential input stage. As a result, the NMOS input pair devices have the largest effective width, while the bias device is relatively small. The PMOS reset devices in the first stage and NMOS reset devices in the second stage consist of a single unit finger in order to reduce the unit comparator cell's power consumption by lowering its parasitic capacitance. However, the output enable switches introduce additional parasitic capacitance that translates into a fixed amount of overhead power consumption.

3.3.4.2 *Expected performance*

Figure 3.23 shows the simulated noise, energy per conversion, and decision speed of the comparator as a function of common-mode input voltage (V_{CM}) and number of enabled unit cells. The input-referred offset characteristics of a single comparator unit cell, obtained via Monte Carlo simulations, are shown in Fig. 3.24. All simulations were performed on a post-layout extraction netlist. Because the large, fixed capacitive load of the enable switches limits the effective noise bandwidth of the unit comparator cell, the input-referred comparator noise scales by much less than a factor of $\sqrt{2}$ when the number of unit comparator cells grows from 1

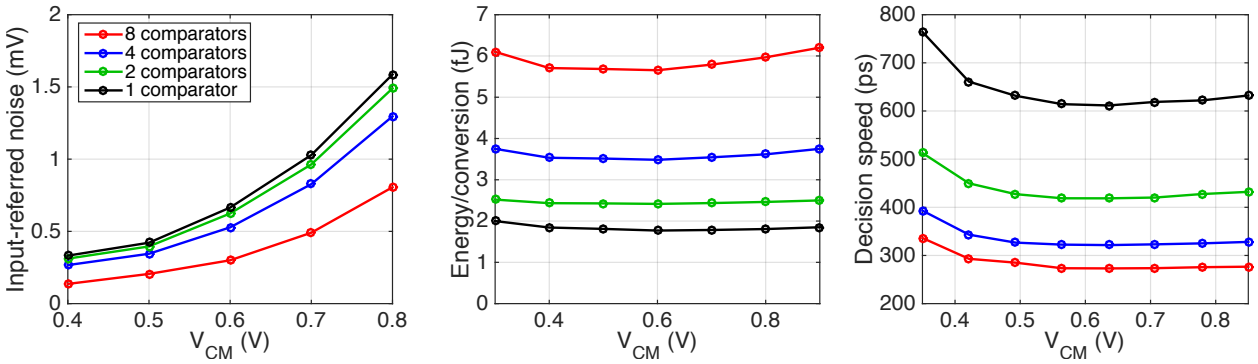

Figure 3.23. Simulated comparator noise, energy per conversion, and speed vs. common-mode input voltage (V_{CM}).

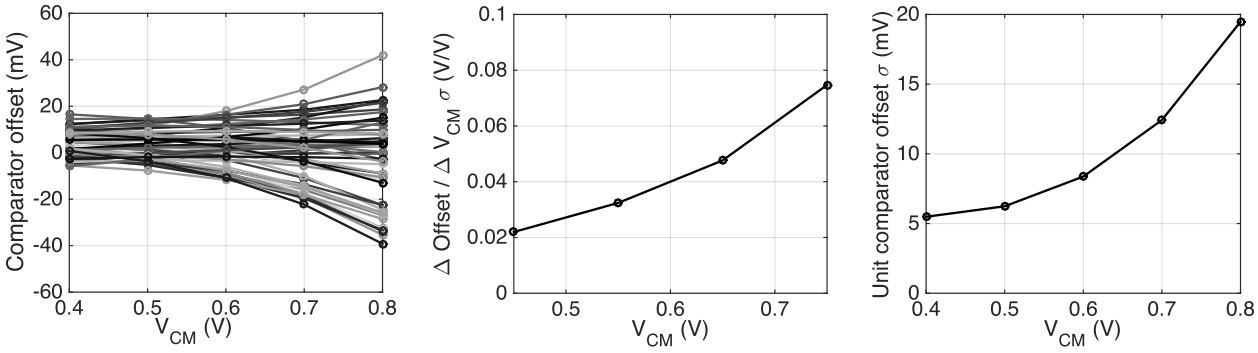

Figure 3.24. Simulated comparator offset characteristics vs. common-mode input voltage (V_{CM}). Monte Carlo simulations of comparator offset (L) were used to estimate offset V_{CM} sensitivity standard deviation (center) and offset standard deviation (right).

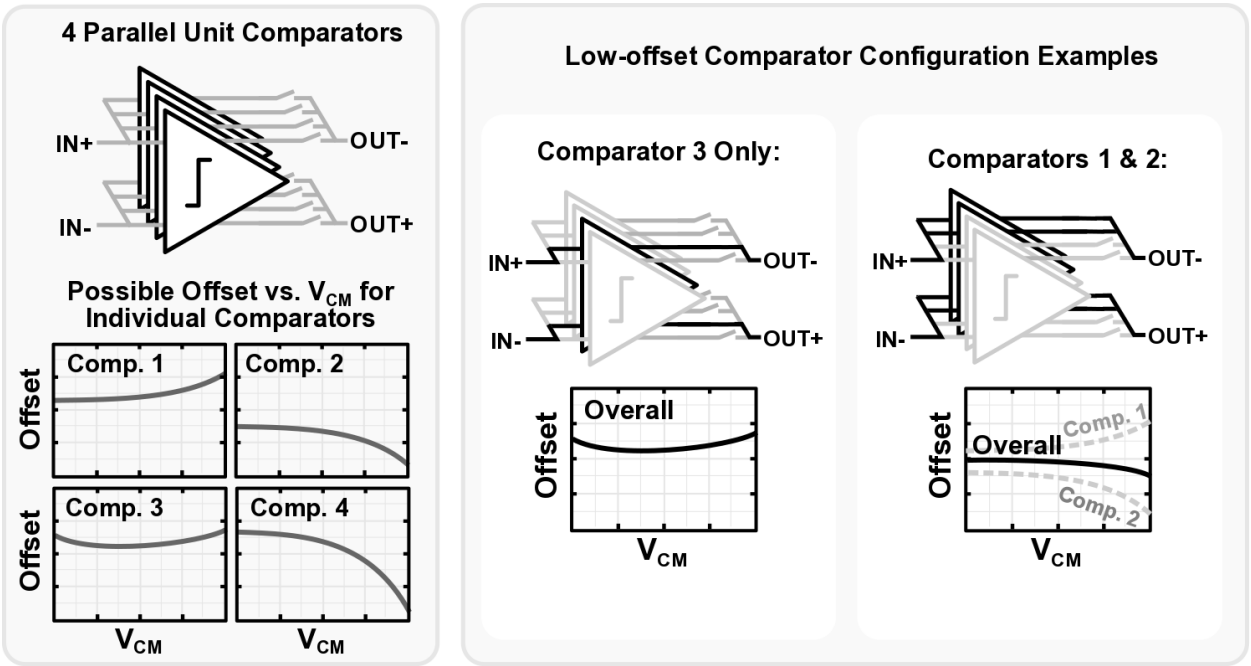

Figure 3.25. Theoretical example of how to choose unit comparator elements to cancel offset dependence on V_{CM}. From the measured offset characteristics that vary due to random fluctuation (left), the comparator can be configured to utilize either the most V_{CM}-independent cell (center), or subsets of cells whose V_{CM} dependencies cancel (right). While a 4-element example is shown, the prototype includes 8 parallel comparators.

to 2. The non-scalable load of the enable switches also causes the single-element comparator to consume more than half of the two-element comparator's energy per conversion. This illustrates the limitations of the configurable comparator's efficiency relative to a minimum-sized comparator optimized for use in a low-power, low-resolution ADC.

While the parasitic overhead of the enable switches limits the comparator's scalability, V_{CM} tuning provides an alternative means of modifying the comparator's input-referred noise and offset. The reduction in noise with V_{CM} is expected, because lowering the input voltage reduces the overdrive voltage on the input stage. In turn, this reduces the speed of the decision, which limits the effective noise bandwidth of the comparator to reduce the total integrated noise per decision at the expense of speed. As shown in Fig. 3.23, however, less than a 50% increase in decision time can lower input-referred noise by a factor of over 2-3x. Similarly, the variance of comparator offset and offset sensitivity to V_{CM} can be lowered by nearly a factor of 4. Because the comparator's energy per conversion is largely dictated by the total parasitic capacitance of the comparator and its supply voltage, it is mostly independent of V_{CM}. The slight increase in energy consumption at low V_{CM} is caused by crowbar current due to the slower comparator decision time. To minimize noise, V_{CM} should therefore be as low as the target conversion speed can allow. Having the flexibility to tune V_{CM} allows the comparator performance to be optimized to account for process technology variability.

3.3.4.3 *Scalability impact*

Figure 3.26 compares the schematic-level design and layout of a scalable and stand-alone comparator. The area overhead of the NOR gate and output switches causes a 220% area increase relative to the size of a single fixed-performance comparator, partly due to the increased width of the overhead components. This increases the internal parasitics of the comparator unit cells, which slows the comparator's decision time and increases the power consumption of a unit comparator cell but also lowers the noise of a minimum-sized comparator.

Figure 3.27 demonstrates that the parasitics introduced by the sampling switches have a fairly significant impact on the comparator's power consumption, speed, and input-referred noise. The total power consumption of the comparator operating at 240 MHz (equivalent of 6-bit conversion for a 40 MS/s ADC) shows that the configurable topology exhibits a substantial 5x power overhead in the minimum-size configuration, which decreases to 2x in the maximum-sized comparator setting. As shown in Fig. 3.27, this increased power is coupled with a longer decision time due to the large parasitics of the configurable unit cell. This increase in decision time translates to a smaller comparator noise bandwidth that reduces the total integrated noise, as also illustrated in the figure. While this is beneficial for high-resolution designs that are limited by comparator noise, the power overhead of the scalable comparator will degrade the efficiency of the ADC in low-power settings. Figure 3.27 also contrasts the comparator figure of merit (FoM) calculated as $(P_{comp} \cdot t_{comp} \cdot v_{n,comp}^2)^{-1}$ achieved by the two designs. Overall, the configurable design's 65-35% reduction in noise is insufficient to compensate for its significant increase in power consumption and reduction in speed, leading to a less efficient design overall in order to afford configurability.

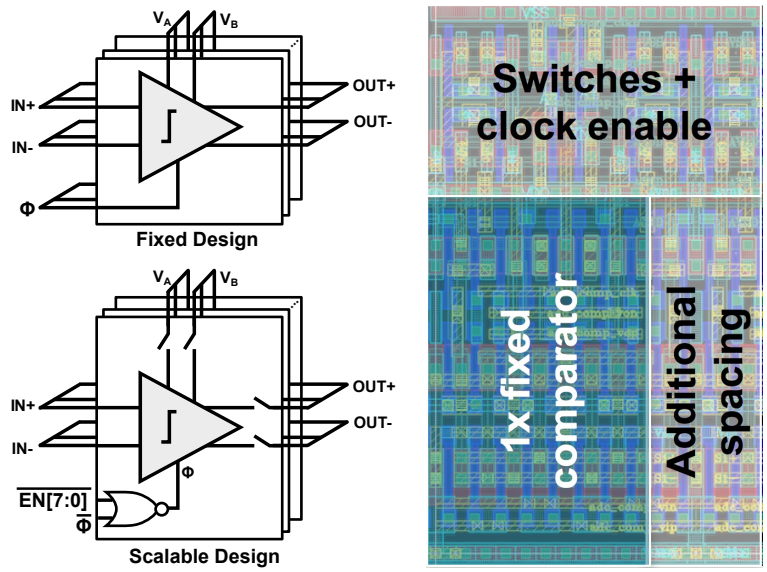

(a) Schematic-level comparison of fixed and configurable comparator unit cells.

(b) Area breakdown of configurable comparator.

Figure 3.26. Comparison of "configurable" and "fixed" comparator topologies.

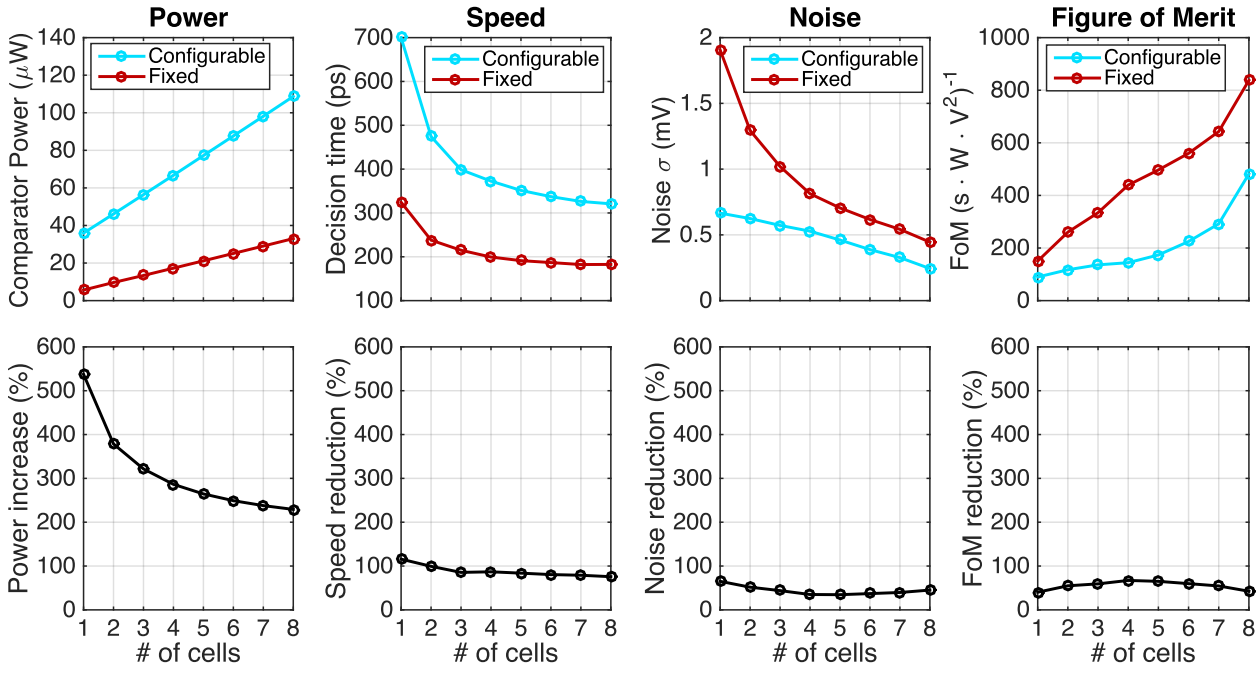

Figure 3.27. Comparison of "configurable" and "fixed" comparator performance.

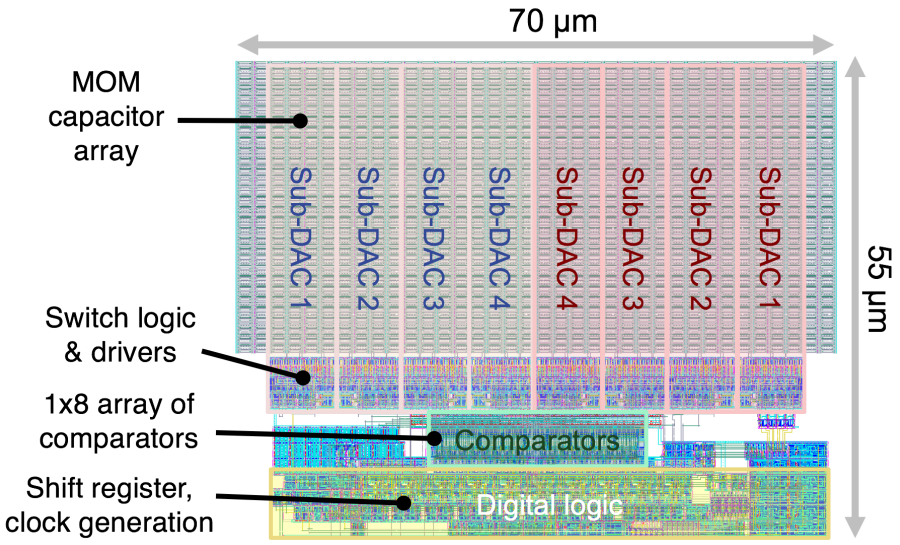

Figure 3.28. Scalable ADC layout.

3.3.5 Top-level design

A prototype of this scalable converter topology was implemented in a 65nm CMOS process. The layout of a single converter is shown in Fig. 3.28; it fits in a compact 70 μm by 55 μm area. The four sub-DAC arrays are located at the top of the converter, and the two halves of the differential DAC are located at the left and right hand side of the layout. Sampling switches, DAC buffer drivers, and enable logic are all located at the base of the DAC array, and a single binary-to-thermometer encoder is included in each DAC half. Immediately underneath the DAC driver logic is the reconfigurable comparator; this placement ensures that the comparator remains close to the DAC. A ground shield is placed along the route from the top plate of the DAC to the comparator input to mitigate coupling of control signals to the comparator input. The digital logic, which operates off of the comparator result, is immediately underneath the comparator. A bank of eleven output registers to store the ADC result is located to the right of the ADC's control logic to interface with the digital baseband circuitry.

Figure 3.30 shows a die micrograph of the chip containing the ADC prototype. The ADC was integrated with a scalable radio transceiver, on-chip supply regulation, and digital test infrastructure developed by other team members [1]. The majority of the chip area is dedicated to these additional blocks; particularly, the digital baseband processing circuitry and the large baseband capacitance shared between the receiver and transmitter. Two converters — one for the in-phase (I) and one for the quadrature (Q) path of the receiver — are contained on the die.

Fig. 3.29 illustrates how the ADC interfaces with the receiver (in blue) and digital infrastructure (in black). A set of on-chip switches enables the ADC input to be taken either directly from the output of the RF receiver or fed from an external source to characterize the stand-alone performance of the ADC. An optional inverter-based buffer can also be used to drive the ADC input at high frequencies when the capacitive load of the converter is too

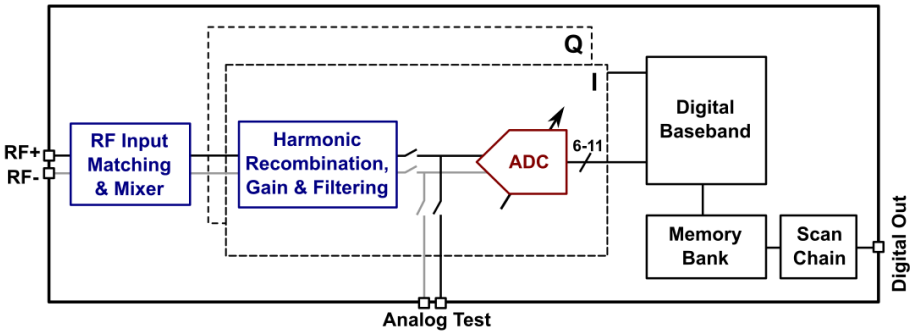

Figure 3.29. Diagram of major blocks (RF receiver, shown in blue, and digital, shown in black) interacting with ADC, and test enable switches.

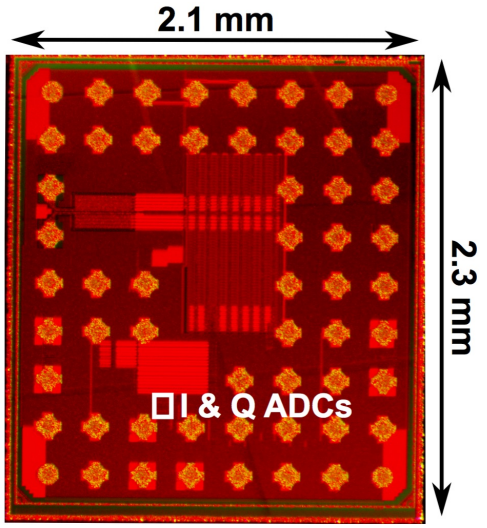

Figure 3.30. 65nm CMOS Die micrograph.

large for the amplifier at the receiver output to drive directly. The output of the ADC is fed to digital baseband processing circuitry that includes an integrated memory bank for ADC characterization, which is read out using a scan chain interface.

3.4 Measurement results

This section describes measurements of a 65nm CMOS prototype of the resolution- and power-configurable ADC. After describing the measurement setup in Section 3.4.1, Section 3.4.2 discusses the achievable resolution and power scaling range of the ADC operating as both a conventional Nyquist-rate 40 MS/s ADC and a 4x oversampled ADC integrated with a radio receiver. In the latter application, the oversampling rate of the ADC relaxes the anti-aliasing filter requirements of the receiver front-end. After reviewing the converter's general capabilities, Section 3.4.3 explores how each of the four tunability mechanisms implemented affect the converter's achievable resolution and power consumption. Building on these details, Section 3.4.5 concludes by using a combination of prototype measurements and post-layout extraction simulations to identify the main sources of power consumption and resolution

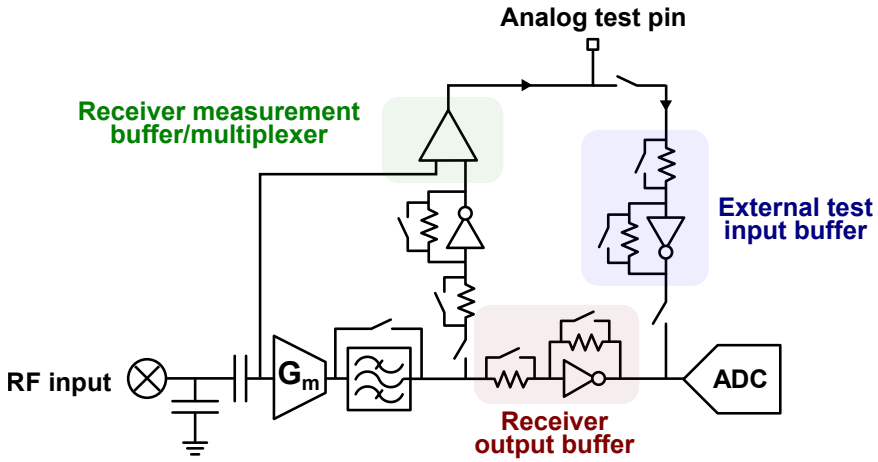

Figure 3.31. Interface between analog test pin, radio receiver, and ADC.

degradation for this converter. These details are used to summarize the limitations of this scalable ADC topology.

3.4.1 Measurement setup

To characterize the stand-alone ADC performance (independent of the receiver front-end), an off-chip test signal can be provided to the ADC through the analog test pin. Figure 3.31 shows the on-chip interface between the analog test pin, RF receiver output, and ADC. Because the analog test pin is also used to measure the receiver output, switches are used to configure the pin as either an input or an output. An inverter-based buffer with a gain of nearly unity is included in the prototype as a means of driving the large input capacitance of the ADC at high frequencies. This buffer can be disabled, however, using additional switches to disconnect the buffer's power supply and bypass its feedback resistors.

For stand-alone ADC testing, the receiver measurement output buffer was disabled and switches along the path from the analog test pin to the ADC input were closed. The external input buffer was also disabled, as it contributed additional noise and nonlinearity to the test signal without significantly improving the bandwidth of the converter's input sampling network. Switches along the path from the receiver output to the ADC input were left open, and the receiver output buffer was also disabled to ensure that only the externally applied test signal was measured. To measure the combined performance of the receiver and ADC, all switches along the path from the analog test pin to the ADC were left open while the switches along the receiver path were all enabled. The receiver output buffer was also disabled, as the final amplifier in the receiver had sufficient strength to drive the capacitive load of the converter.

Figure 3.32 illustrates how the test chip interfaces with external components for measurements. To characterize the overall performance of the ADC, a differential sine wave with a fixed common-mode offset was supplied to the analog test pin. This signal was created using a dual-output waveform generator with programmable common-mode voltage, amplitude, and phase offset. An on-chip frequency divider was used to generate the approximately 40 MHz and 80 MHz ADC sampling clocks from 160 MHz or 320 MHz digital clock signals, respectively. The exact sampling frequency of the converter was chosen to eliminate spectral

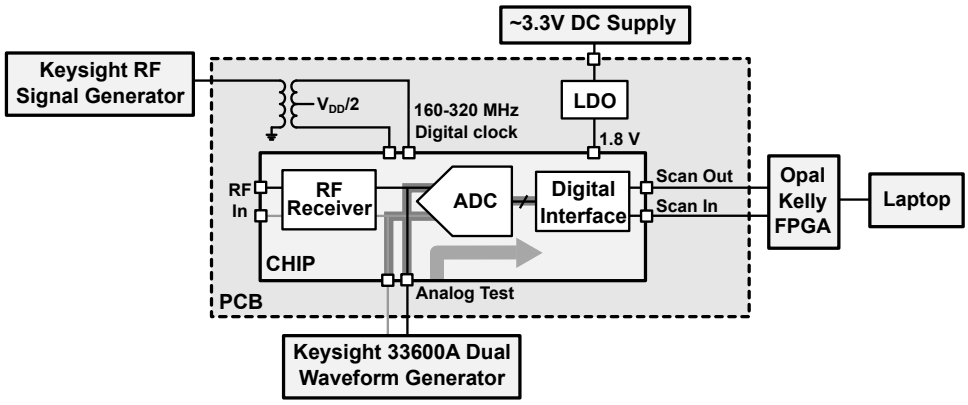

Figure 3.32. ADC test configuration.

leakage in the FFT used to measure the converter's resolution.

The ADC is powered by on-chip LDOs that use an externally supplied reference voltage. This reference voltage is generated using discrete LDOs on the test PCB that regulate the output of a benchtop power supply, as shown in Fig. 3.32. Because the current draw of the analog and digital supplies of the test chip are not measured independently, and the analog and digital supplies of the ADC are shared with other blocks on the chip, the total power consumption of the ADC is measured by taking the difference of total chip-level power when the ADC is enabled and disabled. When the ADC is disabled, it no longer consumes dynamic power; however, the static power consumption due to leakage cannot be isolated. Measurement results suggest that the majority of power consumption is dynamic, particularly at operating speeds of 40-80 MS/s, and the measured power numbers are consistent with post-layout extraction simulation results.

3.4.2 Overall scalability

Figure 3.35 shows the best-achievable SNR, SNDR, and Walden figure of merit (FoM) achievable by the converter for a particular power range. For these measurements, the converter was configured as a 40 MS/s Nyquist-rate ADC. The maximum achievable SNR (obtained by omitting the 10 largest harmonic tones from the calculated noise power) is 53.5 dB as the comparator and DAC sizes are scaled. Significant sampling switch distortion — partially a result of the series resistance added by test switches to configure the functionality of the analog test pin — restricts the maximum achievable SNDR to 49 dB. Because none of the power-scalable components are designed to mitigate the effects of sampling distortion, the FoM increases for higher-power ADC configurations operating in the distortion-limited regime. The FoM also increases for low-power ADC configurations with quantization-limited resolution. In the quantization noise-limited regime, the fixed overhead power of the SAR logic and the finite power scalability of the DAC and comparator severely limit the achievable efficiency of the converter. The ADC's FoM is optimized when the noise and distortion of power-scalable design components such as the DAC and comparator dominate the total SNDR of the converter.

To illustrate the effects of distortion and quantization noise on the achievable ADC resolution, Fig. 3.34 compares the frequency response of the ADC in its lowest-power,

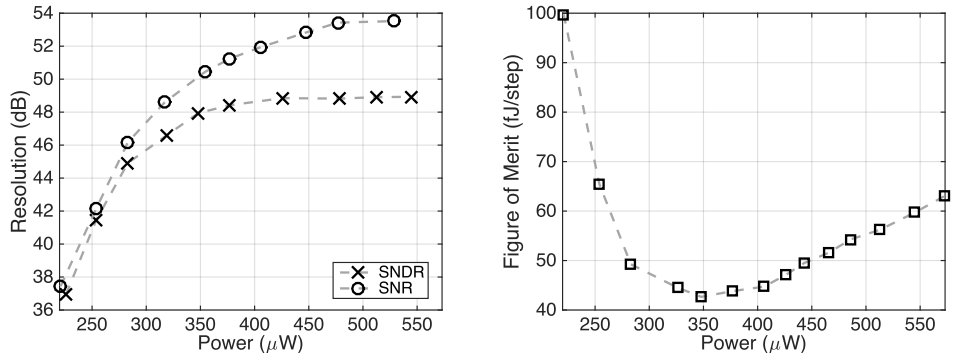

Figure 3.33. Best achievable SNDR, SNR, and Walden Figure of Merit (FoM) as a function of total power consumption when the ADC is configured as a 40 MS/s Nyquist-rate ADC.

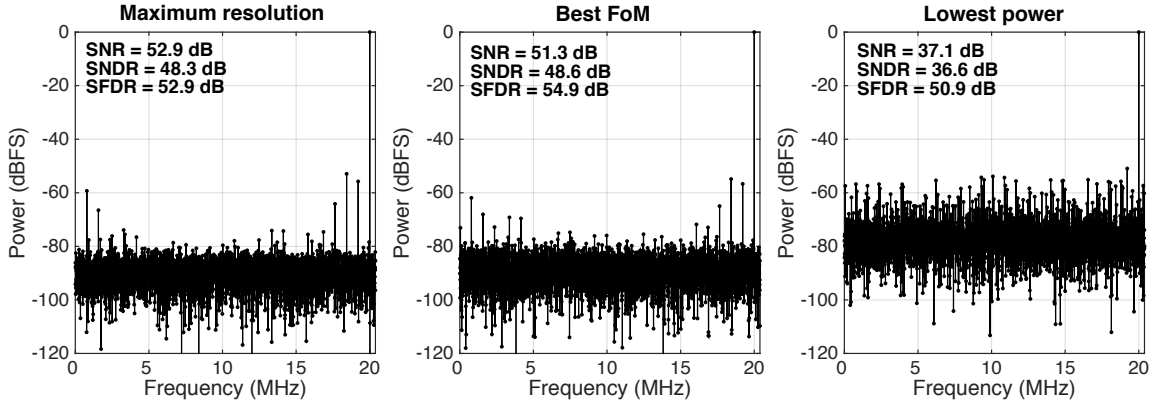

Figure 3.34. FFT of Nyquist-rate 40 MS/s ADC in highest resolution, best FoM, and lowest power configurations.

lowest-FoM, and highest-SNR configurations. In the lowest-power setting, quantization noise dominates and the noise spectrum is relatively flat. In the optimal FoM and highest-SNR settings, sampling distortion introduces harmonics of the input tone that limit the spurious-free dynamic range (SFDR) of the ADC to under 55 dB. The highest-SNR configuration uses the largest DAC capacitance, which increases the effect of distortion due to the sampling switch resistance and results in 52.9 dB SFDR. In the optimal (lowest) FoM setting, the sampling distortion is comparable to the total thermal noise of the ADC, allowing for a more efficient converter.

Alternatively, the ADC sample rate can be increased to 80 MS/s while keeping the target bandwidth fixed to 10 MHz, providing a modest 4x oversampling ratio. When used as an oversampled ADC, the impact of sampling distortion reduces significantly, though the converter's FoM increases. Figure 3.35 summarizes the best SNR, SNDR, and Walden FoM measured for this ADC when configured as an 80 MS/s converter providing a 4x oversampling ratio (targeting a 10 MHz signal bandwidth). These results are reported for a 2 MHz input tone to ensure that the first four harmonics of the input tone fall into the signal bandwidth. Because the effect of sampling distortion is lower, the maximum achievable

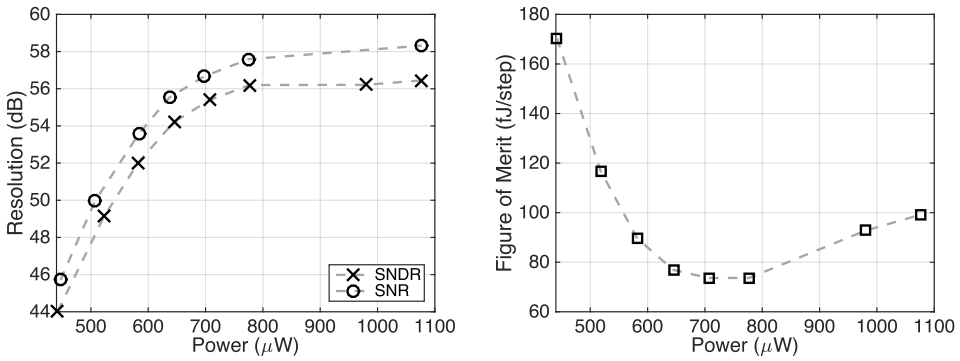

Figure 3.35. Best achievable SNDR, SNR, and Walden Figure of Merit (FoM) as a function of total power consumption when the ADC is configured as an 80 MS/s oversampling ADC with a 10 MHz signal bandwidth (oversampling ratio of 4).

SNR and SNDR are much closer. However, because Walden FoM is inversely proportional to the effective signal bandwidth, the 2x reduction in bandwidth and 2x power increase does not compensate for the converter's improvement in resolution. Despite the degradation in stand-alone ADC performance, the system-level efficiency can be improved using this oversampled configuration due to the relaxed performance requirements of the receiver. In the prototype measurements, the ADCs consume only up to 28% of the total receiver power, operating at a 250 MHz carrier frequency in the minimum-resolution configuration. At a 1.7 GHz carrier frequency in the maximum-resolution configuration, the ADCs consume only 8% of the total receiver power. Because the receiver power is significantly higher than that of the ADC, allocating more power to the ADC to reduce the filter order of the receiver lowers the total system power.

The reduced impact of distortion for the oversampled ADC can be seen in the frequency response shown in Fig. 3.36. In this plot, harmonics generated by sampling distortion are significantly smaller than those generated by operating the converter at Nyquist rate. The SFDR of the ADC in the maximum-SNR configuration increases from 53 dB in the Nyquist-rate case to 66 dB. Harmonics of any input signals at the edge of the signal bandwidth can be filtered digitally, so only input frequencies below 5 MHz or 3.33 MHz will generate second or third harmonic tones, respectively, within the 10 MHz signal bandwidth. While the maximum achievable SNDR of the ADC improves by nearly 8 dB relative to the Nyquist-rate sampling case, the high-resolution scalability of the converter is still limited by additional sources of SNDR degradation such as DAC component mismatch.

Finally, to confirm that the ADC does not significantly degrade the noise performance of the scalable mixer-first receiver, the SNDR measured by the ADC as a function of RF input power is presented in Fig. 3.38. As illustrated in Fig. 3.37, this measurement was performed by feeding the input signal through the receiver. The measured stand-alone receiver noise figure (NF) in the two configurations shown is 13.6 dB (maximum size) and 19.0 dB (minimum size). For the 10 MHz signal bandwidth, this should yield 0 dB SNDR at input powers of -90.2 and -84.8 dBm, which is consistent with the measured results. A slight degradation in sensitivity is expected due to the fact that the NF of a mixer-first receiver degrades at frequencies further from baseband DC (the RF carrier frequency). Due

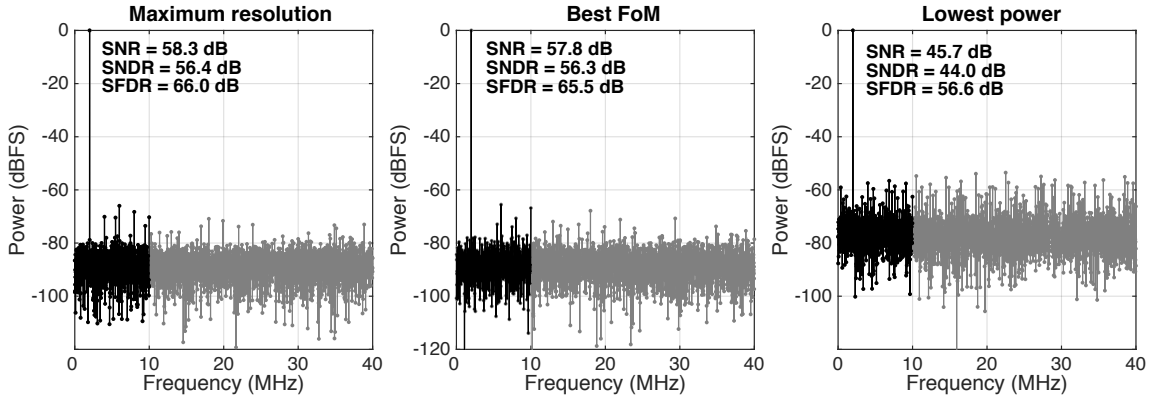

Figure 3.36. FFT of 4x oversampled 80 MS/s ADC in highest resolution, best FoM, and lowest power configurations. The 10 MHz in-band region is highlighted in black.

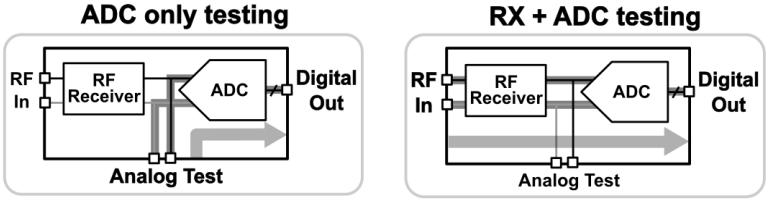

Figure 3.37. Testing configuration for stand-alone ADC measurement (left) and combined receiver and ADC measurement (right).

to the high receiver NF, the maximum ADC SNDR is limited to 30 dB for a full-scale input signal. However, the noise averaging properties of large arrays will improve this resolution by $3\log_2(N_{elem})$ dB, where N_{elem} is the number of elements in the array.

3.4.3 Performance across configuration schemes

This section presents the best-achievable Nyquist-rate SNR, SNDR, and FoM as a function of each configuration setting to illustrate the impact of the four major tuning mechanisms implemented in the prototype — DAC scaling, digital logic scaling, comparator scaling, and switching procedure tuning.

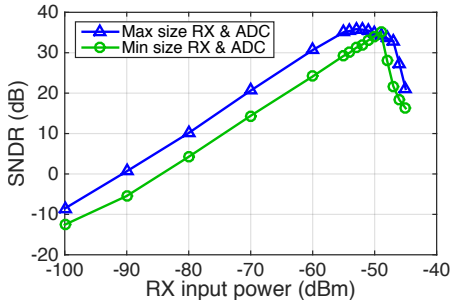

Figure 3.38. SNDR measured by the ADC as a function of receiver input power [1].

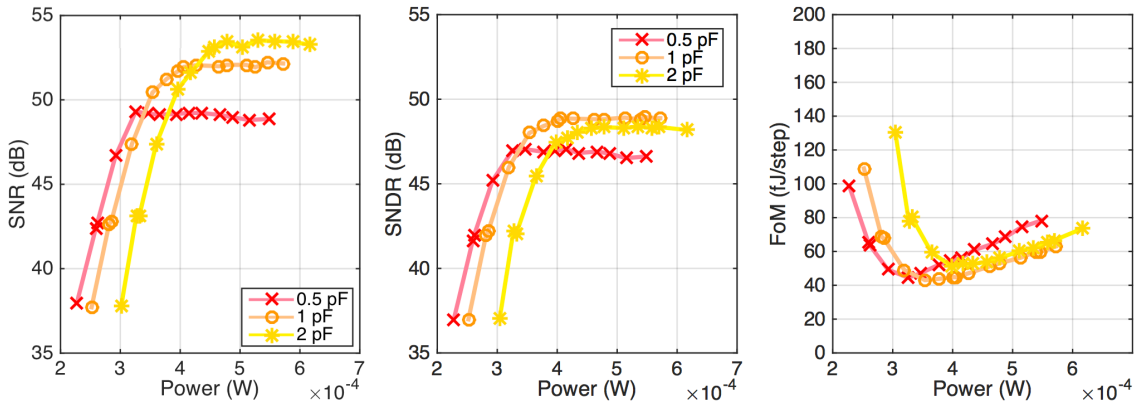

Figure 3.39. SNR, SNDR, and FoM as a function of DAC size.

3.4.3.1 *DAC scaling*

Figure 3.39 illustrates how changing the DAC capacitance affects the SNR, SNDR, and FoM of a 40 MS/s Nyquist-rate ADC. While increasing the DAC capacitance reduces the thermal noise sampled onto the ADC, it also increases the impact of sampling switch distortion due to signal-dependent switch resistance if the sampling switch size is fixed. While the size of the DAC sampling switch within the ADC scales with DAC capacitance, the size of additional switches in the testing path used to configure the usage of the analog test pin do not scale. As a result, Fig. 3.39 shows that while the 2 pF DAC setting improves the maximum achievable SNR of the ADC, the maximum SNDR drops relative to the 1 pF case due to the higher distortion.

Because the 2 pF, 11-bit DAC is constructed from four identical 500 fF, 9-bit sub-DACs, the DAC capacitance is also linked to the number of conversion bits. While this means that the 2 pF DAC has lower quantization noise, the larger DAC capacitance also increases the relative impact of DNL because it requires more unit capacitors to be switched in order to evaluate the most significant bit (MSB) of the ADC.

At low resolutions of 37-48 dB SNDR, Fig. 3.39 illustrates that increasing the DAC capacitance increases the total ADC power consumption but has no impact on the achievable resolution. In these configurations, alternate sources of noise and distortion — particularly quantization noise — dictate the resolution of the ADC. To improve power efficiency at these low resolution levels, the DAC scaling range could be enhanced. However, in the 37 dB SNDR setting, cutting DAC capacitance by 50% from 1 pF to 0.5 pF provides only a 10% reduction in overall ADC power, from 250 μW to 225 μW. While the DAC size does not significantly impact the low-resolution performance of the ADC, it provides an upper limit on the ADC's achievable SNR and SNDR due to distortion and DNL induced by unit element mismatch. The maximum achievable SNDR could be improved by increasing the sampling switch size to reduce the amount of sampling distortion, and by introducing redundancy or calibration to account for DAC mismatch effects.

3.4.3.2 *Logic scaling*

Figure 3.40 shows the effect of scaling the number of conversion bits on the SNR, SNDR, and FoM of a 40 MS/s Nyquist-rate ADC. In the low-resolution settings (6-7 bits), the

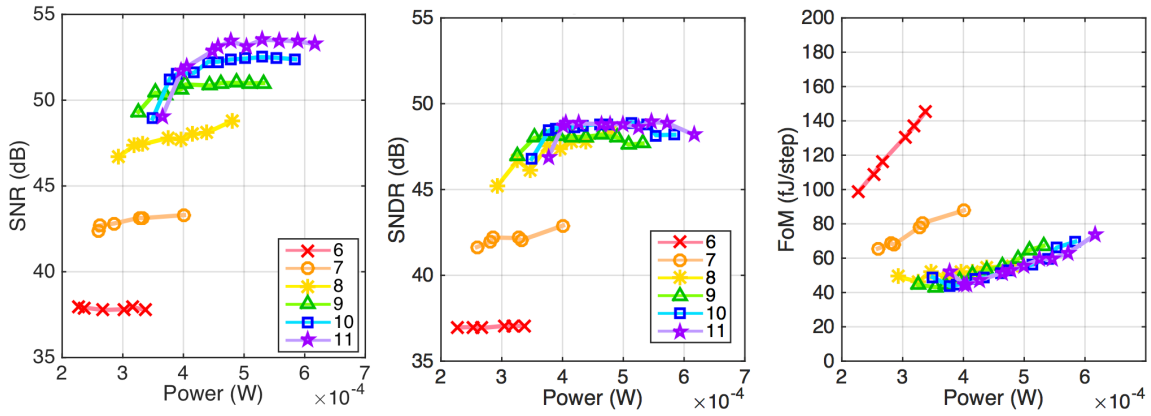

Figure 3.40. SNR, SNDR, and FoM as a function of number of bits.

measured resolution predominantly aligns with the expected quantization-limited SNDR ($6.02N_{bits}$+1.76 dB). As the number of conversion bits increases to 9 bits and above, however, the contribution of quantization noise relative to other noise sources shrinks rapidly, and the achievable SNR improves by only 1-2 dB per added conversion bit. Significant distortion limits the achievable SNDR to 48 dB, causing quantization noise to provide no noticeable improvement in SNDR at the 9-11 bit resolution level.

The optimal ADC FoM is obtained in the 9-bit ADC configuration, where quantization noise is comparable to other sources of noise and distortion within the ADC. The worst-case (highest) ADC FoM is measured in the 6-bit configuration due to the limited power scalability of the digital logic and the overhead power consumption of the minimum-sized comparator and DAC. As discussed in Section 3.2.3, the fact that logic power consumption is a linear function of the number of conversion bits fundamentally limits the achievable efficiency of the converter in the quantization noise-limited regime.

3.4.3.3 *Comparator scaling*

As shown in Fig. 3.41, the maximum SNR and SNDR of the ADC can be obtained when only a single comparator unit cell is enabled, suggesting that thermal noise introduced by the comparator does not have a substantial impact on ADC resolution relative to sampling distortion and DAC mismatch. As a result, increasing the number of parallel unit comparators from one to eight simply increases power consumption by 100 μW without improving resolution. As discussed in Section 3.3.4, the parasitic capacitance and reduction in sampling bandwidth caused by the configuration switches lowers the input-referred thermal noise of a stand-alone comparator.

In Fig. 3.41, common-mode input voltage (V_{CM}) variation due to the single-sided switching procedure does not significantly limit the achievable SNDR because comparator configurations have been chosen specifically to lower the impact of V_{CM}-dependent comparator offset. Figure 3.42 illustrates how this was accomplished by comparing the offset vs. V_{CM} measured for each individual sub-comparator to the offset vs. V_{CM} achievable by averaging the characteristics of a subset of comparators. This figure also presents the distribution of simulated SNDRs achievable for each possible comparator combination. Results are presented for two ADCs. The average digital code measured in response to a sinusoidal input

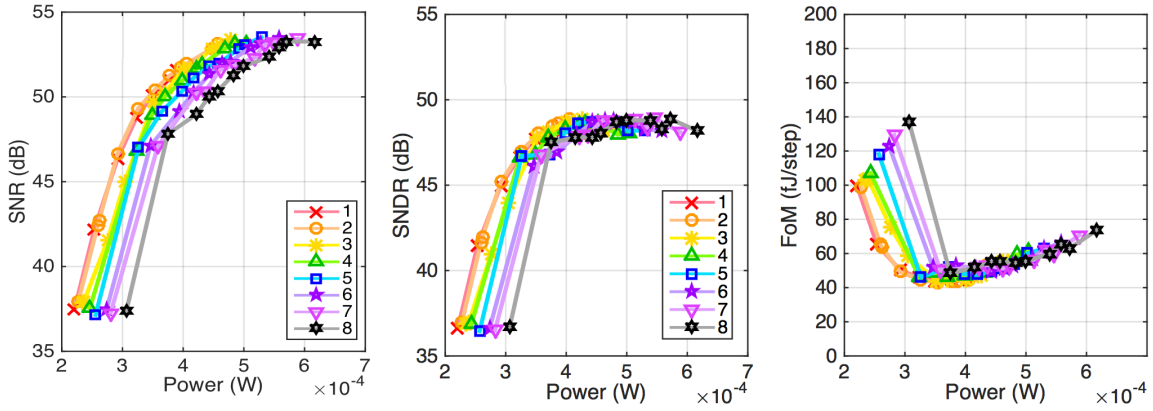

Figure 3.41. SNR, SNDR, and FoM as a function of number of comparators.

tone is used to characterize offset; this code is translated to a voltage using the known input voltage range of the converter.

When k sub-comparators are used in the design, a total of $\binom{8}{k}$ possible combinations of sub-comparators exist, translating to $\binom{8}{k}$ possible comparator offset vs. V_{CM} characteristics. For the low-power monotonic switching procedure that causes V_{CM} to drop from 600 mV to 350 mV, averaging can enable nearly constant offset over the relevant V_{CM} range, as seen in the center plot of Fig. 3.42(a). Moreover, disabling cells with a large degree of mismatch allows smaller comparator sizes to achieve a lower offset dependence on V_{CM} than the fixed 8-comparator implementation. To examine how this impacts performance, the rightmost plots in Fig. 3.42 present the cumulative distribution function (CDF) of SNDRs obtainable when all $\binom{8}{k}$ comparator configurations are considered, for $k = 1$ to 8. To isolate the effect of comparator offset and noise on ADC performance, the SNDRs in Fig. 3.42 were obtained via MATLAB simulation that incorporates the measured offset vs. V_{CM} characteristics and simulated noise vs. V_{CM} behavior of the comparator. The simulation models a monotonically decreasing SAR ADC switching procedure. Overall, the CDF demonstrates that the increased number of comparator characteristics obtained by individually selecting unit cells improves the comparator's performance relative to a larger stand-alone design.

If comparator offset dependence on V_{CM} was not reduced by selecting specific unit cell combinations, the number of unit cells could have a more significant impact on the ADC's achievable SNR and SNDR. The effect of V_{CM} dependence can be emphasized further by considering only the SNR and SNDR measured using the monotonically increasing DAC switching procedure, which biases the comparator at a high V_{CM} that significantly increases its noise and offset. Figure 3.43 summarizes the effect of comparator size when only this switching procedure is used. In this case, the number of comparators has a significant impact on the ADC's resolution. The single-comparator setting limits the achievable SNDR to 40 dB. The highest resolution can be reached only in the 4-comparator configuration, which has the most constant offset as a function of bias point. This plot demonstrates that if the switching procedure is not chosen properly, the converter's resolution is primarily constrained by the noise and offset variation of the comparator.

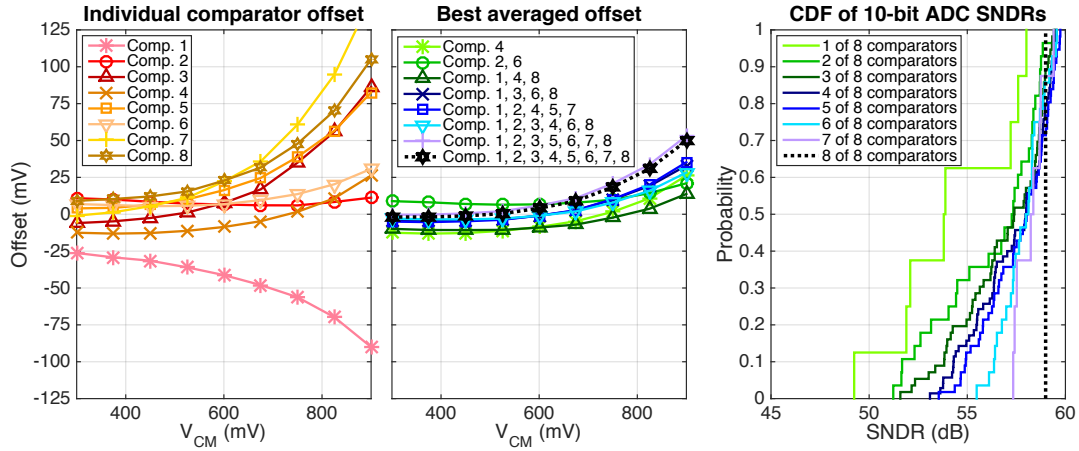

(a) Performance comparison of comparator measured on I ADC.

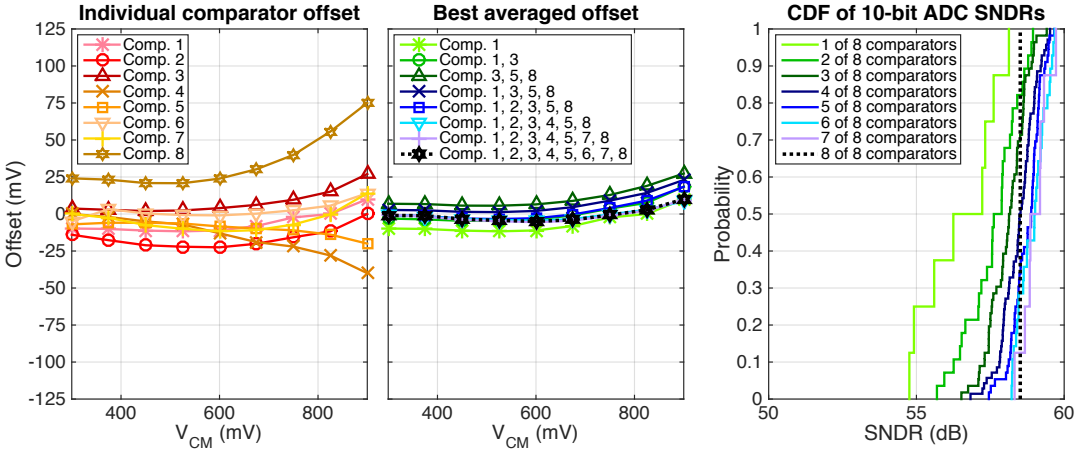

(b) Performance comparison of comparator measured on Q ADC.

Figure 3.42. Comparison of comparator offsets measured for I and Q ADCs on a sample die (left), best-achievable offset characteristics (center) and simulated distribution of and SNDRs obtained by averaging all possible subsets of comparators (right). Random variability introduces different comparator offset dependencies on V_{CM}.

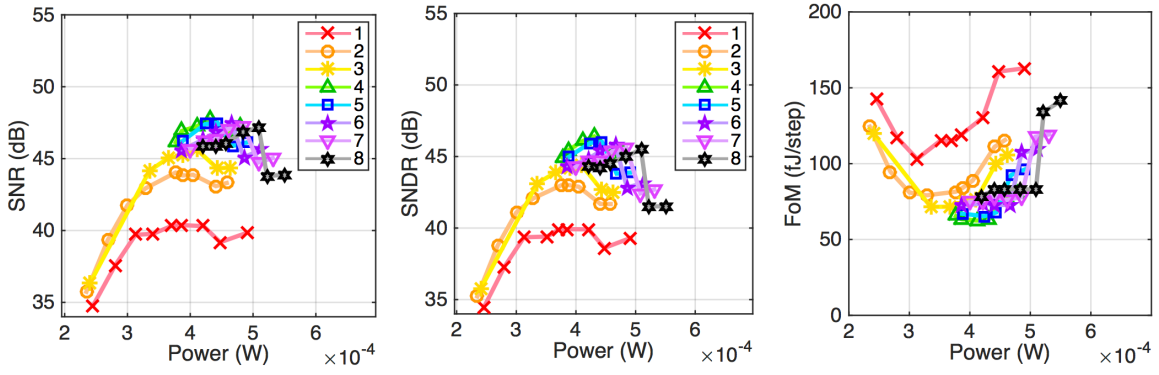

Figure 3.43. SNR, SNDR, and FoM as a function of number of comparators for only the high-V_{CM} switching procedure.

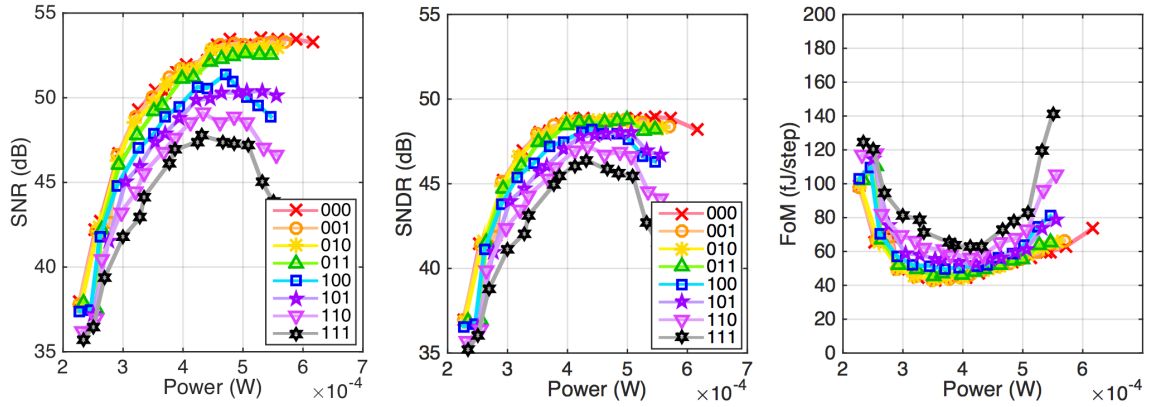

Figure 3.44. SNR, SNDR, and FoM for a 40 MS/s Nyquist-rate ADC as a function of switching scheme. The "000" code corresponds to the monotonically decreasing V_{CM} scheme, while the "111" code corresponds to the monotonically increasing V_{CM} scheme.

3.4.3.4 *Switching procedure tunability*

As seen from the measured dependence of ADC performance on the number of comparators, the common-mode input voltage to the comparator (V_{CM}) greatly affects its input-referred noise and offset. Figure 3.44 shows the measured SNR, SNDR, and FoM of the ADC as the DAC switching procedure is modified to adjust the comparator bias. As expected, the maximum achievable resolution occurs for the switching procedure corresponding to the lowest possible V_{CM}. This biases the comparator at the lowest noise and offset setting in the final iterations of the SAR algorithm, where the DAC residue is smallest. While specific comparator configurations were chosen to minimize offset sensitivity to V_{CM}, the noise sensitivity still significantly reduces the achievable resolution of the ADC in the high V_{CM} operating mode.

To illustrate how comparator biasing influences ADC performance, Figs. 3.45 and 3.46 show the measured differential nonlinearity (DNL) and frequency response (respectively) of the ADC as a function of switching procedure for comparator configurations with high and low V_{CM} sensitivity. As shown in Fig. 3.24, increasing the comparator's V_{CM} from 450 mV to 750 mV causes nearly a 4x increase in offset sensitivity to V_{CM}. From Fig. 3.45, it is clear that this sensitivity significantly degrades the converter's DNL. Even using the comparator with low V_{CM} sensitivity, the ADC's DNL increases from 2.6 LSBs using the low V_{CM} switching procedure to nearly 16 LSBs using the high V_{CM} switching procedure. While even the 2.6 LSB DNL obtained using the low V_{CM} switching procedure is high for the 11-bit ADC setting, it satisfies the 9 effective bit resolution range targeted by this system. Using the most sensitive comparator, the best-achievable DNL of the ADC is only 17.9 LSBs, with a worst-case DNL of 58 LSBs with the high V_{CM} switching procedure. This illustrates the importance of minimizing the sensitivity of comparator offset to V_{CM} when a single-sided DAC switching procedure is used.

The measured FFTs presented in Fig. 3.46 illustrate how the nonlinearity highlighted in Fig. 3.45 impacts the resolution of the converter. When DNL is large, missing ADC codes create discontinuities in the sampled data that generate sizeable spurs in the con-

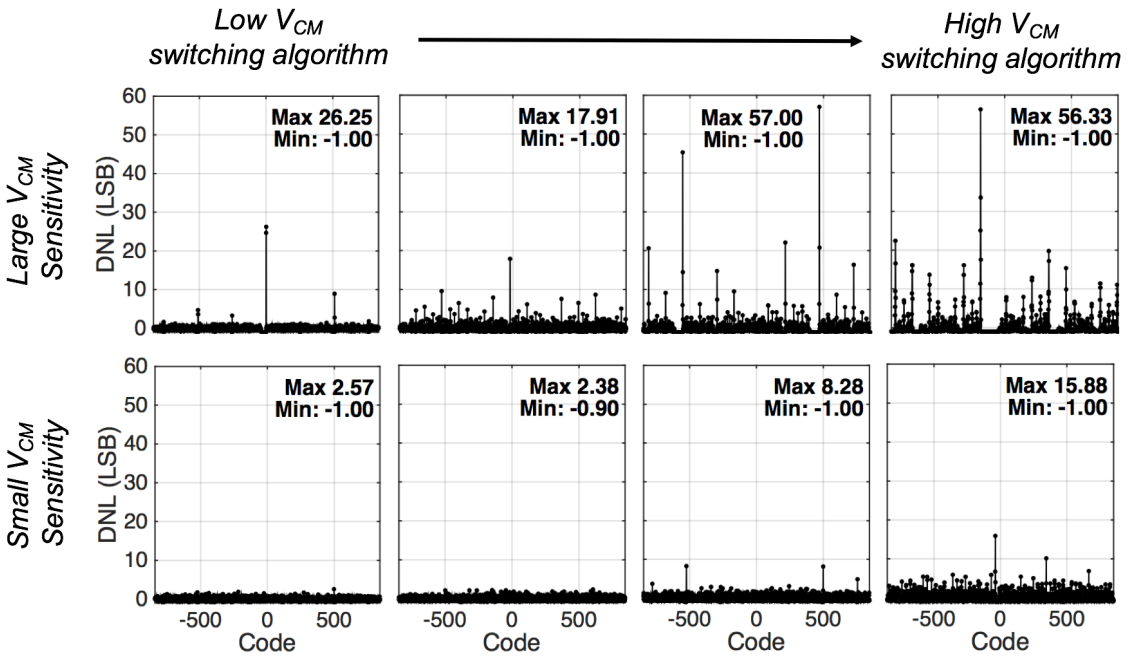

Figure 3.45. Measured DNL for various switching procedures. The top measurements use a comparator with high V_{CM} sensitivity, and the bottom measurements use a comparator with low V_{CM} sensitivity.

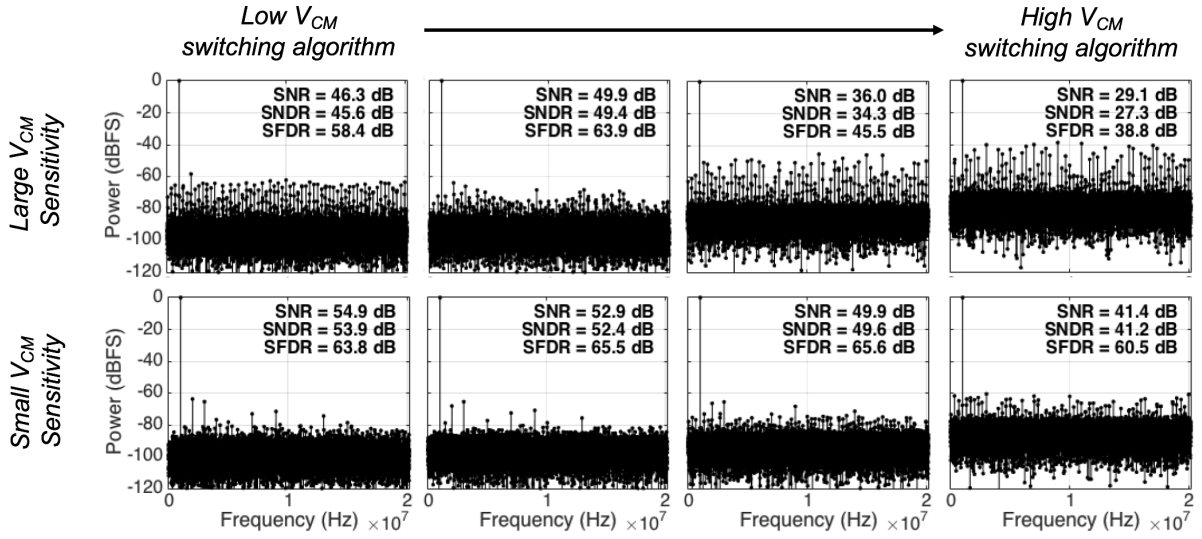

Figure 3.46. Measured FFTs for various switching procedures. The top measurements use a comparator with high V_{CM} sensitivity, and the bottom measurements use a comparator with low V_{CM} sensitivity.

verter's frequency response. The SFDR of the ADC using the high V_{CM} switching procedure and high-sensitivity comparator is limited to 38.8 dB, with an overall SNDR of 27.3 dB. In contrast, the SFDR and SNDR of the ADC using the low V_{CM} switching procedure and low-sensitivity comparator are 63.8 dB and 53.9 dB, respectively. Spurs in this configuration correspond to harmonics of the input tone, which is consistent with sampling switch distortion or DNL due to DAC mismatch. Figure 3.46 also illustrates that regardless of offset sensitivity, raising V_{CM} significantly increases the comparator's input-referred noise. From the low to high V_{CM} switching procedure, the noise floor grows by nearly 10 dB.

While the minimum V_{CM} switching procedure provides the lowest input-referred comparator noise, Fig. 3.46 shows that using the mid-range V_{CM} switching procedure can improve SNDR when the comparator's V_{CM} sensitivity is high. Figure 3.47 explains this by summarizing the SNDR achievable using the two comparator configurations as a function of switching procedure. The measurements are compared to the SNDR predicted by a MATLAB model incorporating the comparator's measured offset sensitivity to V_{CM} and simulated noise sensitivity to V_{CM}. The illustration on the horizontal axis shows the progression of V_{CM} during the first three SAR comparisons for each switching procedure; the leftmost scheme monotonically decreases, minimizing V_{CM}, while the rightmost scheme monotonically increases, maximizing V_{CM}. The SNDR of the ADC using a low-sensitivity comparator drops with increasing V_{CM}, suggesting that its resolution is thermal noise-limited. Using the high-sensitivity comparator, SNDR is maximized when the switching procedure minimizes V_{CM} variation, suggesting that is resolution is limited by offset fluctuation.

Overall, the MATLAB model and measured results in Fig. 3.47 match closely, except in the three lowest V_{CM} settings where the simulated SNDR is higher than the measurements. This shows that alternate sources of noise and distortion not captured in the simulation — such as DAC mismatch — restrict the achievable SNDR. Therefore, while comparator noise and offset sensitivity can significantly limit ADC performance, properly designing the comparator and selecting an appropriate DAC switching procedure can mitigate this impact.

3.4.4 Comparison to prior work

To compare the performance of this prototype to previously published ADC designs, Fig. 3.48 [81] shows the Walden figure of merit (FoM) as a function of SNDR for both designs published at ISSCC/VLSI through 2017 and performance-scalable ADC designs. These works are compared against the prototype ADC configured as either an oversampled converter (80 MS/s, 10 MHz bandwidth) or a Nyquist-rate converter (40 MS/s, 20 MHs bandwidth). While the FoM of the ADC operating at Nyquist is better than its FoM in oversampled mode, oversampling relaxes the receiver filter requirements to improve performance at the system level. Figure 3.48 shows that the converter's FoM is comparable to prior SAR ADCs with Nyquist sampling rates above 40 MS/s. The SAR architecture provides superior (lower) FoM than other architectures, though the peak achievable SNDR is lower than pipelined or oversampled $\Delta\Sigma$ converters.

The right plot in Fig. 3.48 similarly shows that the SAR architecture affords better FoM than other designs. While the prototype's FoM is higher than some previously published scalable SAR designs, the SAR ADCs with better FoM [61,63,64] operate at lower sampling rates ($\leq$4 MS/s) to afford improved energy efficiency. The table in Fig. 3.49 compares the performance of the prototype to the two previously shown power-scalable SAR designs

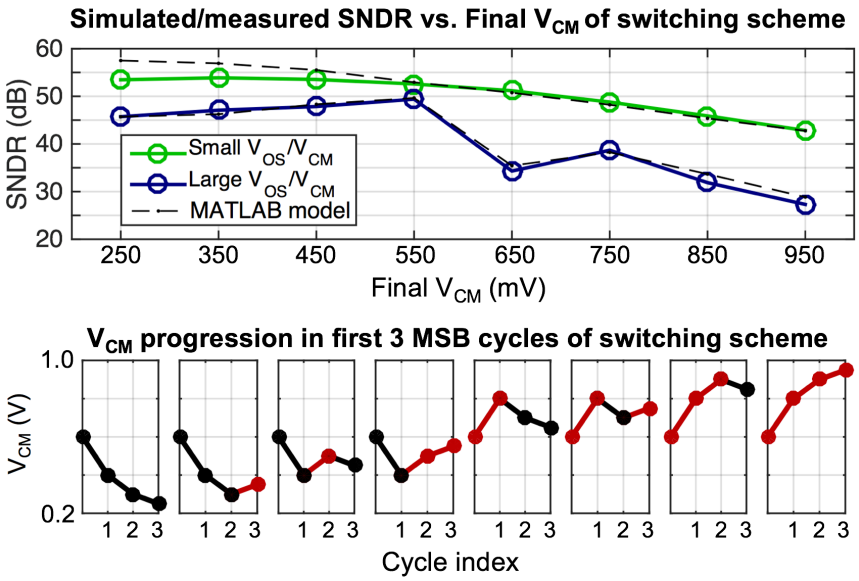

Figure 3.47. Measured and simulated SNDR as a function of switching procedure V_{CM} for high-sensitivity (blue) and low-sensitivity (green) comparator configurations.

described (note that [64] and [63] are variations of the same design). The work in [61] uses a low-leakage process technology and low supply voltage to minimize power consumption. The work in [64] uses a larger 90nm process technology node, which also has lower leakage. These process nodes are suitable for low-speed ADCs, but are not suitable for higher-speed wireless applications targeted in this work. The main degradation in achievable FoM is caused by the overhead power consumption described in the next section. The design presented here has no additional calibration circuitry, leading to a very small total area but also lower peak SNDR.

3.4.5 Analysis of scalability limitations

To identify some of the limitations of developing an energy-efficient, resolution-scalable converter, Fig. 3.50 compares the power and noise breakdown of the converter in the low-power, optimal FoM, and highest-resolution configurations. This illustrates that optimal energy efficiency is obtained when a component's power contribution is comparable to its noise contribution. For the prototype ADC, optimal FoM is obtained when the scalable DAC and comparator have the largest relative noise contribution. The ADC is inefficient in the maximum resolution setting because the comparator consumes 44.7% of the total power but contributes only 0.5% of the noise. In the lowest power setting, the majority of the converter's power consumption is set by components that do not scale — the fixed overhead power of the logic, comparator, and DAC — which also causes the converter to be inefficient.

The power breakdown from Fig. 3.50 is obtained via post-layout extraction simulations. Both the DAC and comparator power are divided into "scalable" and "fixed" components— the scalable portion scales exponentially with the number of bits (e.g. DAC capacitance switching energy), and the fixed portion either does not scale (e.g. leakage power) or scales at best linearly with the number of bits (e.g. minimum-sized comparator switching energy).

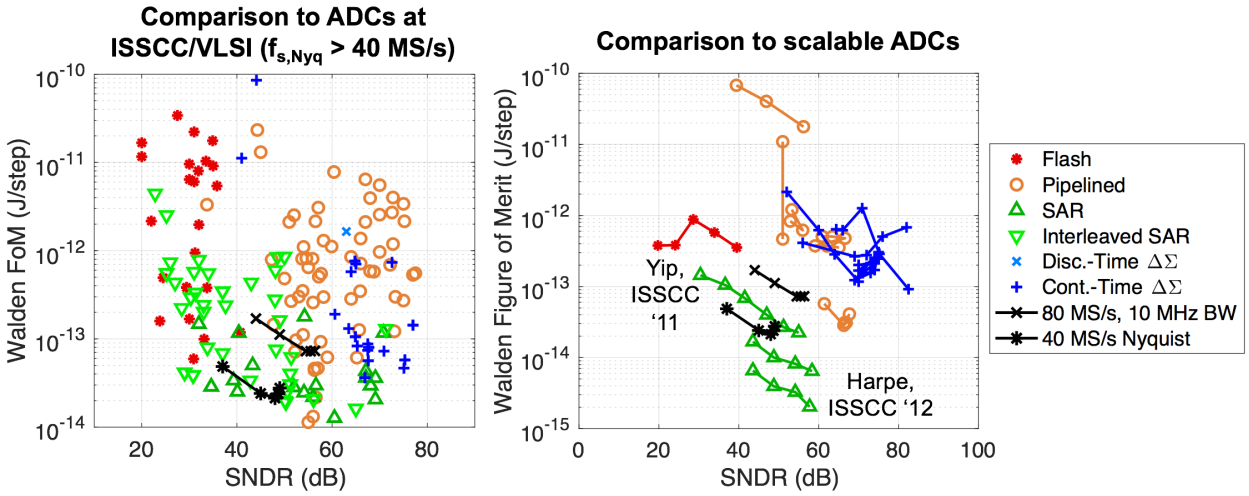

Figure 3.48. Comparison of figure of merit (FoM) to stand-alone designs published at ISSCC/VLSI [81] with Nyquist sampling frequency > 40 MS/s (left) and to previously published scalable converters (right).

	Yip, ISSCC '11			Harpe, ISSCC '12		This Work					
Architecture	SAR			SAR		SAR					
$f_{s,nyq}$ (MS/s)	0.06	0.02		4		20 (f_s=80 MS/s)			40		
SNDR (dB)	30.4	36.6	55.0	43.7	56	44.0	55.5	56.4	37.0	48.0	49.0
Power (μW)	0.234	0.116	0.206	8.22	17.44	440	705	790	225	350	510
Walden FoM (fJ/step)	143	105	22.4	16.4	6.5	170	72.4	73.2	97.2	42.6	55.4
Supply (V)	0.5	0.55		1.1		1.2 / 1.0					
Technology	65nm (low-leakage)			90nm		65nm (general purpose)					
Area (mm^2)	0.212			0.047 (w/ decoupling)		0.004					

Figure 3.49. Comparison to previously published scalable SAR ADC designs.

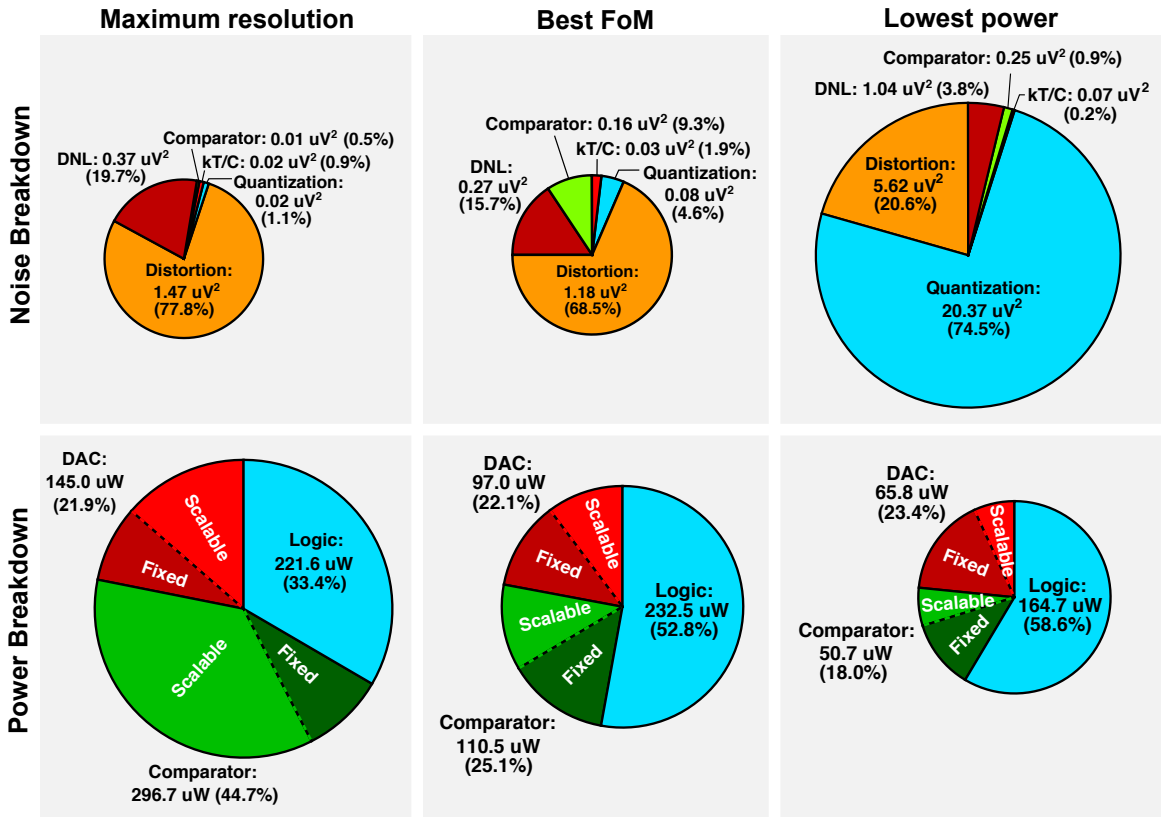

Figure 3.50. Breakdown of noise (top) and power consumption (bottom) between different aspects of the ADC for maximum resolution, best FoM, and lowest power configuration settings.

The noise breakdown in Fig. 3.50 is estimated from measurements and calculations. The total noise variance is calculated from the measured SNDR. The power of the largest harmonic tones in the FFT are used to calculate the distortion. Comparator noise is estimated via transient noise simulations of the post-layout extracted netlist, while sampling noise variance is calculated as $4kT/C_{DAC}$. Quantization noise variance is calculated from the peak differential input voltage $v_{in,pk}$ of 0.6 V and number of ADC bits N_{bits} as:

$$v_{n,quant}^2 = \frac{\frac{1}{2}v_{in,pk}^2}{10^{(6.02N_{bits}+1.76)/10}}$$

The remaining noise budget is allocated to differential nonlinearity (DNL) induced by capacitor mismatch and comparator offset variation.

From Fig. 3.50, it is clear that two non-scalable components of ADC noise — DNL and distortion — are major sources of SNDR degradation in this prototype. This fundamentally limits the achievable efficiency of the scalable ADC by dictating its maximum resolution. In future designs, the distortion can be reduced by lowering the sampling switch resistance or by adding a bootstrap circuit, though these modifications will increase the ADC's overhead power and reduce its minimum-resolution efficiency. While the sampling switch resistance can be scaled with the required resolution to minimize excess power consumption, the leakage power of devices used to drive larger switches will not scale. The large fraction of distortion due to DNL is likely due to capacitance mismatch. Redundant ADC decisions and calibration could help lower this noise, again at the expense of additional overhead power. Moreover, restricting the full-scale input range of the ADC using parasitic capacitance within each unit cell worsens the unit-element mismatch by lowering the size of the switched unit-element capacitance. Using a lower reference voltage to scale the ADC input range instead could improve mismatch, but requires generating an additional supply voltage.

Figure 3.50 also illustrates that the switching energy of the SAR logic consumes a substantial amount of power. To improve the converter's efficiency in low-resolution applications, future implementations of power-scalable SAR ADCs may require a more streamlined logic approach as presented in [79], which uses only separate comparators and delay cells. To avoid independently tuning comparator offsets, this design could be modified to include a single main comparator or preamplifier. Figure 3.50 also reiterates that the noise of the minimum-sized power-scalable comparator is sufficiently low that comparator noise is never a significant fraction of the total ADC noise. To address this issue, a stand-alone low-power comparator could be incorporated into the design for low-resolution operation. In the DAC, leakage current drawn by unused buffers consumes a fixed amount of power. Because the buffers must be large enough to settle the DAC capacitance within an iteration of the SAR algorithm, little can be done to reduce this overhead. However, redundant conversion steps used to improve the converter's DNL can also relax the timing requirements on the DAC buffers, enabling the use of smaller, lower-leakage devices. Overall, many of the limitations on power and resolution scalability highlighted in Fig. 3.50 can be addressed to improve the efficient scaling range of the converter.

3.5 Conclusion

This work illustrates that while circuit-level design configurability can afford power savings, the efficient tuning range is restricted by both power overhead (e.g. fixed logic power con-

sumption) and resolution limitations (e.g. DAC mismatch requiring calibration). The ADC will be most efficient in the operating range where thermal noise or mismatch dominates. These are components that allow the ADC power consumption to be reduced exponentially with resolution, maintaining a constant Walden FoM. In the quantization noise-limited regime, linear power scaling of the digital logic limits the achievable resolution. As a result, any design component which increases the power consumption of logic (e.g. level shifters to move between digital and analog supplies) will reduce the power efficiency that the converter can achieve. While the total DAC capacitance can nominally be reduced by a factor of two for each bit reduction in required resolution, this only applies to the portion of power consumed by switching the capacitor array. Any overhead associated with generating the DAC control signals will either scale linearly with the number of bits (e.g. comparator result decoding), or remain a fixed power overhead (e.g. binary-to-thermometer encoding circuitry). To maintain efficiency, the control logic power overhead, which scales only linearly with the number of bits, must be minimized. This is easier to accomplish in scaled process technology nodes.

Similar to the way in which overhead power consumption dictates the minimum power consumption of the converter regardless of resolution, any fixed sources of noise or distortion will limit the maximum achievable resolution of the ADC. For instance, one major source of SNDR degradation in high to moderate speed applications is distortion due to sampling switch nonlinearity. The size of the sampling switches should scale with the DAC capacitance size, or include a bootstrapping circuit to maintain an input voltage-independent switch resistance. Given that capacitor mismatch is constrained by the LSB size, a mismatch-constrained DAC scalability technique can only cut total DAC capacitance by a factor of two for each bit reduction in required resolution. Either the DAC should not be sized for thermal noise limitations, or a robust calibration scheme should be in place to tolerate this mismatch. Calibration will introduce additional power overhead that increases the minimum power consumption of the converter. Maintaining performance efficiency over a broad tuning range is difficult, as the additional circuitry required to increase the maximum achievable performance introduces overhead that degrades the efficiency of the converter in low-resolution operation.

Chapter 4

Architecture-configurable RF-to-digital receiver

The previous chapter explores the limitations of designing a performance-configurable ADC for a conventional ADC-last receiver architecture. This chapter describes an alternative digital-intensive receiver that uses a reconfigurable architecture to achieve a broader range of performance tunability for improved power savings. Additionally, the CMOS scaling-friendly nature of the proposed receiver can enable higher levels of design complexity for fully integrated systems-on-a-chip. Directly digitizing the RF input signal also enables more signal processing to be performed in the digital domain, coming closer to the software-defined radio paradigm proposed over two decades ago [33].

As discussed in Chapter 2, a conventional radio receiver front-end includes a low-noise amplifier (LNA), mixer, and active filtering prior to the ADC. The LNA provides impedance matching and provides gain that relaxes the design requirements of subsequent stages in the receiver, but is often difficult to optimize. The proposed design eliminates an explicit LNA and active filter in favor of a hybrid successive approximation (SAR) and voltage-controlled oscillator (VCO) based ADC that provides both high linearity and sensitivity competitive with alternative RF-to-digital receiver designs. The receiver can be configured in a hybrid SAR+VCO mode, a SAR-only mode, or a VCO-only mode to scale linearity and sensitivity without the need for an explicit programmable gain amplifier or variable attenuator. Moreover, the dynamic range is programmable in order to afford power savings when large blockers are not present. Reconfigurable discrete-time filtering is implemented in the SAR sampling process, eliminating the need for an active gain stage and affording a high degree of digital filter configurability.

This chapter explores the design and implementation of this architecture-configurable RF-to-digital receiver. To provide context, prior efforts to realize software-defined radio are first discussed in Section 4.1. Subsequently, Section 4.2 discusses the major design considerations required to implement the hybrid SAR and VCO-based ADC topology, providing an overview of how filtering is integrated into the SAR sampling process and how the VCO is designed to minimize input-referred noise. Subsequently, specific design details of the 16nm CMOS prototype are provided in Section 4.3. Measurements of this prototype are presented in Section 4.4 before Section 4.5 reviews final conclusions.

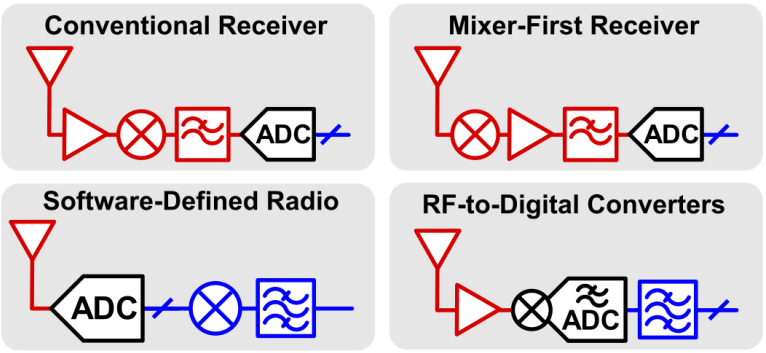

Figure 4.1. General comparison of typical receiver architecture (top left), mixer-first receiver architecture (top right), fully software-defined radio receiver (bottom left), and RF-to-digital receivers (bottom right).

4.1 Prior software-defined radio work

The concept of software-defined radio, in which a high-performance ADC is used to directly digitize the RF input signal and implement filtering and down-conversion in the digital domain, has existed for decades [33, 34]. Figure 4.1 compares the software-defined radio architecture against some alternative receiver designs. While the high dynamic range and fast sampling rate required of the ADC has limited the progress made towards realizing fully software-defined radio receivers, partially software-defined RF-to-digital receiver architectures and wideband, digitally-configurable mixer-first receiver architectures have emerged to support this vision.

This section describes some general efforts to implement direct RF-to-digital receiver designs, as well as some of major innovations that serve as the foundation for the proposed architecture. Section 4.1.1 in particular describes some of the receiver and ADC topologies that can facilitate direct RF-to-digital conversion. Section 4.1.2 provides a brief overview of the mixer-first receiver architecture, which can provide widely tunable RF performance. Finally, Section 4.1.3 describes the integrated SAR filtering approach that has been integrated into conventional receiver designs to provide additional passive baseband filtering within an ADC.

4.1.1 Direct RF-to-digital receiver design

Direct RF-to-digital converters can eliminate the need for conventional analog baseband filtering to implement blocker rejection in the digital domain or directly within an ADC. In some cases, even RF down-conversion can be integrated into the sampling process of the ADC without the need for a separate mixer. Oversampled ADCs are typically used to construct RF-to-digital receivers, which allows some filtering to be integrated in the loop filter of a $\Delta\Sigma$ converter and further filtering to be performed in the digital post-processing of the input signal. Alternatively, high-speed, high-resolution ADCs have emerged in recent years that come closer to enabling software-defined radio.

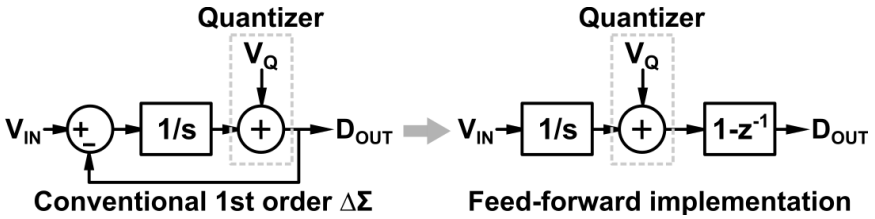

Figure 4.2. Comparison between a typical feedback-based first-order $\Delta\Sigma$ converter (left) and its equivalent feed-forward implementation (right).

4.1.1.1 $\Delta\Sigma$ *based topologies*

The most common architecture for direct RF-to-digital receivers is a $\Delta\Sigma$ based topology. In a conventional $\Delta\Sigma$ ADC, the difference between the input signal and a coarse digital approximation of the signal is fed to a (typically low-pass) loop filter, as previously shown in Fig. 3.2. Because the quantization noise is introduced at a different location than the input signal, the signal and noise experience different transfer functions. When a low-pass loop filter is used, the quantization noise is pushed to higher frequencies while the signal remains unchanged at low frequencies, improving the signal-to-quantization noise ratio (SQNR) within a small signal bandwidth. As a result, these converters can tolerate high levels of quantization noise (relaxing the required resolution and accuracy of the quantizer) while still achieving high SQNR, provided the sampling rate of the ADC is high.

In a $\Delta\Sigma$ based wireless receiver, the RF input signal is amplified and combined with a feedback DAC to generate the difference between the analog input and the digital approximation. Due to the oversampled nature of a $\Delta\Sigma$ ADC, this down-conversion can operate at the RF carrier frequency (f_{LO}) to integrate the feedback directly with the mixer and generate an error signal at baseband frequencies [82–84]. This front-end is typically followed by additional integrator and feedback stages in order to construct an appropriate loop filter that can provide both quantization noise shaping and blocker filtering. The highly oversampled nature of the converter relaxes the filtering requirements, as blockers can be filtered digitally provided the dynamic range of the converter is sufficiently large. In [82] and [83], passive integrators are used to provide high linearity and dynamic range at the expense of noise figure. Comparatively, improved noise figure is obtained in [48, 85] using active integrators and a carefully designed loop filter that provides additional blocker filtering, at the expense of design complexity and power consumption.

Overall, this $\Delta\Sigma$ based direct down-conversion technique can enable high linearity and dynamic range. The design is typically integrated with a wideband LNA and can support a broad range of carrier frequencies in order to easily configure the operating frequency of the receiver. However, the implementation of this receiver architecture still relies on an ability to design high-precision feedback DACs and an accurate integrator and summing stage, which can make it challenging to implement in scaled process technologies.

4.1.1.2 *VCO based topologies*

An alternative $\Delta\Sigma$ converter architecture uses a voltage-controlled oscillator (VCO) based ADC immediately after a mixer and LNA. A VCO-based ADC operates as a feed-forward implementation of a first-order $\Delta\Sigma$ converter, recognizing that the feedback loop of a $\Delta\Sigma$

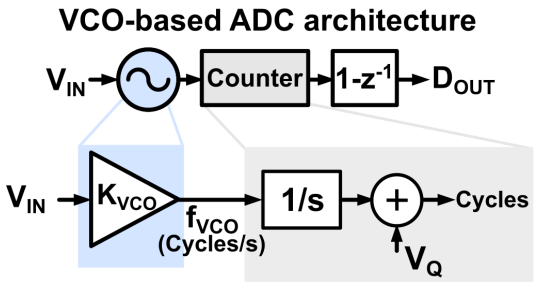

Figure 4.3. VCO-based ADC concept.

design can be rewritten as the equivalent feed-forward model illustrated in Fig. 4.2. To implement this structure, a VCO translates the input voltage into a frequency, which can then be integrated and digitized using a counter (or alternative phase quantizer), as shown in Fig. 4.3. The counter output can be differentiated in the digital domain to complete the feed-forward $\Delta\Sigma$ converter. Because this ADC topology requires only a voltage-controlled oscillator, counter, and digital subtractor, it is highly suitable for scaled CMOS process nodes. Prior work has demonstrated that the full receiver can be implemented in a hardware description language, though an active LNA and mixer is still required [86].

The main limitation of the VCO-based ADC technique is the highly nonlinear voltage-to-frequency transfer characteristic of the VCO. This makes the receiver topology suitable for small input signals, but linearity-constrained as the input amplitude increases. Various calibration techniques exist to correct for these nonlinearities, from digital calibration techniques [59, 87, 88] to carefully tuning the gain control from the LNA to manage the VCO nonlinearity [89]. Different modulation schemes can be used to relax these linearity constraints, such as interpolating between only two VCO frequencies during operation [90]. Alternatively, VCO-based ADCs have been used as digital-intensive quantizers in more conventional $\Delta\Sigma$ converters [91]. For wireless receiver applications, the low linearity makes it difficult for the receiver to operate in the presence of large blockers. The natural sinc response of the VCO-based ADC can be used to provide some inherent blocker rejection [89], but the poor in-band linearity of this architecture may not be suitable for digital beamforming applications with large in-band blockers.

4.1.1.3 *High-speed ADCs*

A fully digital radio baseband requires an ADC with dynamic range high enough to both detect small signals and tolerate large blockers, and a sampling rate high enough to convert the full RF spectrum so that down-conversion and channel selection can be performed in the digital domain. As discussed in Chapter 2, the dynamic range requirements are dictated by the blocker mask outlined in a particular wireless standard. After incorporating margin for modulation schemes and estimation, an ADC of at least 60 dB dynamic range would be desirable for software-defined radio systems. The converter must be able to maintain this resolution even at fast sample rates. In a fully software-defined radio where down-conversion occurs in the digital domain, the Nyquist sampling theorem requires that the ADC sample rate must be at least twice the maximum supported RF carrier frequency. This requires sample rates in the GS/s range to support 0-2.4 GHz wireless systems.

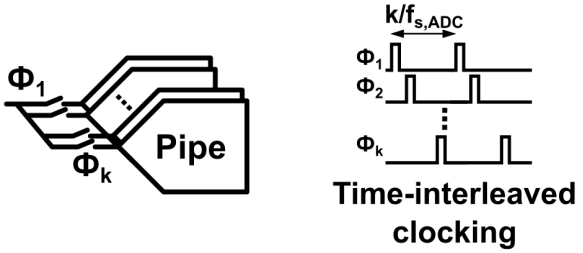

Figure 4.4. Sample time-interleaved pipelined ADC architecture.

While constructing ADCs to meet the performance requirements of software-defined radio systems is difficult, advances in CMOS scaling have made it possible for time-interleaved pipelined ADCs (illustrated in Fig. 4.4) to achieve performance suitable for these applications. As described in Section 3.1, pipelined converters can enhance the achievable speed and resolution of an ADC by splitting the conversion into many smaller stages. Time interleaving allows a high-speed converter to be constructed from many lower-speed ADCs. The work in [45] uses a pipelined ADC architecture to achieve 56 dB SNDR at a sampling rate of 4 GS/s, while the work in [46] uses a time-interleaved pipelined-SAR design to achieve 57.3 dB SNDR at 4 GS/s. The time-interleaved pipelined design in [47] achieves 58 dB SNDR at 5 GS/s, suitable for 2.4 GHz Wi-Fi applications.

This level of performance comes with a substantial power overhead. To achieve high performance in existing designs, calibration techniques are required correct for errors introduced through time interleaving, providing improved performance at the expense of power consumption. The ADCs in [45], [46], and [47] draw 300 mW, 513 mW, and 709 mW, respectively, relative to an optimized wireless receiver that would typically consume under 100 mW. These ADCs would also require an LNA to detect RF input signals below -50 dBm, given that the SNDR reported by those designs is characterized assuming a full-scale input signal of over 1V (+4 dBm). As process technologies shrink further to afford faster sampling rates and more complex calibration schemes with lower power overhead, these converters may become increasingly suitable for wireless applications. However, existing designs are not suitable for low-power receivers.

4.1.2 Mixer-first receiver design

The presence of a radio frequency low-noise amplifier (LNA), which is often designed carefully to provide impedance matching at a specific frequency, presents a challenge in making conventional receivers suitable for wideband operation. As an alternative, mixer-first receivers [12, 92] eliminate the explicit RF LNA stage by placing a passive mixer immediately after the antenna, using a baseband amplifier in feedback as shown in Fig. 4.5 to provide input matching, filtering, and amplification. Because the RF carrier frequency is determined by the mixer clock frequency instead of the LNA response, this architecture can support wideband performance. Moreover, the required low-pass response of the baseband amplifier is typically easier to implement than a wideband LNA, which needs to support GHz bandwidths.

While mixer-first receivers can achieve wideband operation using a relatively simple architecture, design techniques must be applied to reject unwanted signals at multiples (har-

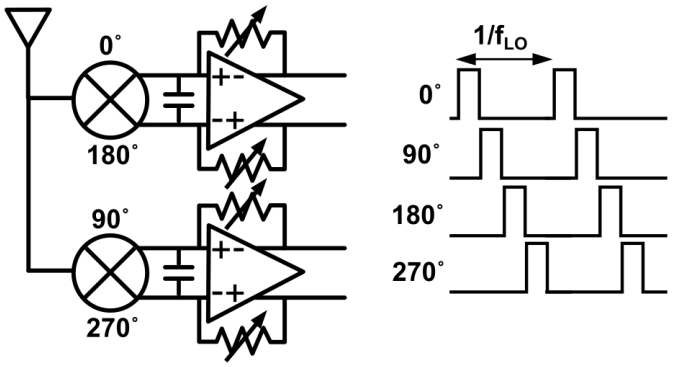

Figure 4.5. Sample 4-phase mixer-first receiver architecture.

monics) of the LO frequency. As illustrated in Fig. 4.5, multiple copies of the receiver are typically placed in parallel and driven by different LO phases to support harmonic rejection. The example shown in Fig. 4.5 illustrates a four-phase mixer-first receiver design suitable for canceling the third LO harmonic; higher levels of blocker rejection can be obtained by including more phases, as described in [12] and [93].

Finally, mixer-first receivers require a tradeoff between noise figure and design complexity. Though mixer-first receiver designs often have higher noise figure (> 5 dB) than LNA-first architectures, careful baseband amplifier design [94] or noise cancellation techniques [14, 95] can be applied to obtain < 3 dB noise figure. Because placing a passive mixer at the receiver input can provide higher linearity than a typical LNA-first design, carefully-designed mixer-first receivers can be competitive with high-performance wireless systems [14, 94]. However, though mixer-first architectures eliminate the need for an RF LNA, they typically still require carefully designed feedback amplifiers that can be difficult to implement in scaled process technologies using low supply voltages.

4.1.3 Filtering SAR ADC

As an alternative to fully redesigning a conventional receiver, ADCs with integrated filtering can simplify receiver design by relaxing the required baseband filter order. In a typical SAR ADC, the input signal is sampled and held on the capacitive DAC at the ADC sampling frequency, as shown in Fig. 4.6; the input buffer and sampling switches are designed to ensure that the input signal fully settles during the sample period. Because no frequency-dependent attenuation is incorporated into the sampling network, high-frequency tones may fully alias into the signal bandwidth. Alternatively, the work in [24,25] proposes a specialized SAR ADC sampling scheme that implements discrete-time filtering. As shown in Fig. 4.6, this is accomplished by sampling the input signal onto many smaller capacitors prior to the ADC conversion. Before executing the successive approximation algorithm, a set of switches is closed to calculate the weighted average of these samples using charge sharing, which gives the output of a finite impulse response (FIR) filter.

While this integrated filtering technique is conceptually simple, it increases design complexity, area, and power consumption. Because the FIR filter response must attenuate signals at multiples of the ADC sampling rate, the sampling frequency must be carefully chosen to optimize rejection for a particular filter configuration. Practically speaking, time must be

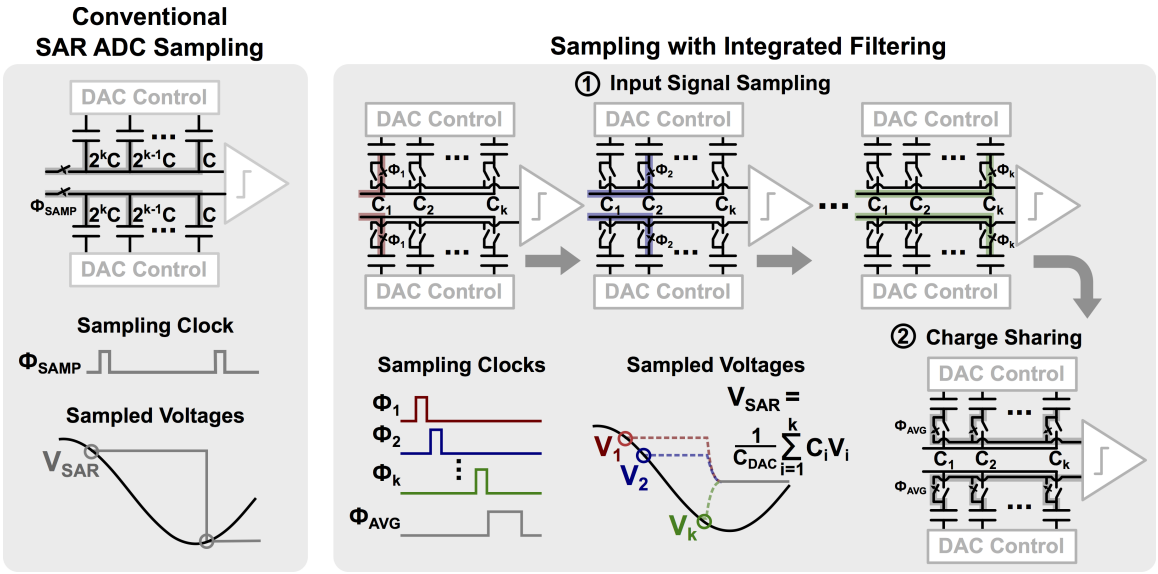

Figure 4.6. Comparison between conventional SAR ADC sampling (left) and FIR/SAR sampling (right).

allocated to execute the SAR algorithm, which favors placing multiple DAC banks in parallel to support a faster sample rate. This translates to a substantial area overhead in designs where the DAC capacitance must be large to minimize thermal sampling noise. Finally, the maximum achievable filter rejection is limited by non-idealities such as charge injection and capacitance mismatch. Despite these challenges, the integrated filter affords highly reconfigurable filtering well-suited to implementation in scaled CMOS. In this work, the filter topology is utilized to provide anti-alias filtering in an RF-to-digital converter. A more detailed discussion of this architecture's limitations and the FIR filter design procedure is described in the next section.

4.1.4 Limitations of prior work

While prior RF-to-digital receivers have been proposed to implement software-defined radio, they are often either difficult to implement in scaled CMOS process technologies (e.g. $\Delta\Sigma$ RF-to-digital receivers), performance-limited (e.g. VCO-based ADCs) or consume a significant amount of power (e.g. high-resolution RF bandwidth ADCs). Mixer-first receiver architectures have gained popularity in the past decade as an analog-intensive approach to configurable radio design that eliminates the need for a high-performance LNA. However, such receivers still require carefully designed baseband amplifiers that may be difficult to build in deeply scaled process nodes with complex design rules and small supply voltage headroom.

To address these issues, this work proposes an alternative RF-to-digital receiver architecture that combines SAR and VCO-based conversion techniques to achieve scaling-friendly operation without sacrificing dynamic range. Using a mixer-first architecture, the proposed receiver uses a SAR ADC with integrated filtering to afford high linearity and a VCO-based ADC to convert the small-amplitude SAR residue as proposed in [96]. With this design, the receiver architecture can also be configured to trade performance for power consumption:

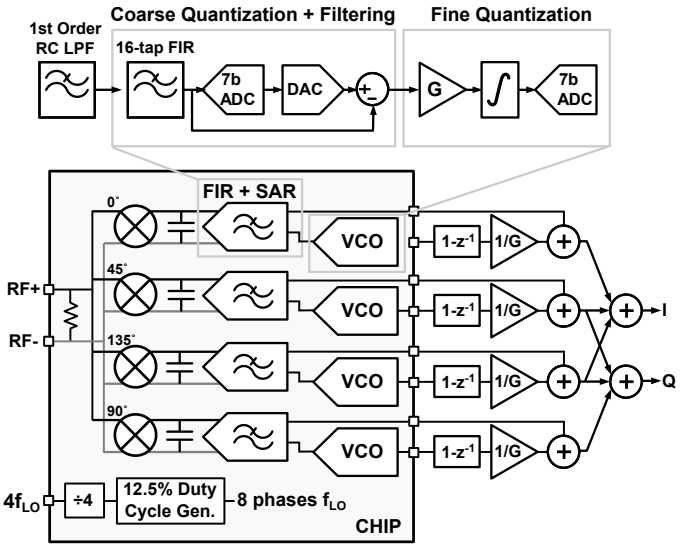

Figure 4.7. Diagram of implemented RF-to-digital receiver topology.

both the SAR ADC and VCO can be enabled to maximize dynamic range, the VCO-based ADC can be used by itself to improve sensitivity when high dynamic range is not required, and the SAR ADC can alternatively be used for high-linearity, high-noise figure applications. The next section discusses the design of the core building blocks of this architecture — namely, the SAR ADC with integrated filtering and the VCO-based ADC — in further detail.

4.2 Design considerations

For compatibility with advanced CMOS process technologies, the design uses both SAR and VCO-based ADC architectures to eliminate the need for high-performance analog amplifiers. The proposed receiver consists of an 8-phase mixer-first receiver with four differential-input sub-ADCs, as illustrated in Fig. 4.7. The multi-phase mixer enables harmonic rejection and noise averaging between sub-ADCs, but each sub-ADC can be disabled to facilitate power savings when harmonic rejection is not required and higher noise is tolerable. Each sub-ADC consists of two stages. The first stage incorporates a coarse-resolution 7-bit SAR ADC with FIR filtering integrated into the sampling process. The capacitive input sampling provides both high blocker tolerance and digitally-configurable filtering. The fine-resolution VCO-based ADC in the second stage operates as a high-resolution first-order $\Delta\Sigma$ ADC, using the integrating properties of a VCO to eliminate the need for an active analog integrator. The FIR+SAR stage can be bypassed to lower power consumption and noise when high linearity is not required, or the VCO-based ADC can be disabled if high noise figure is not required.

This section elaborates on the design of core components of the proposed architecture-configurable RF-to-digital receiver. To be suitable for wireless applications, the integrated SAR/FIR filter must be constructed to provide substantial anti-alias filtering, and the effective input-referred noise of the full ADC must be low to boost the receiver's sensitivity. Section 4.2.1 first discusses the implementation of the FIR filter integrated within the SAR ADC, providing an overview of the filter topology choice and its implications on the ar-

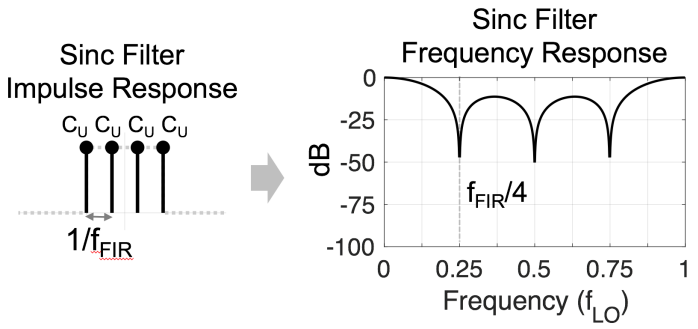

Figure 4.8. Example 4-tap sinc filter impulse response (left) and frequency response (right).

chitecture of the receiver. It also explores the effect of component mismatch and circuit non-idealities on the achievable filter rejection. Subsequently, Section 4.2.2 discusses the noise and linearity of a VCO-based ADC.

4.2.1 Filtering SAR ADC design

As outlined in Chapter 2, a receiver must be able to attenuate any blockers at multiples of the ADC sampling frequency that fold into the signal bandwidth, and other interferers must still be attenuated to fit within the dynamic range of the ADC. As presented in [24, 25], the capacitive sampling process of a SAR ADC can be configured to integrate discrete-time filtering. While prior work uses this technique to simplify the filter design in a conventional LNA-first receiver, this work samples the RF input signal directly onto the filtering SAR ADC concurrently with passive mixing. Without an additional active filtering stage, this sampling technique must provide most of the receiver's alias rejection. To explore how significant passive filtering can be obtained, this section first discusses how the FIR filter weights can be chosen to implement anti-alias filtering and how this impacts the ADC implementation. From there, it discusses how mismatch and sampling non-idealities may impact the performance of the filter. Finally, the section concludes by discussing how the SAR ADC design affects the achievable sensitivity of the receiver.

4.2.1.1 *FIR filter response*

The passive mixer provides a modest single pole of filtering, but more attenuation is required to tolerate large blockers. As a result, the integrated FIR filter must have a very sharp anti-aliasing response. Unlike a conventional low-pass or band-pass filter, an anti-alias filter must place notches at all frequencies that will fold into the signal bandwidth. A natural topology choice to address this requirement is a sinc filter, as illustrated in Fig. 4.8, which places nulls at frequency intervals determined by the total sample duration.

The sinc filter is easy to implement using the proposed sampling scheme due to its simple boxcar impulse response. For instance, an N tap FIR filter can be constructed from N capacitors of equal weights by sampling the input signal onto each of these capacitors sequentially prior to the SAR ADC conversion. More generally, if the filter is constructed from N FIR filter taps sampled at a frequency f_{FIR} so that the total sample duration is $t = N/f_{FIR}$, the corresponding filter will place nulls at f_{FIR}/N. The main challenge in applying this sampling scheme is that practically speaking, additional time must be allocated

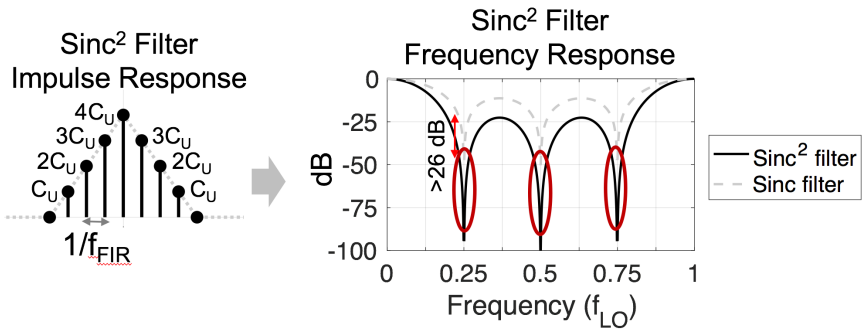

Figure 4.9. Example 4-tap sinc^2 filter impulse response (left) and frequency response (right).

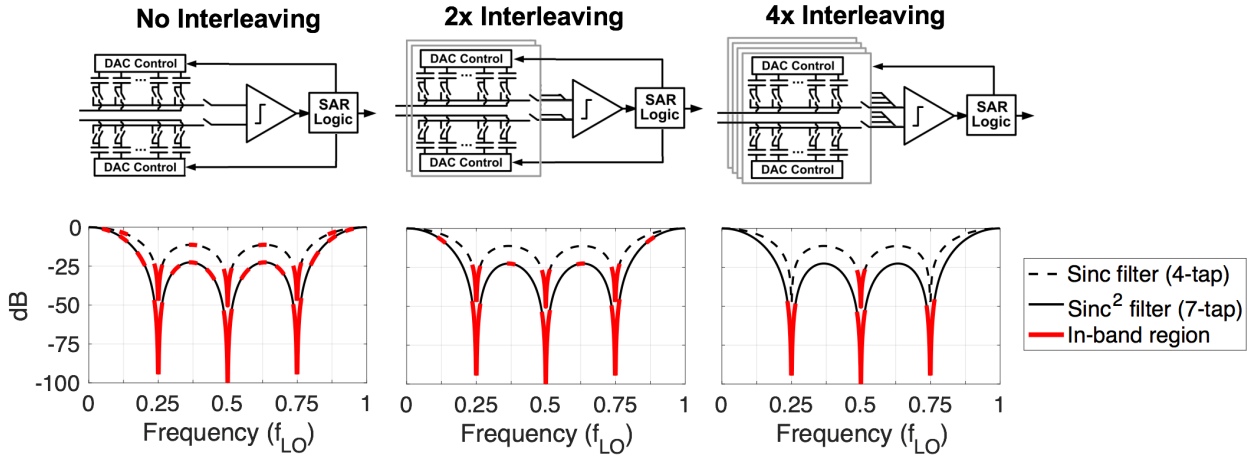

Figure 4.10. Effect of DAC interleaving on frequencies that alias in-band for sinc and sinc^2 filter configurations.

to the SAR ADC conversion. For instance, if an additional N clock cycles were allocated to the conversion, the slower ADC sample rate of $f_{FIR}/(2N)$ would alias frequency bands that are not significantly attenuated by the FIR filter. Therefore, as discussed in [24, 25], multiple DAC banks can be interleaved to support a faster sample rate at the expense of area, complexity, and power. The work in [97] also discusses the implementation of discrete-time switched-capacitor filters in further depth.

Two interleaved DAC banks are sufficient for the work in [24, 25], which uses the filtering SAR ADC as the final stage of a conventional receiver with additional active filtering preceding the ADC. In this work, however, the passive FIR filter must provide the bulk of the receiver's alias rejection. As a result, improved anti-alias filtering can be integrated into the ADC using a higher degree of interleaving. As shown in Fig. 4.9, a sinc^2 filter can be used to improve the achievable rejection by over 20 dB. This can be implemented using 4-way interleaving, illustrated in Fig. 4.10. This level of interleaving is used in the prototype design to balance design complexity (area and power overhead due to incorporating multiple capacitor banks) and achievable filter rejection. In general, a filter of order sinc^N can be incorporated by using a factor of 2^N interleaving in the ADC.

One key advantage of the digitally-configurable filter is the flexibility to place nulls specifically to target large blockers, as alternative filter topologies can be used to place

notches closer to the carrier. If the digital baseband can detect the presence of strong blockers (e.g., by comparing the measured signal with and without the anti-alias filter), a digital adaptation loop could be integrated with the design to modify the tap weights to minimize the dynamic range of the receiver and its in-band interference. This work has not studied the power and area overhead associated with such back-end processing.

4.2.1.2 *Noise considerations*

In a conventional SAR ADC, the input-referred noise and offset of the comparator significantly impact the achievable resolution of the converter. In this application, however, only the accuracy of the voltage residue held on the DAC after the SAR conversion affects the overall resolution. Incorrect comparator decisions due to input-referred noise or offset fluctuations may cause this residue voltage to exceed the size of the converter's nominal LSB, but will not degrade the ADC resolution if this residue does not saturate the input range of the fine-resolution VCO stage. As a result, the main source of noise contributed by the SAR ADC will be thermal (kT/C) sampling noise. Moreover, an accurate calibration scheme is required to compensate for DAC mismatch, given that the SAR conversion residue must be accurately estimated to within the desired resolution level of the overall converter.

The modified sampling procedure, in which the full DAC capacitance is split into smaller unit cells sampled at programmable time instants, does not alter the total thermal noise in each sample. This is because the higher noise is averaged via charge sharing prior to conversion. To derive this result, we can consider the total thermal noise generated by sampling onto N capacitors of size C_{DAC}/N, and then averaging the sampled noise via charge sharing. In this scenario, the variance (σ^2) of thermal noise sampled onto each unit capacitor is $\sigma^2(v_{n,samp}) = NkT/C_{DAC}$. Using the fact that the variance of the sum of uncorrelated random variables is the sum of their individual variances, we can calculate the variance of the mean of these samples:

$$\sigma^2\left(v_{n,tot}\right) = \sigma^2\left(\frac{1}{N}\sum_{i=1}^{N} v_{n,samp}\right) = \frac{1}{N^2} \times N \times \sigma^2\left(v_{n,samp}\right) = \frac{N^2}{N^2}\frac{kT}{C_{DAC}} = \frac{kT}{C_{DAC}}$$

Therefore, the total thermal sampling noise-limited DAC capacitance (C_{DAC}) is a function of the target minimum detectable signal level ($P_{min,dBm}$), termination impedance (R_{match}), target signal bandwidth (f_{BW}), ADC sampling frequency ($f_{s,ADC}$) and other sources of noise in the converter ($v^2_{n,other}$), including the input-referred noise of the VCO-based ADC. The relationship between these variables is summarized in the following expression:

$$P_{min,dBm} = 10\log_{10}\left(\left(\frac{4kT}{C_{DAC}} + v^2_{n,other}\right)\left(\frac{1}{R_{match}}\right)\left(\frac{1}{1mW}\right)\left(\frac{2f_{BW}}{f_{s,ADC}}\right)\right)$$

If half of the thermal noise budget is allocated to sampling noise, $R_{match} = 50\Omega$, and a modest oversampling ratio of 4 is used ($f_{s,ADC}/2 = 4f_{BW}$), C_{DAC} must be 12 pF to maintain $P_{min,dBm} = -78$ dBm for each sub-ADC. Allocating half of the noise to other sources provides margin for the input-referred noise of the VCO, trading off between the required DAC capacitance and the static power draw of the VCO. If the SAR ADC provides 7 bits of coarse resolution, each of the 128 unit cells in the DAC requires a unit capacitance of roughly 50 fF, which is much larger than a conventional design. To provide more insight,

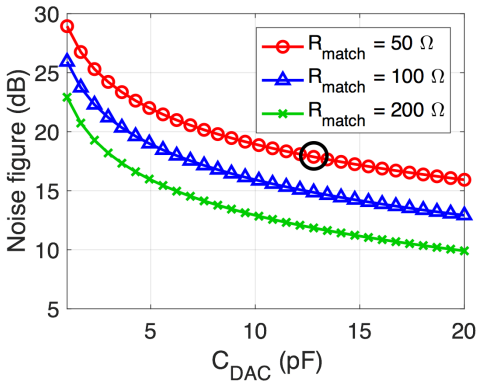

Figure 4.11. Expected best-achievable noise figure due to DAC sampling noise vs. DAC size.

Fig. 4.11 summarizes the kT/C noise-limited noise figure achievable from the receiver as a function of matching resistance and DAC size, assuming a 1700 MHz LO frequency and 10 MHz signal bandwidth. The circled point marks the configuration used in this prototype, demonstrating that kT/C noise alone limits the achievable noise figure to 17 dB. Adding passive gain by increasing the matching resistor size and using a transformer at the receiver input can help improve noise figure, subject to the feasibility of constructing a transformer with a large turns ratio.

Without an active gain stage, the thermal noise sampled by the SAR ADC fundamentally limits the achievable sensitivity of the receiver. This degradation in noise figure may be acceptable if the SAR ADC is used only in large blocker scenarios requiring high linearity and dynamic range, but not low sensitivity. If this is not acceptable, either the DAC capacitance or the ADC sampling rate must be increased to boost sensitivity without an active amplification stage. For each 3 dB improvement in sensitivity, either the capacitance or the sampling rate must be doubled, resulting in an exponential growth in either area (larger C_{DAC}) or power consumption (higher $f_{s,ADC}$). Increasing C_{DAC} will also correspondingly increase the power consumed by switching the DAC array, though a low-power switching algorithm can be used to ensure that this is a relatively small fraction of the total receiver power.

The typical disadvantage of using a large DAC capacitance in a SAR ADC is the power consumed by the buffer used to drive this capacitance to a high degree of accuracy at a fast sample rate. This assumes the ADC is sampling at Nyquist, and that the driver must settle each new sample fully to within half of the full LSB during the sample period. However, in a mixer-first receiver design, a large baseband capacitor (up to 100s of pF) is typically used to provide a narrow (10-20 MHz) cutoff frequency using the translational first-order pole set by the RC time constant of this baseband sampling capacitor and the small (50Ω) matching resistor. In this case, the large baseband capacitor can be used to store the bulk of the signal charge, while charge sharing samples the input signal onto the smaller fixed capacitor in the DAC with minimal attenuation and without requiring an amplifier with a high gain-bandwidth product.

One additional aspect of SAR ADC performance that may influence the combined achievable SAR+VCO resolution is kickback noise from the comparator. While this noise

can couple back onto the SAR ADC residue voltage, it is significantly attenuated by the capacitive divider ratio between the comparator's input pair and the large DAC capacitance. For instance, if the gate to source capacitance (C_{GS}) of the comparator's input devices is 1 fF, a 6 pF single-ended DAC capacitance will attenuate a 0.8V swing at the comparator output to 133μV. If the full-scale input range of the comparator is 600 mV (corresponding to a 0 dBm input signal), this voltage step corresponds to 0.9 LSBs. Because the comparator does not need to be sized for low noise, small input devices can be used that will minimize the amount of kickback. Alternatively, the comparator kickback noise can be reduced using a multi-stage topology with preamplifier such as the double-tail comparator described in the previous chapter. A final aspect of SAR ADC performance that may influence the overall resolution is mismatch, which will be discussed in the next section.

4.2.1.3 *Mismatch effects*

Both the SAR residue accuracy and the achievable FIR filter rejection depend on the accuracy of the unit capacitors in the DAC, which are susceptible to mismatch due to process variability. Given that the variance of random mismatch in a capacitor is inversely proportional to its size, the large unit cell capacitance required to meet the thermal noise requirements somewhat relaxes the expected variability of the weighting factors due to random variations. However, the large size of each unit cell and relatively large number of unit cells requires the cells to be spread out across a wide area, potentially making the design susceptible to large-scale mismatch due to process gradients.

To model the effects of mismatch, the capacitance of each unit cell can be modeled as a Gaussian random variable. The mean capacitance of C_{unit} is multiplied by a factor of $1+\alpha$, with $\alpha \sim \mathcal{N}(0, \sigma^2)$. Figure 4.12 shows an example of how the FIR filter frequency response changes due to mismatch. In the ideal frequency response, sharp notches provide a large amount of rejection at the desired frequencies, but the maximum rejection obtainable with these notches steadily degrades using FIR filter coefficients generated with higher levels of mismatch. Figure 4.13 illustrates how the peak consistently achievable rejection (top dashed line in Fig. 4.12) and minimum notch rejection (bottom dashed line in Fig. 4.12) varies as the level of mismatch grows. The histogram in Fig. 4.13(a) illustrates a higher variance of notch depths than peak rejection levels. The peak achievable rejection drops from 60 dB using a mismatch coefficient of $\sigma = 0.005$ to 35 dB assuming a mismatch of $\sigma = 0.05$. The dependence of rejection achievable to certain confidence intervals on mismatch σ is summarized in Fig. 4.13(b). With a mismatch σ of 0.05/0.01, the notch depth is 42/28 dB below the peak in a 99% confidence interval.

In addition to poorer filter rejection, DAC mismatch contributes errors in the SAR ADC conversion residue because it influences the accuracy of the feedback DAC. As a result, the capacitor weights must be calibrated to the full level of resolution desired by the converter. One means of accomplishing this is to use a least mean squares (LMS) algorithm, as outlined in [98], on input data that is known to be sinusoidal. This models each capacitance value as an unknown weight on each bit of the digital ADC output. For a k-point sinusoidal fit $\boldsymbol{f} \in \mathbb{R}^k$ to a sequence of N-bit samples $\boldsymbol{X} \in \mathbb{R}^{N \times k}$, the optimal weighting coefficients $\boldsymbol{w} \in \mathbb{R}^N$ can be calculated as:

$$\hat{\boldsymbol{w}} = (\boldsymbol{X}^T \boldsymbol{X})^{-1} \boldsymbol{X}^T \boldsymbol{f}$$

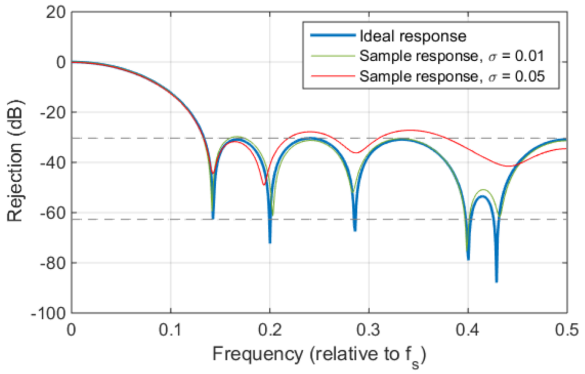

Figure 4.12. Sample effect of mismatch on 11-tap FIR filter frequency response with weights [1 2 3 4 5 5 5 4 3 2 1].

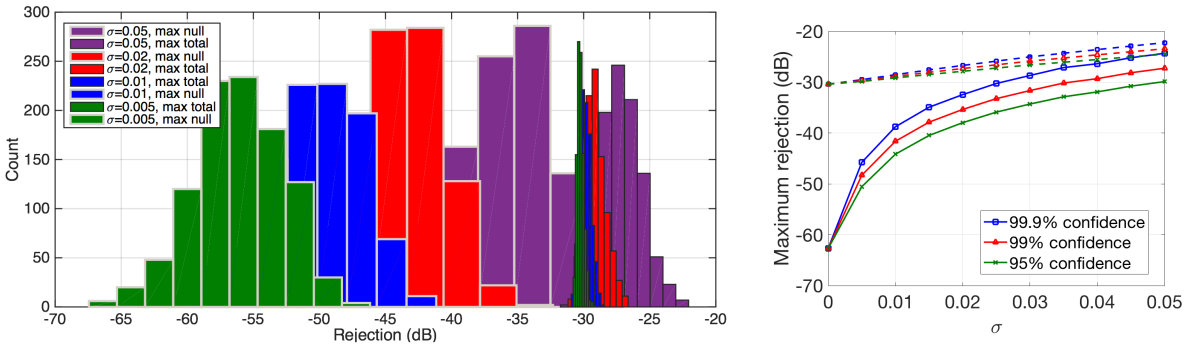

(a) Histogram of maximum rejection and minimum notch depth obtained using normally distributed FIR filter coefficient weights with the variances shown.

(b) Summary of max. rejection and min. notch depth achievable in certain confidence intervals vs. mismatch variance.

Figure 4.13. Summary of mismatch effects on FIR filter rejection.

This one-time computation to measure the capacitor weights can be done in the digital domain. While it consumes power and may require area overhead depending on the system requirements, $\hat{\boldsymbol{w}}$ needs to be calculated only once for each DAC array. For this application, calibration of the DAC coefficients is required to obtain high SNDR in SAR+VCO mode because the coefficient weights must be accurate to within the full accuracy of the converter.

4.2.2 VCO-based ADC design

While the SAR architecture can afford high linearity, input-referred noise from the comparator makes it difficult to design high-resolution SAR ADCs while maintaining low power consumption. This is not suitable for wireless applications, where the minimum detectable signal of the receiver must be very small. As a result, the proposed design uses the SAR ADC as a coarse conversion stage, while a voltage-controlled oscillator (VCO) based ADC serves as a fine conversion stage in a subranging configuration to resolve the sampled conversion residue. As discussed in the previous section, a VCO-based ADC operates as a first-order $\Delta\Sigma$ converter that can be constructed from only a VCO, digital counter, and digital differentiator (Fig. 4.3). Because the VCO-based ADC is oversampled, the SAR stage can be disabled when high linearity is not required to enable the VCO-based ADC to serve as a stand-alone receiver that can tolerate only a single pole of anti-alias filtering from the passive mixer. The VCO design has significant implications on the sensitivity of the receiver. This section summarizes some of the major considerations in designing a VCO-based ADC for low noise presented in [99], and discusses some of the factors limiting the achievable linearity of this design.

4.2.2.1 *Noise analysis*

Understanding the resolution limitations of the VCO-based ADC requires understanding its operation. Conceptually, a VCO-based ADC first translates the input voltage to the frequency of a digital signal using a VCO. This frequency be estimated using the difference between the counter output taken at successive time intervals. As discussed in the previous section, this is a feedforward implementation of a first-order $\Delta\Sigma$ ADC. Its performance will be a function of the resolution of the phase detector (counter), oversampling rate, thermal noise, and any errors in the output sample accuracy due to metastability. Figure 4.14 presents a block diagram that models the operation of the VCO-based ADC and indicates where sources of noise are introduced in the converter. At the input, v_n^2 models the input-referred voltage noise of any preamplifier used in the VCO. In the oscillator, f_n^2 models frequency jitter. After the oscillator frequency is integrated to determine its phase, q_n^2 models quantization noise from the counter (or other phase detection scheme), and s_n^2 models any additional noise introduced during sampling (e.g. metastability errors). A full analysis of noise sources in a VCO-based ADC is presented in [99]. This section summarizes the main results of this work, and applies the analysis more generally to VCO architectures with preamplifiers.

Preamplifier noise An active buffer is typically required to drive an oscillator with the charge residue stored on the capacitive DAC of a SAR ADC. An amplifier can both provide buffering and improve the sensitivity of the VCO, but the input-referred noise of the preamplifier will contribute noise. The input-referred noise of the preamplifier can be added directly to noise from other sources in calculating the total input-referred noise of the VCO. If a fully complimentary amplifier topology with an effective transconductance of G_m is con-

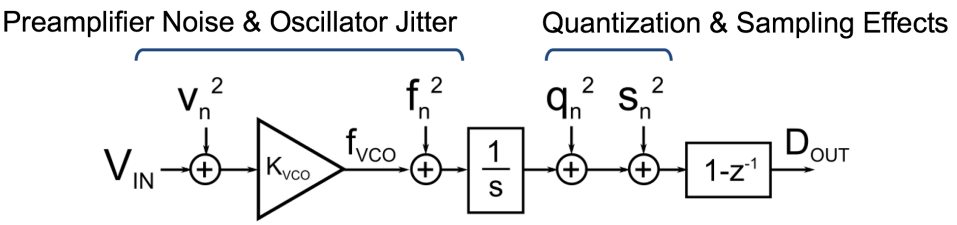

Figure 4.14. Block diagram of the VCO-based ADC, including noise sources.

sidered, the total input-referred voltage noise variance can be calculated as $4kT\gamma f_{BW}/G_m$, where γ is a technology-dependent parameter ($\gamma \approx 2/3$ in long-channel devices) and f_{BW} is the total signal bandwidth. The preamplifier should therefore have large G_m to minimize its input-referred noise.

Input-referred jitter Thermal noise introduces random variation in the delay of each oscillator element (jitter). Because this uncertainty occurs within the VCO signal path, it is not noise-shaped like quantization noise and contributes directly to the input-referred noise of the converter. As a VCO is typically characterized by its voltage-to-frequency gain (K_{VCO}), calculating the input-referred timing jitter requires calculating an effective frequency variance from the inverter delay variance. Modeling a single inverter delay as a fixed mean value of t_{inv} with an uncertain variation $\alpha \sim \mathcal{N}(0, \sigma_{inv}^2)$ and using $var(x)$ to denote the variance of x, the variance of the effective inverter delay frequency f_{inv} is given by the following:

$$var(f_{inv}) = var\left(\frac{1}{t_{inv}(1+\alpha)}\right)$$

Assuming that $\alpha \ll 1$, it can be shown that the variance of f_{inv} approaches σ_{inv}^2/t_{inv}^4. Because the sampling jitter between delay elements is uncorrelated, the effective VCO frequency variance in each ADC sample will scale proportionally with the number of individual delays that accumulate in the total sample duration of $t_{s,ADC}$:

$$var(f_{VCO}) = \frac{t_{s,ADC}}{t_{inv}} var(f_{inv})$$

Given this effective frequency variance, the total input-referred voltage variance due to jitter $v_{n,jitter}^2$ can be calculated using K_{VCO} and the previously discussed parameters:

$$v_{in,jitter}^2 = \left(\frac{\sigma_{inv}}{t_{inv}^2}\frac{1}{K_{VCO}}\right)^2 \frac{t_{s,ADC}}{t_{inv}}$$

To understand how the VCO should be designed to minimize $v_{in,jitter}^2$, σ_{inv}^2 can be expressed in terms of physical VCO design parameters, as derived in [100]. Assuming each delay element drives a capacitive load C_L, contains devices with a threshold voltage of V_{th} (for simplicity, NMOS and PMOS devices have equal V_{th}), and operates on a supply voltage of V_{DD}, the delay variance normalized to a unit delay (σ_{inv}^2/t_{inv}^2) can be calculated as follows:

$$\frac{\sigma_{inv}^2}{t_{inv}^2} = \frac{kT}{C_L(V_{DD}-V_{th})^2}(4\gamma+1)$$

By this analysis, jitter can be minimized by increasing V_{DD}, choosing the lowest available V_{th} for devices used in the delay elements, and increasing C_L, given that γ, k, and T are fixed parameters under given operating conditions. However, because K_{VCO} represents the oscillator frequency change in response to a certain input voltage swing, it is inversely proportional to C_L. To maintain high K_{VCO} while increasing C_L, more current must be provided to the delay elements. Additionally, maximizing t_{inv} can reduce input-referred jitter by reducing the number of delays that accumulate in each sampling period.

Quantization noise The relationship between number of quantization levels N_{lvls}, ADC sampling frequency $f_{s,ADC}$, and signal bandwidth f_{BW} in a typical first-order $\Delta\Sigma$ converter can be summarized as follows:

$$SQNR_{dB} = 20\log_{10}(N_{lvls}) + 30\log_{10}\left(\frac{f_{s,ADC}}{2f_{BW}}\right) - 12.4$$

In this converter architecture, quantization noise increases by 20 dB for each decade increase in $f_{s,ADC}$, leaving a small in-band quantization noise component if the signal bandwidth (f_{BW}) is much lower than $f_{s,ADC}/2$. When combined with the 10 dB/decade SQNR improvement afforded by spreading quantization noise over a wider bandwidth, the SQNR of a typical first-order $\Delta\Sigma$ ADC improves by 30 dB/decade with $f_{s,ADC}$. However, SQNR only improves by 10 dB/decade in a VCO-based ADC because the quantizer resolution (N_{lvls}) is also tied to $f_{s,ADC}$. Intuitively, integrating the VCO phase over a shorter time window makes the total integrated phase smaller and therefore more difficult to detect. For instance, in a digital counter quantizer, the resolution will be set by the number of detectable transitions in the sample period. A longer integration time can detect more transitions for the same input signal. Quantitatively, N_{lvls} can be obtained by dividing the frequency tuning range by $f_{s,ADC}$. This tuning range is the product of K_{VCO} and the input voltage swing $V_{pp,in}$, which allows the VCO-based ADC SQNR to be expressed as:

$$\begin{aligned} SQNR_{dB} &= 20\log_{10}\left(\frac{V_{pp,in}K_{VCO}}{f_{s,ADC}}\right) + 30\log_{10}\left(\frac{f_{s,ADC}}{2f_{BW}}\right) - 12.4 \\ &= 20\log_{10}\left(V_{pp,in}K_{VCO}\right) + 10\log_{10}\left(f_{s,ADC}\right) - 30\log_{10}\left(2f_{BW}\right) - 12.4 \end{aligned}$$

From a design perspective, this SQNR analysis indicates that a high-resolution quantizer should be used to reduce the baseline quantization noise as much as possible. While the simplest means of detecting the VCO phase is to place a digital counter at the output of a single element of the oscillator, this approach neglects other delay cells in the oscillator that contain additional phase information. As discussed in [99], counters can be placed at the output of each delay element and then added digitally to improve resolution. Further details about the VCO decoding scheme in this work are presented in Section 4.3.4.

Sampling noise Any uncorrelated noise introduced during the output sampling process will be first-order shaped because it is directly differentiated at the ADC output without being integrated by the oscillator. This includes errors such as metastability in the digital circuitry decoding the output of the VCO, therefore simplifying the digital design requirements. If the VCO sampling clock is generated from a clean reference such as an off-chip signal source, any edge-to-edge jitter due to phase noise or on-chip buffer noise will also be

first-order shaped. Because this work uses an external sampling clock, noise due to sampling jitter is therefore a less significant concern than input-referred VCO jitter and the input-referred noise of the preamplifier. As discussed in [99], however, any clock source that accumulates jitter will contribute integrated noise that is not shaped by the VCO. If an on-chip oscillator is used to sample the VCO output, care should be taken to minimize its input-referred noise. The work in [99] derives the effective sampling noise-limited SNR in this scenario.

Summary Table 4.1 presents a summary of how major design variables impact the SNR of dominant noise sources. As discussed, doubling sampling frequency (f_s) provides only a 3 dB improvement in quantization noise and sampling error due to metastability because of the relationship between sampling frequency and quantizer resolution. As illustrated by the block diagram in Fig. 4.14, these are the only two noise sources that are first-order shaped; in a conventional first-order $\Delta\Sigma$ ADC, doubling f_s would improve the effective SNR by 9 dB. Because the signal bandwidth (f_{BW}) is fixed, the most effective means of improving VCO resolution is to increase K_{VCO} while keeping the VCO free-running frequency (f_{VCO}) low. The only means of improving the input stage SNR is to burn enough power to increase its transconductance or sample the ADC more quickly.

Variable change	**SQNR**	**Jitter SNR**	**Preamplifier noise SNR**	**Metastability SNR**
$f_s \times 2$	+3 dB	+3 dB	+3 dB	+3 dB
$K_{VCO} \times 2$	+6 dB	+6 dB	-	+6 dB
$f_{BW} \times 0.5$	+9 dB	+3 dB	+3 dB	+9 dB
$f_{VCO} \times 0.5$	-	+3 dB (if constant $\frac{\sigma_{inv}}{t_{inv}}$)	-	-
$\frac{\sigma_{inv}}{t_{inv}} \times 0.5$	-	+6 dB	-	-
$G_m \times 2$	-	+6 dB if $K_{VCO} \propto G_m$	+3 dB	-

Table 4.1. Summary of how design variable changes affect signal to noise ratio of various noise/error sources in VCO-based ADC. A '-' indicates that the design variable does not influence SNR for that particular noise source.

4.2.2.2 *Linearity limitations*

The linearity of the VCO-based ADC is restricted by both the linearity of the preamplifier and the linearity of the oscillator's voltage-to-frequency transfer function. Additionally, mismatch between delay elements can introduce spurious tones in the VCO output spectrum.

Delay element mismatch Each oscillator is comprised of multiple delay cells, which may have different delays due to variation in routing parasitics, device capacitances, and threshold voltages. As described in [99], this can introduce spurious tones in the VCO output spectrum. These tones are first-order noise shaped, and will occur at multiples of the VCO free-running frequency, as transitions between variable inverter element delays will occur at this rate. This source of nonlinearity does not significantly impact the VCO performance if the VCO frequency is much higher than the signal bandwidth, causing these

spurs to fall out of band. Nevertheless, care should be taken in the VCO layout to minimize systematic mismatch between delay elements.

Voltage-to-frequency transfer curve The VCO has a nonlinear voltage-to-frequency transfer characteristic due to the way the preamplifier performance and VCO sensitivity varies with the bias point of each element in the amplifier and oscillator. While feedback amplifiers can provide linear gain insensitive to fluctuations in op-amp gain, a VCO must operate in a feed-forward fashion that prevents such techniques from being applied to improve linearity. Sources of nonlinearity in the VCO can be understood by deriving the K_{VCO} of a generic amplifier. Assuming the input stage provides a transconductance (voltage to current gain) of G_m and small-signal output resistance of R_{out}, and that individual delay cells have a small-signal input resistance of R_{in}, threshold voltage of V_{th} and drive a capacitive load of C_L, K_{VCO} can be approximated as follows:

$$K_{VCO} = G_m \frac{R_{out}}{R_{in} + R_{out}} \frac{1}{C_L V_{th}}$$

This expression shows that K_{VCO} is set by many bias-dependent parameters, especially the small-signal G_m, R_{out}, and R_{in} terms. Even C_L may be bias-dependent if it is dominated by parasitic capacitances within a transistor, which heavily depend on a transistor's bias current and operating region. While these systematic nonlinearities affect the converter largely in the same way as the nonlinear voltage transfer characteristic of an LNA will affect the linearity of a receiver, high-frequency distortion products generated by VCO nonlinearity will be filtered by the sinc response of the VCO integrator [99]. The systematic component of this nonlinearity can be calibrated using look-up-tables optimized digitally using on-chip nonlinearity estimators [59]. In this work, using the VCO as a SAR residue amplifier as in [96] ensures that the amplitude of the VCO input signal is small enough to avoid substantial performance degradation due to nonlinearity.

4.3 Prototype implementation

To demonstrate the suitability of this subranging and filtering architecture for deeply scaled process technologies, a prototype RF-to-digital converter was designed and implemented in a 16nm FinFET process. It consists of an eight-phase mixer-first receiver topology constructed from four parallel sub-ADCs. Each sub-ADC contains a hybrid SAR and VCO-based ADC, using the SAR ADC as a coarse conversion stage to afford high linearity (+17 dBm in-band IIP3) and a VCO-based residue quantizer for improved dynamic range (60 dB SNDR). Either the SAR ADC or the VCO-based ADC can be disabled to configure performance. A passive finite impulse response (FIR) filter is integrated into the SAR ADC sampling to provide configurable filtering. The prototype is designed to detect bandwidths up to 20 MHz with RF carrier frequencies from 700 MHz to 1.9 GHz at moderate power consumption levels (9-45 mW).

This section discusses details of the prototype implementation. First, a high-level overview of the design is provided in Section 4.3.1. Next, Section 4.3.2 describes the local oscillator (LO) clock divider, matching resistor, and mixer. Design of the hybrid SAR ADC and FIR filter is then described in Section 4.3.3, followed by a discussion of the fine-resolution VCO-based ADC in Section 4.3.4. Finally, Section 4.3.5 describes the off-chip

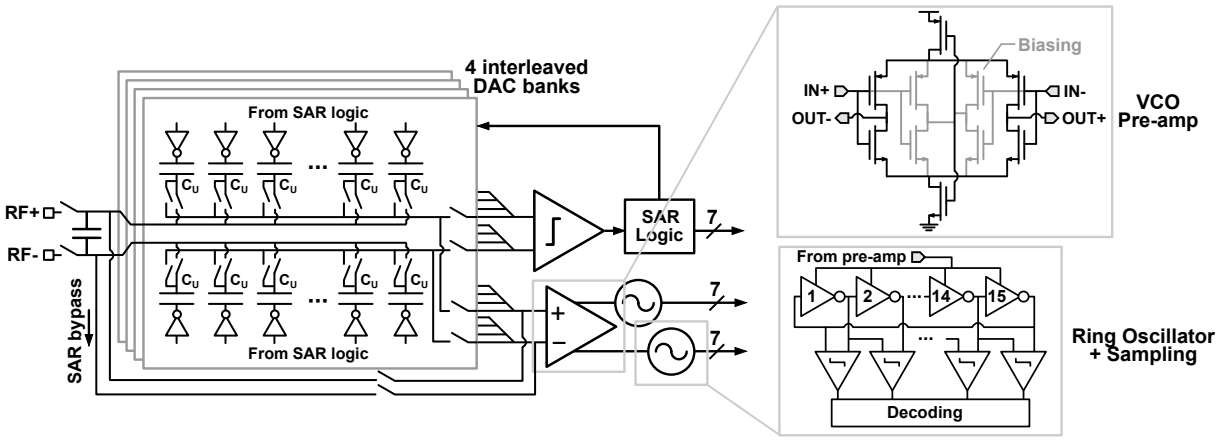

Figure 4.15. Sub-ADC implementation.

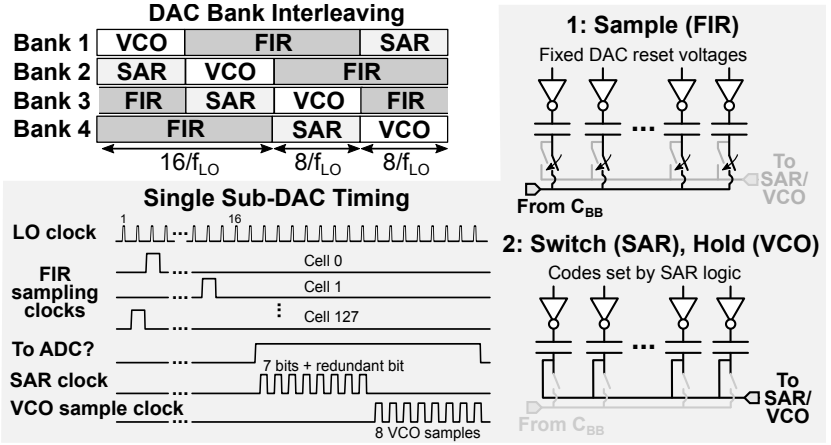

Figure 4.16. Sub-ADC timing diagram.

post-processing procedure used to combine the SAR and VCO signals, and to calibrate for static non-idealities in the circuit.

4.3.1 Overview

The full eight-phase receiver architecture was shown previously in Fig. 4.7. A detailed diagram of a single sub-ADC is shown in Fig. 4.15. In each sub-ADC, a large baseband capacitor sampled at the LO frequency (f_{LO}) down-converts the RF signal to baseband and provides first-order low-pass filtering. The baseband capacitance is programmable, allowing the input bandwidth to be adjusted. The SAR ADC, operating at $f_{LO}/8$, contains four sub-DACs built from 128 unit capacitors that can be individually enabled or disabled to configure the SAR ADC weights and FIR filter response. Figure 4.16 illustrates the interleaved sub-DAC operation, as each DAC alternates between 16 LO cycles of FIR sampling, 8 LO cycles of SAR conversion, and 8 LO cycles of VCO operation. As each unit capacitor is sampled at f_{LO}, interleaving capacitor banks enables up to 16 FIR filter taps. The 16-tap filter can be configured to place nulls at $f_{LO}/8$, providing sharp anti-alias filtering. One SAR comparator and logic block is shared between all sub-DACs to avoid comparator

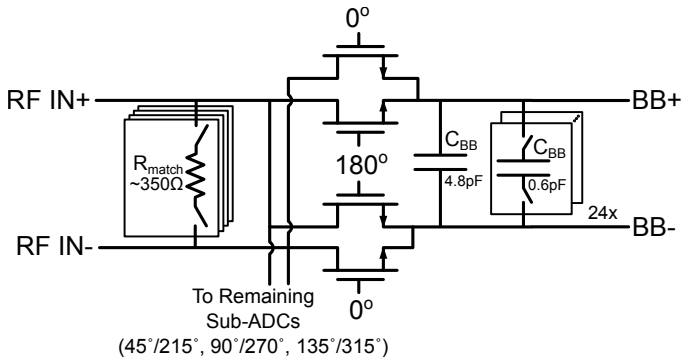

Figure 4.17. Input sampling network for single sub-ADC.

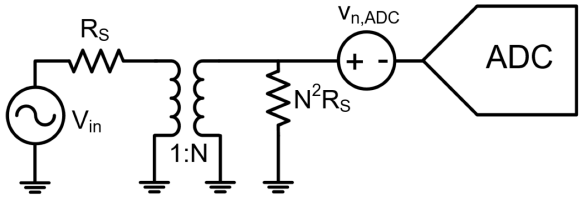

Figure 4.18. Single-ended impedance transformation example.

offset mismatch between banks. The VCO, which is also shared between sub-DACs to mitigate mismatch concerns, is constructed from an open-loop, self-biased preamplifier and two inverter-based ring oscillators to provide two 7-bit digital outputs sampled at f_{LO}.

4.3.2 Input sampling

As shown in Fig. 4.17, the input sampling network consists of a tunable matching resistor, differential NMOS sampling switches, and an 8-phase local oscillator (LO) pulse generator. While the matching resistor is shared between sub-ADCs, one pair of sampling switches and a differential baseband capacitor is contained in each sub-ADC. The LO pulse generator creates 8 phases of the LO (operating at f_{LO}) from an external differential clock signal operating at $4f_{LO}$.

4.3.2.1 *Input matching*

In many mixer-first receivers, input matching is implemented with a feedback resistor in an amplifier immediately after the mixer [12, 94]. However, the baseband sampling capacitor in this design presents a high impedance to the RF port at f_{LO}, so an explicit polysilicon resistor is used in this prototype (as shown in Fig. 4.17) to provide impedance matching. While this resistor limits the achievable noise figure (NF) of the receiver to 3 dB, its noise contribution is significantly lower than the expected input-referred noise of the VCO and SAR ADC.

Because the fixed input-referred ADC noise dominates the total noise of the receiver, an off-chip transformer can be used to provide passive voltage gain that will lower the receiver's effective NF using the structure in Fig. 4.18. While the 3 dB NF limit due to the impedance match will still hold, passive voltage gain will lower the noise contribution of

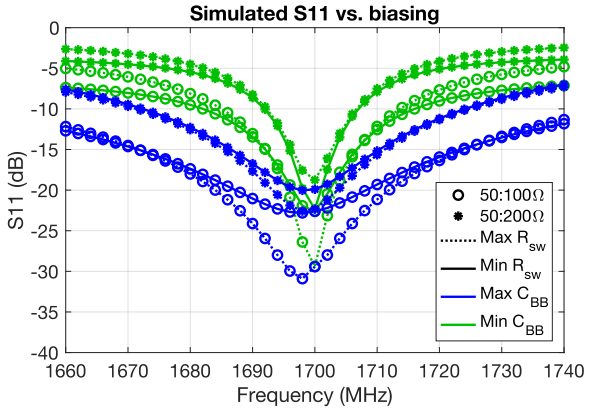

Figure 4.19. Simulated S11 of input sampling network under various biasing conditions and configurations for $f_{LO} = 1700$ MHz.

the ADC relative to the noise generated by the antenna source. While a typical differential termination requires a 100Ω resistor, a 400Ω resistor can be used with a 1:2 impedance ratio transformer to provide 3 dB of passive voltage gain. To support both test configurations, the matching resistor is constructed from four parallel polysilicon resistors designed to provide a fixed resistance of $\approx 350\Omega$. Each resistor segment can be enabled with NMOS switches that provide an additional series resistance of $\approx 20 - 50\Omega$, depending on the common-mode input voltage and supply voltage. Figure 4.20 shows how this resistance is affected by these biasing conditions. By configuring the number of parallel resistors, the total matching resistance can be set to roughly 110Ω, 140Ω, 200Ω or 370Ω.

Figure 4.19 shows the simulated S11 of the input network at $f_{LO} = 1700$ MHz using the maximum and minimum baseband capacitor (C_{BB}) sizes. In a 10 MHz bandwidth, the S11 is kept below -7 dB using the 200Ω matching resistor setting with the largest C_{BB} size (smallest bandwidth). Decreasing C_{BB} to extend the input bandwidth improves the S11 to -16 dB. In the 100Ω setting, the S11 is consistently under -13 dB. The bias voltage sensitivity of the resistor enable switches (R_{SW}) affects the minimum achievable S11, but does not significantly alter the matching bandwidth. While the baseband capacitor (C_{BB}) appears as an open circuit at exactly f_{LO}, causing the matching resistor to set the S11, the translational first-order low-pass filter reduces the impedance of C_{BB} at frequencies offset from f_{LO} according to the filter bandwidth.

4.3.2.2 *Mixer and baseband capacitor*

The mixer is constructed from pairs of low-resistance, ultra low-threshold NMOS sampling switches driven by appropriate phases of the LO, as shown in Fig. 4.17. Each switch is designed to have an ON resistance of $\approx 20 - 30\Omega$, and the input sampling clock is locally buffered to sharpen the sampling edge. The low ON resistance of the mixer switches helps lower noise and reduce bandwidth mismatch between sub-ADCs, since the larger matching resistor is shared among all ADCs. Figure 4.20 illustrates how the switch resistance varies with bias point. To facilitate differential sampling, each mixer output phase is driven by a pair of switches that sample alternating phases of the RF input during alternating clock phases, as shown in Fig. 4.17.

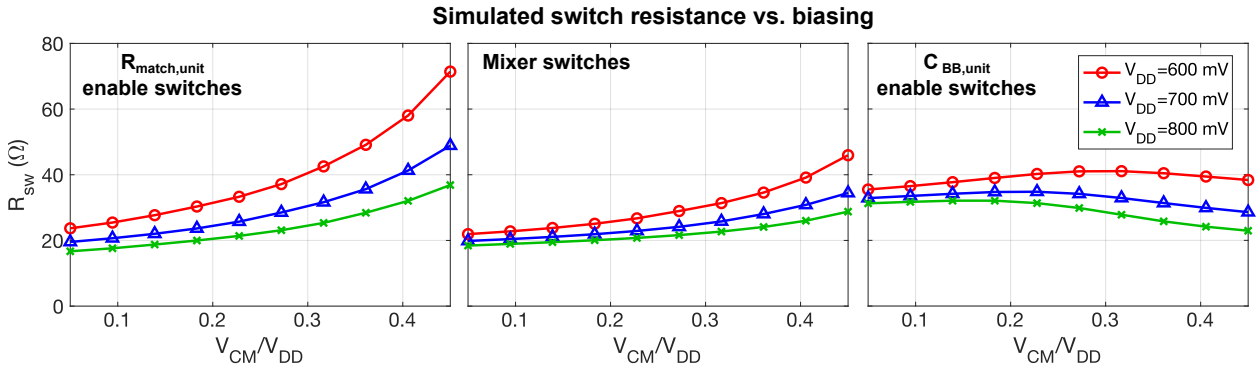

Figure 4.20. Simulated resistance of matching resistor enable switches (left), mixer switches (center), and baseband capacitor enable switches (right) as a function of common-mode input voltage (V_{CM}) and supply voltage (V_{DD}).

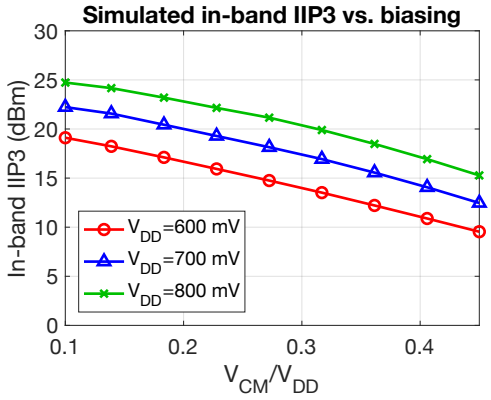

Figure 4.21. Simulated IIP3 of mixer under various biasing conditions.

The differential baseband capacitor (C_{BB}) is constructed from metal-oxide-metal (MOM) unit cells that can be individually enabled to scale the total capacitance from roughly 5 to 20 pF in steps of 0.6 pF. While MOM capacitors have a lower capacitance density than MOS-based gate capacitance, MOM capacitors can be made fully differential, which halves the total required capacitance. CMOS transmission gate switches with an ON resistance of $\approx 30 - 40\Omega$ are used to enable each unit 0.6 pF unit cell. The simulated switch resistance vs. bias point is summarized in Fig. 4.20. Because each side of C_{BB} is sampled for two of the eight total phases of the LO, scaling C_{BB} from $5 - 20$ pF will tune the receiver's 3 dB bandwidth from roughly 20-80 MHz, assuming a 100Ω matching resistor. The partial sampling phase increases the effective RC time constant by a factor of 4 because the signal is only integrated onto C_{BB} 1/4 of the time (two of the eight sampling phases).

For design simplicity, no clock coupling or bootstrapping is used to enhance the linearity of the sampling switches in this prototype. As a result, the achievable linearity of the receiver will be limited by the signal-dependent sampling switch resistance. Figure 4.21 shows the simulated in-band IIP3 of the mixer as a function of input common-mode voltage (V_{CM}) and supply voltage (V_{DD}). The maximum achievable in-band IIP3 receiver is +25 dBm using a minimum V_{CM} and high V_{DD}. The IIP3 decreases roughly linearly as V_{CM} increases and V_{DD}

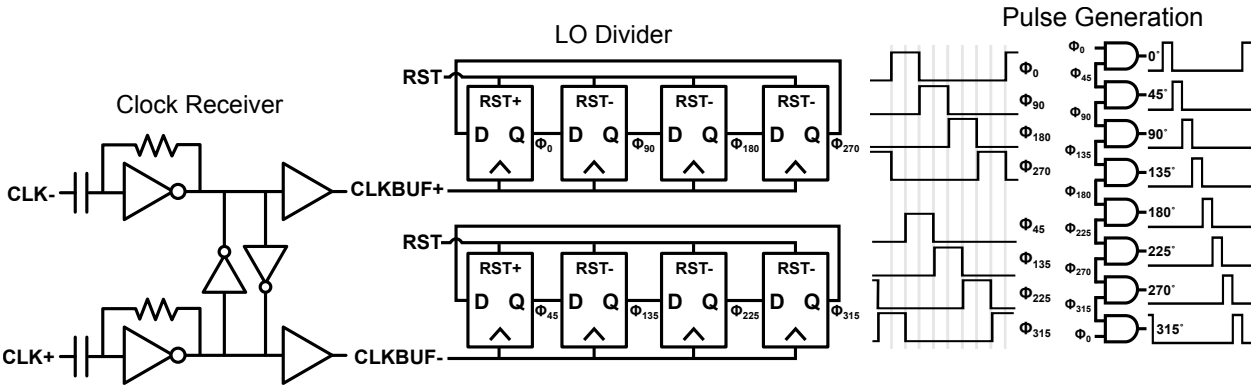

Figure 4.22. Clock receiver and LO divider schematic.

is lowered. To improve the achievable receiver linearity in future designs, thicker oxide devices could be used to increase the V_{GS} of the sampling switches and minimize V_{GS} variation, or a bootstrapped sampling technique could be used to maintain signal-independent V_{GS}.

4.3.2.3 *Clock receiver + LO divider*

Figure 4.22 shows a schematic of the clock receiver and the 8-phase LO divider. The clock receiver is pseudo-differential, and assumes that the CLK+ and CLK- signals are 180° out of phase and operate at $4f_{LO}$. Both CLK+ and CLK- are AC coupled to the input of an inverter-based first stage amplifier with resistive self-biasing. The initial gain stage is followed by latch to help maintain a 180° phase relationship between CLKBUF+ and CLKBUF-. A sequence of standard cell clock buffers to provides additional gain. In simulation, the clock receiver provides a gain from CLK to CLKBUF of > 20 V/V for input frequencies from 50 MHz to 15 GHz.

Two 4-element cyclic shift registers clocked by CLKBUF+ and CLKBUF- are used to generate the phase-shifted LO pulses. In each shift register, one flip-flop is initialized high while the remaining flip-flops are initialized low so that the input clock of $4f_{LO}$ generates four pulses of width $1/(4f_{LO})$ offset by $1/(4f_{LO})$. The 180° phase shift between CLKBUF+ and CLKBUF- provides a phase offset of $1/(8f_{LO})$, so the two shift registers will generate eight pulses offset by $1/(8f_{LO})$ as illustrated in Fig. 4.22. The final eight clock pulses of width $1/(8f_{LO})$ are generated by a series of AND gates. The LO divider and pulse generation circuitry are constructed from maximum size, ultra low threshold standard cells to improve drive strength and reduce sampling jitter.

The functionality of this clock receiver assumes that the CLK+ edge rises before the CLK- edge. If the CLK- edge rises first, the input signals will be 90° out of phase and only four of the eight sampling phases will be generated. The clock receiver in this design implementation does not guarantee this relationship between CLK+ and CLK-, but this could be fixed using additional control logic in future designs. Moreover, if the input clock is applied continuously while the clock divider is reset, the reset signal must have a sufficiently sharp edge to guarantee that all four shift registers are initialized appropriately before the first rising edge of the input clock.

Periodic noise simulations estimate that the LO divider will introduce 420 fs sampling

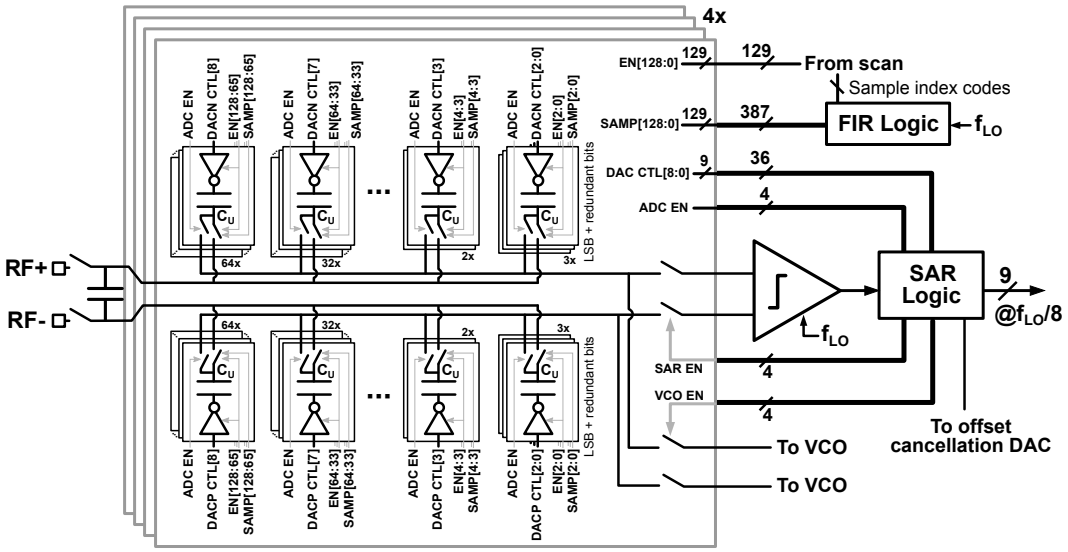

Figure 4.23. Detailed schematic implementation of FIR filter / SAR ADC.

jitter for $V_{DD} = 0.8$ V and 480 fs sampling jitter for $V_{DD} = 0.6$ V, assuming an input clock amplitude of 50 mV peak-to-peak. MATLAB-level simulations of sampling jitter indicate that jitter below 10 ps will not degrade the achievable receiver noise figure by over 0.5 dB. Because the sampling clock is fed externally, noise will contribute cycle-to-cycle jitter but will not have a cumulative effect. Because the LO divider phases are used to generate a clock within each sub-ADC that drives the FIR sampling, SAR operation, and VCO sampling, jitter in the LO divider will contribute to input-referred ADC noise that degrades the receiver NF.

4.3.3 SAR + FIR filter

As shown in Fig. 4.23, the hybrid FIR filter and SAR ADC is constructed from four capacitive DACs, a comparator, SAR logic, and an additional logic block to generate the FIR sampling pulses. The capacitive input sampling both enhances the linear operating range of the receiver and provides a high degree of digitally-configurable passive filtering. To guarantee that the SAR conversion residue falls within the linear input range of the VCO, two redundant bits are added to the seven nominal bits in the SAR ADC. The four DACs operate in an interleaved fashion to implement a 16-tap FIR filter sampled at f_{LO}, which can place nulls at the SAR ADC sampling frequency of $f_{LO}/8$ to prevent aliasing.

4.3.3.1 *DAC design*

To support both FIR filtering and SAR conversion while maintaining a high degree of configurability, each sub-DAC is constructed from 129 identical unit cells, as illustrated in Fig. **??**. As shown in Fig. 4.16, each unit cell is sampled during one of 16 consecutive LO cycles. This sampling cycle can be programmed independently to configure the FIR filter response. Incorporating 129 unit cells provides 7 bits of resolution and two redundant bits to ensure that the SAR residue falls within the linear input range of the VCO. To modify the DAC radix and provide different levels of redundancy, the SAR ADC weights can be modified by disabling cells.

Relative to a conventional SAR ADC, the total DAC capacitance (C_{DAC}) in this RF-to-digital converter must be very large to keep sampled thermal noise below the acceptable noise limit of the receiver. Without a low-noise amplifier (LNA), the input-referred kT/C noise is not attenuated, and each differential DAC sample contributes a voltage noise variance of $4kT/C_{DAC}$ as discussed in Section 3.2.2. While C_{DAC} is divided into unit cells that sample higher noise, these uncorrelated noise samples are averaged by charge sharing prior to SAR conversion to generate a total noise variance of $4kT/C_{DAC}$. This is the dominant noise source during SAR operation because noise generated by the input sampling network can be kept low. The first-order RC filter formed by the baseband capacitor (C_{BB}), impedance matching resistor, and input sampling switches has a bandwidth much smaller than the duration of the sampling pulse, so noise contributed by the mixer is smaller than kT/C_{BB}.

A total C_{DAC} of 13 pF (unit cell capacitance of 50 fF) was chosen to balance the receiver sensitivity requirements against the area and power overhead incurred by using a large sampling capacitor. This keeps the total DAC area to 250 μm by 100 μm per sub-ADC, while limiting the sensitivity (0 dB SNR input power level) to -86 dBm per sub-ADC for signal bandwidths of 10 MHz and $f_{LO} = 1700$ MHz. Each 3 dB improvement in sensitivity requires doubling the DAC capacitance, and therefore area. The sensitivity degrades to -82 dBm for $f_{LO} = 700$ MHz, as the SAR ADC sampling frequency ($f_{s,SAR}$) is a function of f_{LO}. Overall, the relationship between the matching resistor R_{match}, bandwidth (f_{BW}), SAR sampling frequency ($f_{s,SAR}$), and C_{DAC} is given by the following formula:

$$P_{min} = 10\log_{10}\left(\frac{(1mW)(R_{match})}{4kT/C_{DAC}} \times \frac{f_{s,SAR}}{2f_{BW}}\right)$$

To minimize area overhead, the 50 fF unit capacitor is implemented as a 7-layer MOM capacitor. While MOS capacitors provide higher density, their single-ended nature and poor linearity are not well-suited to this application. To reduce the required FIR sampling switch resistance, the switch is transparent for the full LO cycle, and not simply the sampling pulse duration ($1/f_{LO}$ instead of $1/(8f_{LO})$). Compact transmission gate switches with a worst-case on resistance ($V_{DD} = 0.6$V, $V_{CM}/V_{DD} = 0.4$) of 900Ω can settle to 80 dB accuracy in 410 ps, enabling f_{LO} up to 2.4 GHz. For $V_{DD} = 0.8$V, the switch resistance is under 450Ω, which can maintain the same settling accuracy for f_{LO} up to 4.8 GHz. Because the down-converted RF signal is held relatively constant, the design is resilient to sampling clock jitter.

Figure 4.24 shows a schematic of each unit DAC capacitor. The top plate of the capacitor can connect to either the mixer output (during sampling) or the SAR comparator input (during conversion). When a unit cell is disabled, both switches are held open. The bottom plate is driven by a tristate inverter that is high impedance when the cell is disabled, and otherwise driven by the DAC control signals generated from the SAR logic. To afford configurability, logic within each cell generates the switch control signals.

The full layout of the DAC is illustrated in Fig. 4.25. Each sub-DAC is implemented as a vertical array of unit cells, which allows FIR sampling control signals to be routed horizontally to multiple sub-DACs. For uniformity, dummy columns of unit cells are included at the edges of the DAC arrays. Dummy columns are also included in the middle of the array to facilitate routing simplicity. Because the parasitic signal routing resistance combines in series with the sampling switch resistance, it should be minimized using straps of thick metals to avoid increasing settling time. Sharing traces between sub-DACs as shown in Fig. 4.24

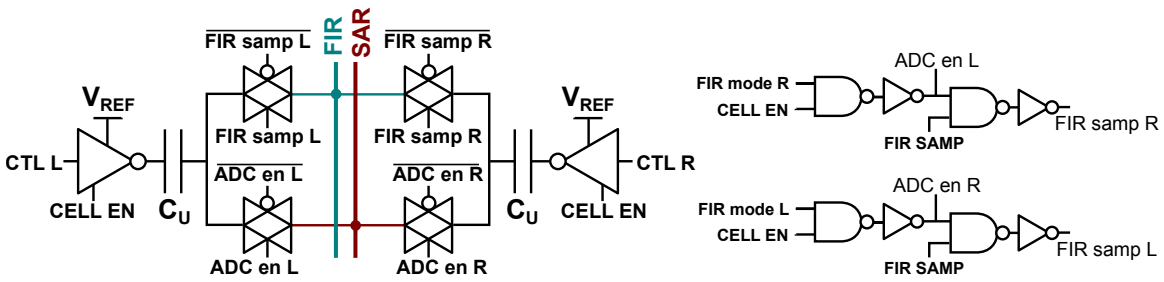

Figure 4.24. DAC unit cell implementation. Unit cells for pairs of sub-DACs are combined together in order to reduce layout area.

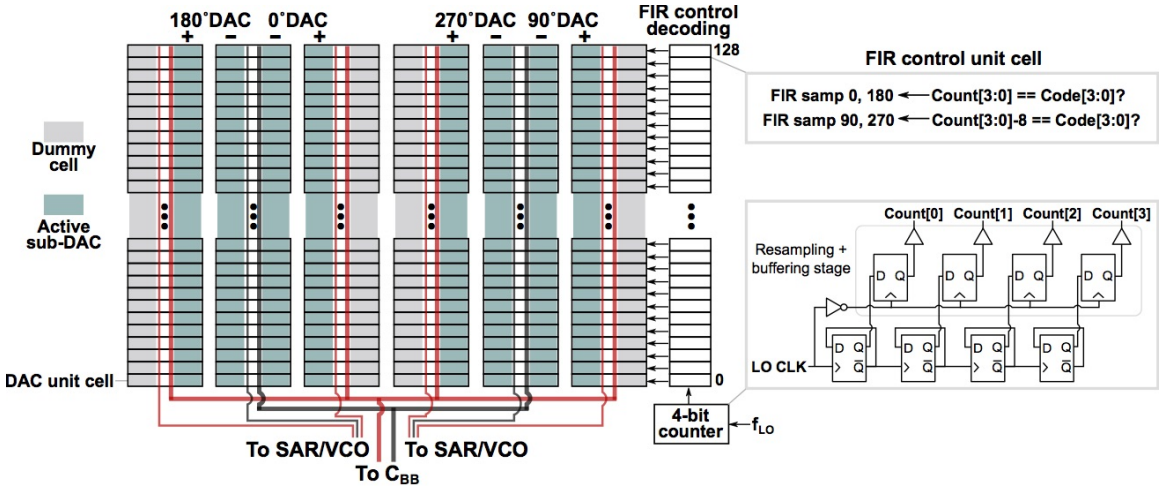

Figure 4.25. Overall DAC and FIR sampling cell implementation.

reduces the area overhead of using thick traces. However, mismatch in routing parasitics and the orientation of unit cells in alternating sub-DAC arrays must be calibrated digitally. For instance, the four columns of DAC+ signal routing relative to two columns of DAC- routing will translate to a higher fixed parasitic capacitance that creates gain mismatch between DAC+ and DAC-.

4.3.3.2 *FIR timing*

Each unit capacitor in the DAC bank is sampled during one of the 16 LO cycles. To generate the sampling pulses, the LO phase is tracked using a 4-bit counter driven by the LO sampling clock, as shown in Fig. 4.25. In the timing cell for each unit element, the counter result is compared to a programmable 4-bit index that generates a sampling pulse of width $1/f_{LO}$ when the index and counter result are identical. Because f_{LO} is much larger than the expected RC bandwidth of the input signal, jitter introduced by the clock generation circuitry will not significantly degrade noise. Moreover, to ensure that all clock edges transition simultaneously, the counter output is re-sampled using four registers with a common clock. This buffering stage occurs once per group of 32 unit cells.

While control signals can be shared between sub-DACs that operate during opposite clock phases (e.g. 0° and 180°), separate control signals must be used for sub-DACs with

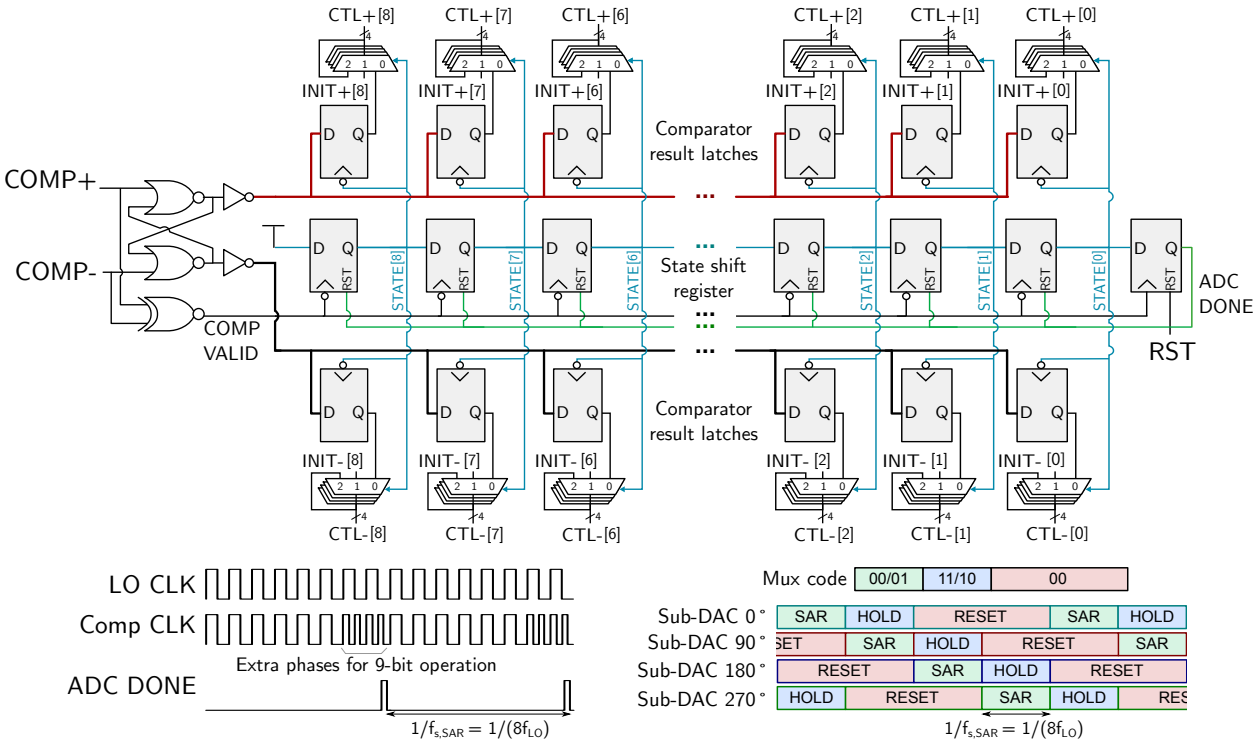

Figure 4.26. SAR ADC logic implementation and clocking scheme.

overlapping phases (e.g. 0° and 90°). To implement this 90 degree phase shift, the 4th bit of the counter result is inverted to generate an additional control pulse that is routed to the appropriate sub-DACs, as illustrated in Fig. 4.25.

Routing the FIR control signal along the 210μm vertical height of the DAC requires long metal wires with a large surface area, introducing significant parasitic capacitance that increases the power consumption of the logic control block. Overall, the estimated power draw of this block is 1.7 mW per sub-DAC at $V_{DD} = 0.8V$ and $f_{LO} = 1700$ MHz. This is 22% of the total sub-DAC power consumption and 31% of the digital power. While this power can be lowered by reducing the number of independently-selectable unit cells, doing so would reduce the configurability of the FIR filter; the power of this logic block is a fixed overhead cost required to implement the configurable filter.

4.3.3.3 *SAR logic and timing*

The SAR ADC control logic determines the bottom-plate capacitance of all four sub-ADCs. In a conventional ADC, the logic resets the DAC bottom-plate voltages during sampling and then sequentially switches the control voltages for bits of decreasing significance as the algorithm progresses. For each sub-DAC in this design, however, the initial reset voltages are fixed for 16 LO cycles (FIR sampling), then the DAC control voltages are switched according to the SAR algorithm for 8 LO cycles (SAR conversion), and then those control voltages are held for the remaining 8 LO cycles (VCO residue amplification). These three modes of operation must be interleaved between the four DAC banks, as shown in Fig. 4.26.

Figure 4.26 also shows how the logic is constructed from standard cell flip-flops and

multiplexers. A clock generated from the LO and SAR logic (bottom left of Fig. 4.26) drives the comparator for synchronous operation, and the XNOR of the two comparator outputs (COMP+ and COMP-) generates a signal to indicate that the comparator result is valid. The sequential rising edges of the shift register drive another set of registers that store the current comparator result, as in a conventional SAR ADC. Each comparator result latch drives four multiplexers that control the bottom-plate voltages of one of the four sub-DACs. During the FIR sampling phase, the DAC bottom-plate voltage is reset to a programmable initial value that dictates the final common-mode output voltage of the DAC. During the SAR operating phase, the multiplexer is used to switch between the initial DAC voltage and the digital logic output. When the VCO is enabled, the mux is configured as a latch to store the previous DAC bottom-plate code, maintaining a constant residue voltage without the area overhead of a full flip-flop. After 9 phases are complete, the digital logic is reset by a final register that generates a 'conversion DONE' pulse. This DONE signal then drives a 2-bit digital counter to generate control codes for the multiplexers driving specific sub-DACs (bottom right of Fig. 4.26).

The SAR ADC logic is also used to cycle the DAC reference voltage between one of three off-chip low-dropout regulators (LDOs), because the reference noise requirements vary between operating modes. The reference voltage noise must be kept low during the residue conversion phase (VCO operation) and sampling phase (FIR operation), but high noise is tolerable during the SAR conversion phase. The SAR conversion phase also generates a large amount of supply noise when the large bottom-plate capacitance is switched during the conversion procedure. As a result, three separate references are employed to isolate the switching noise of the SAR phase from the noise-sensitive VCO phase.

The digital logic timing requirements are greatly relaxed by the subranging ADC architecture. Conventional SAR ADCs require the DAC to settle within half an LSB at each decision cycle, but this architecture requires complete settling only by the beginning of the VCO conversion phase. This affords the large MSB capacitors the longest settling window. The 50 fF unit capacitor and $< 860\Omega$ driving switch will safely settle to within 80 dB accuracy in 400 ps, so even the LSB will settle fully using $f_{LO} > 2.5$ GHz. All but the two redundant unit capacitors will settle to > 100 dB accuracy in this time frame. Incorrect comparator decisions resulting from incomplete DAC settling are tolerable provided the final SAR residue is within the VCO input range; the two redundant SAR bits help correct for such errors.

Because DAC settling error can be kept small, the main limitation on logic performance is that the comparator decision and logic delay must complete within each LO cycle due to the synchronous operation of the ADC. Figure 4.27 shows the simulated logic decision delay vs. comparator residue amplitude for various supply voltages, using the post-layout extraction netlist. This indicates that the SAR logic should operate reliably up to a frequency of 3.8 GHz with $V_{DD} = 800$ mV and 2.1 GHz with $V_{DD} = 600$ mV.

4.3.3.4 *SAR switching procedure*

While constant-V_{CM} switching procedures are often used to minimize nonlinearity caused by comparator offset variation, this subranging architecture is resilient to comparator decision errors, as discussed above. Moreover, while the common-mode voltage of the RF input signal must be kept low to reduce sampling nonlinearity in the ADC (Fig. 4.21), the input to the

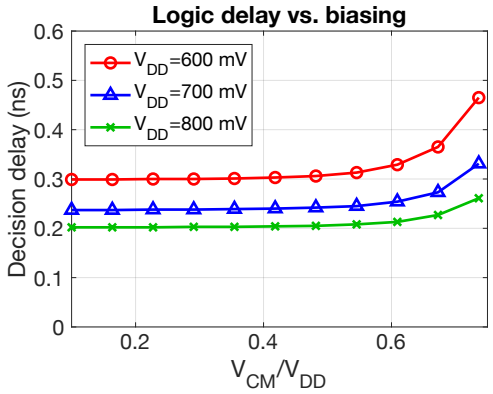

Figure 4.27. Simulated SAR ADC logic decision delay (comparator clock rising edge to DAC control signal rising edge) vs. bias point. A residue input of 100 μV is used.

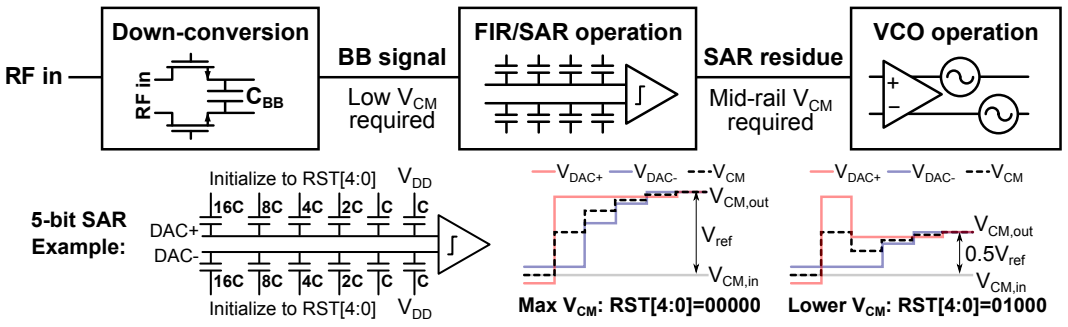

Figure 4.28. The NMOS input sampling switches require a low input common-mode voltage (V_{CM}), while the VCO preamplifier requires a mid-rail V_{CM}. As a result, the programmable-V_{CM} switching procedure described in the previous chapter (illustrated here) is used.

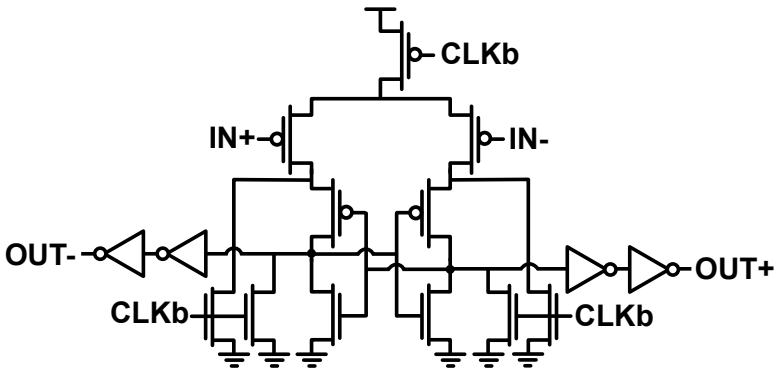

Figure 4.29. Strong-arm latch comparator used for the SAR ADC.

VCO preamplifier should be near mid-rail to maximize its transconductance and provide low input-referred noise. A low-power switching procedure is also desirable due to the large DAC capacitance required to keep thermal sampling noise low. As a result, the same switching procedure used in the resolution-configurable SAR ADC of Chapter 3 (discussed in Sec. 3.3.1) is also used in this work to integrate programmable common-mode biasing directly into the SAR sampling process (Fig. 4.28). The initial bottom-plate voltage of the DAC unit cells determines the common-mode output voltage.

4.3.3.5 *Comparator*

Because this architecture is resilient to comparator decision errors, the input-referred noise and offset of the comparator can be relatively high. As a result, a compact PMOS-input strong-arm latch shown in Fig. 4.29 is used for the SAR ADC. PMOS input devices are used because the common-mode input voltage is low to minimize distortion from the NMOS sampling switches. While this results in a larger comparator input overdrive that increases input-referred noise and offset as discussed in the previous chapter, it also lowers the comparator decision time. No offset calibration techniques are integrated into the comparator itself. Because the single comparator is shared between sub-ADCs, offset can be removed digitally as a fixed DC component of the output signal.

Figure 4.30 shows the simulated noise and offset of the comparator, obtained (respectively) via transient noise simulations and Monte Carlo simulations. The input-referred noise standard deviation is under 5 mV when V_{CM} is low, but easily falls below 2 mV as V_{CM} increases. For an 800 mV full-scale input value, this is under half the 7-bit LSB of 6.25 mV. While the comparator offset variance is high, it will largely be a static value that can be corrected digitally. The main performance limitation of this design is the comparator decision time, which may limit the achievable speed of the converter. However, the results in Fig. 4.27 include the comparator conversion delay and indicate that LO frequencies up to 2.1 GHz can be supported down to operating voltages of 600 mV.

4.3.3.6 *Expected conversion residue*

While the input-referred comparator noise can be significantly larger than the target noise floor of the converter, it must still be smaller the input range of the VCO-based ADC. Similarly, nonlinearity induced by comparator offset V_{CM} dependence (as discussed in the previous chapter) and DAC mismatch must be smaller than the VCO's linear input range.

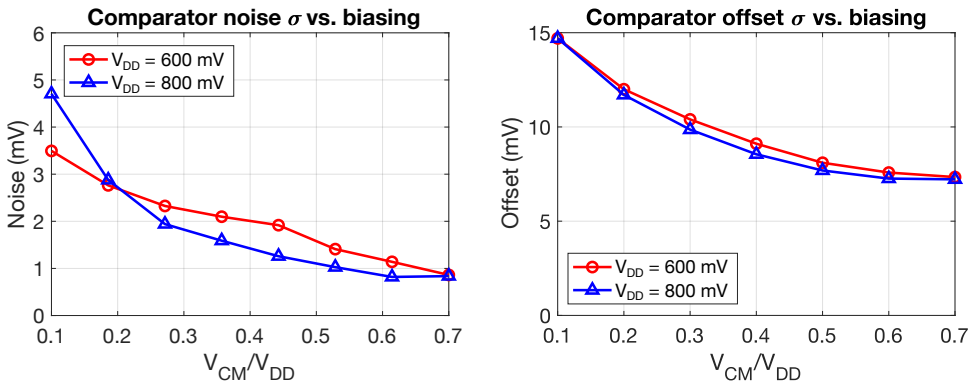

Figure 4.30. Simulated input-referred noise and offset standard deviation of PMOS-input comparator. Noise σ was estimated via transient noise simulations, while offset σ is obtained from Monte Carlo simulations

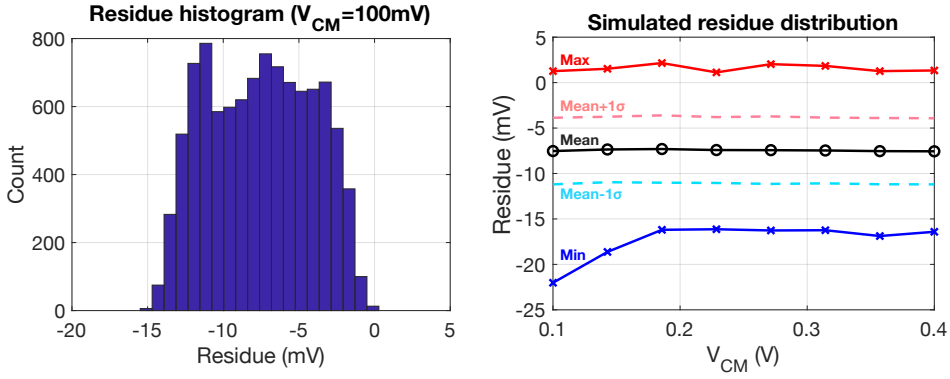

Figure 4.31. Simulated distribution of SAR ADC residue voltages. An example residue histogram for $V_{CM} = 100$ mV is shown at the left (assuming $V_{DD} = 800$ mV), and the summarized distribution statistics as a function of V_{CM} are shown at the right.

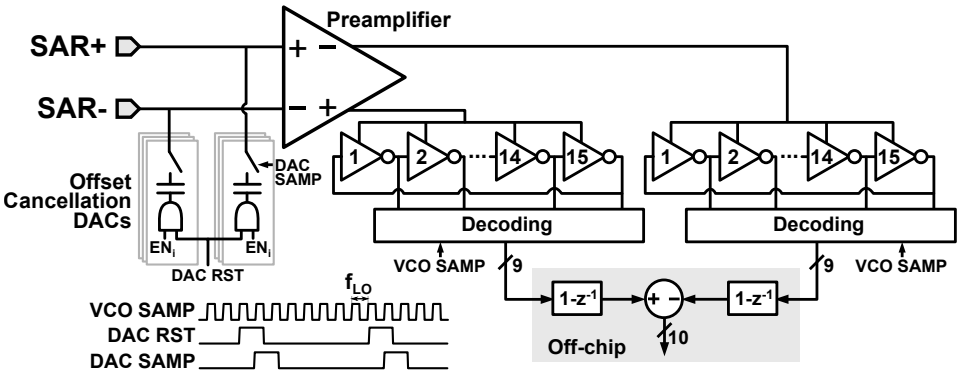

Figure 4.32. VCO-based ADC architecture.

Figure 4.31 compares the expected distribution of SAR ADC residue voltages, assuming a DAC unit cell mismatch of 5%. This estimate is obtained using a MATLAB model of a SAR ADC that assumes a switching procedure with monotonically-increasing V_{CM}. The comparator characteristics (noise and 1σ offset) from Fig. 4.30 are included in the model. From the sample histogram in Fig. 4.31, it can be seen that the residue experiences a nearly uniform distribution with a Gaussian tail due to the comparator noise. The summarized distribution statistics as a function of initial V_{CM} illustrate that the residue distribution is relatively independent of V_{CM}, and nearly all of the ADC residues fall safely within a 30 mV window.

4.3.4 VCO-based ADC

The VCO-based ADC operates as an open-loop implementation of a first-order $\Delta\Sigma$ ADC, in which the input signal is integrated using the frequency/phase relationship of an oscillator, and then digitally differentiated. As shown in Fig. 4.32, the VCO-based ADC implemented in this work consists of an open-loop preamplifier to reduce noise and buffer the capacitively sampled SAR residue, two VCOs operating in a pseudo-differential fashion, and decoding logic to translate the VCO phase into a digital signal. Because the input-referred noise of the VCO-based ADC adds directly to the noise of the matching network and input sampling structure, both thermal noise and in-band quantization noise must be kept low to avoid degrading the receiver sensitivity. Quantization noise requirements are relaxed due to the first-order noise shaping properties of the VCO. While high linearity is not required due to the SAR stage, the linear input range of the VCO must cover the full SAR ADC residue range. This section discusses in further detail the implementation of each block required to build this fine-resolution ADC.

4.3.4.1 *Preamplifier design*

The VCO preamplifier buffers the capacitively sampled SAR voltage residue from the ring oscillator and maintains low input-referred noise to achieve high sensitivity. Figure 4.33 shows the differential self-biased preamplifier topology chosen for this prototype to maximize transconductance (G_m) and afford implementation simplicity. A fully differential structure is chosen to reduce the second-order harmonic distortion of the VCO. Moreover, the amplifier is relatively simple (few transistors) to be compatible with increasingly complex design rules

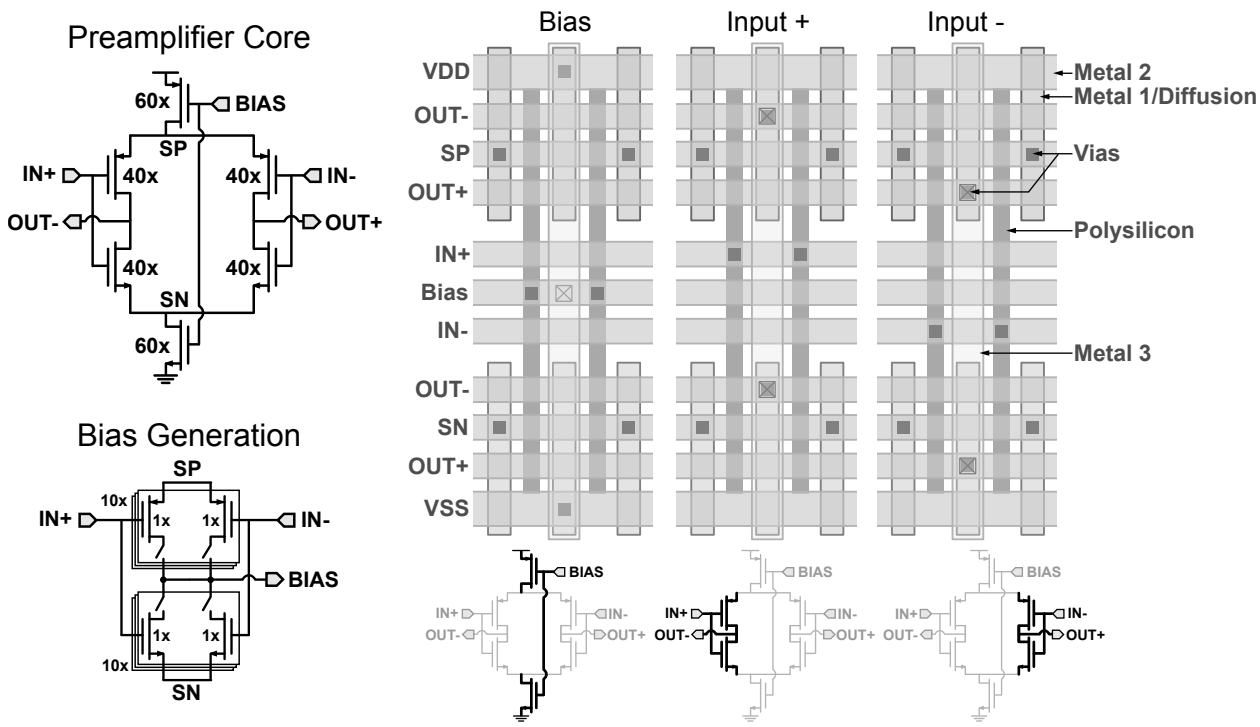

Figure 4.33. VCO preamplifier schematic (left) and layout unit cell implementation (right).

in scaled process nodes.

Relative to a pseudo-differential amplifier built from two independent inverter-based amplifiers, this architecture provides improved common-mode voltage (V_{CM}) rejection to tolerate V_{CM} variation induced by common-mode noise or differential gain imbalance. The complementary amplifier design also lowers the input-referred oscillator noise relative to a single differential input pair driving the two ring oscillators. As discussed in Section 4.2.2, the oscillator noise is inversely proportional to the ratio between Δf induced by the input signal (a function of the input swing and G_m) and the free-running VCO frequency (a function of the fixed bias current). As a result, the fully differential input topology affords a high input transconductance that helps to minimize the input-referred noise of the design.

As shown in Fig. 4.33, a tunable self-biasing replica circuit adjusts the DC output voltage of the preamplifier to control the free-running VCO frequency. To adjust the output voltage, the relative PMOS/NMOS strength can be tuned by modifying the number of parallel devices in the replica input pair. If the PMOS drive strength is larger, the bias voltage is pulled high to lower the output voltage, which lowers the VCO speed. Conversely, the free-running VCO speed can be increased by using more parallel NMOS devices to lower the amplifier's bias voltage and raise its output voltage.

The highly symmetric amplifier structure also facilitates a straightforward physical layout, as shown in Fig. 4.33. All transistors in the core of the amplifier can be constructed from identical two-finger unit cells. In the bias devices, the inner shared source/drain region can connect to V_{DD} or V_{SS}, while the outer source/drain region connects to the common source node of the input pair. The outermost node of the input pair devices connects to this

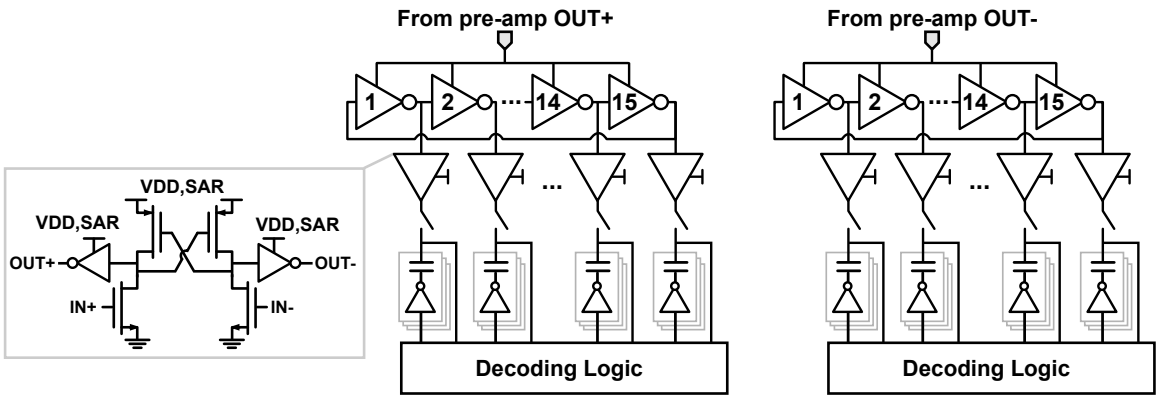

Figure 4.34. Ring oscillator implementation.

same node, while the innermost node will connect to one of the two preamplifier outputs. These unit cells can be interleaved to improve matching and tiled to easily configure the relative strength of the input devices and bias transistors.

Device dimensions are chosen to maximize VCO gain while maintaining low input-referred noise. Using many parallel fingers to implement the input devices helps maximize the effective G_m of the input stage for a fixed power (drive current). However, to maximize the proportion of differential output current ($G_m v_{in}$) fed to the ring oscillator, the output resistance of the preamplifier ($R_{out,amp}$) must be higher than the input resistance of the VCO ($R_{in,VCO}$):

$$\Delta i_{VCO} = G_m \Delta v_{in} \frac{R_{out,amp}}{R_{in,VCO} + R_{out,amp}}$$

For a given VCO, larger Δi_{VCO} for a fixed Δv_{in} will provide a higher VCO gain (K_{VCO}) that improves ADC sensitivity, so G_m and $R_{out,amp}$ should be maximized. In this design, a total of 60 unit bias cells are used relative to the 40 input devices to provide an overall input transconductance of 22 mS. At a supply voltage of 0.8V, the preamplifier consumes an estimated 730 μW, which falls to 210 μW on a 0.6V supply.

Finally, the input-referred offset of the preamplifier can be canceled using a capacitive DAC at the preamplifier's input, as shown in Fig. 4.32. While the offset of the preamplifier may be small due to the large input device size, the input-referred offset of the SAR ADC comparator will effectively be added to the SAR conversion residue, and may be large enough to saturate the input range of the VCO. This 5-bit DAC has a total capacitance of $\approx$ 0.6 pF that can be combined with the 6.5 pF single-ended DAC capacitance to tune offset over a range of $\pm$30 mV with 0.6 mV of resolution. The bottom plate of the DAC is initialized to a known state prior to conversion during the reset phase, and a fixed voltage change is applied during the sampling phase. These control signals are generated from the SAR logic.

4.3.4.2 *Oscillator topology*

The oscillator must provide low input-referred noise, have sufficient phase resolution to minimize quantization noise, and be simple to implement in scaled CMOS process nodes. As shown in Fig. 4.34, each VCO is constructed from a 15-element inverter ring whose supply voltage is controlled by the output of the preamplifier. The output of the ring oscillator is

buffered by an additional stage of inverters powered by the SAR ADC voltage reference, and then decoded digitally. The ≈ 400 mV SAR ADC reference voltage is used for the buffer supply to account for the low preamplifier output voltage.

As discussed in Section 4.2.2 and detailed in [99], the input-referred VCO thermal noise is a function of the VCO gain (K_{VCO}), inverter delay (t_{inv}), and jitter introduced by each delay element (σ_{inv}). When the inverter delay standard deviation is much smaller than the unit inverter delay ($\sigma_{inv}/t_{inv} \ll 1$), the input-referred noise can be approximated as follows:

$$v_{n,in}^2 = \left(\frac{\sigma_{inv}}{t_{inv}^2}\right)^2 \frac{t_{samp}}{t_{inv}} \left(\frac{1}{K_{VCO}}\right)^2$$

This expression shows that the input-referred jitter can be minimized by increasing t_{inv}, lowering σ_{inv}, or increasing K_{VCO}. As presented in [100] and discussed in Section 4.2.2, these parameters are related; for instance, increasing the capacitive load at the output of each inverter (e.g., by increasing device length) will increase t_{inv} and lower σ_{inv}, but also lower K_{VCO}. Meanwhile, for a fixed sample rate, the quantizer resolution will be inversely proportional to t_{inv}, so the delay element sizing presents a tradeoff between the input-referred jitter of the ring oscillator and the first-order shaped quantization noise.

To address these design considerations, the inverters in the VCO are constructed of low-threshold, minimum-length transistors to reduce area and optimize for phase resolution. Each transistor is constructed from multiple device fingers to increase capacitive loading, thereby reducing σ_{inv}/t_{inv} while keeping t_{inv} small to maintain phase resolution. Placing multiple low-threshold devices in parallel also lowers the effective input resistance of the oscillator to boost K_{VCO}, as discussed in the previous section. Low threshold devices maximize the voltage swing of each inverter, which reduces σ_{inv}/t_{inv} as detailed in [100].

The number of delay elements in the oscillator impacts the low-frequency noise, nonlinearity, and decoding complexity of the ADC. As described in [100], using many delay elements provides a higher degree of low-frequency noise averaging within the oscillator. However, the analysis in [99] illustrates that using more unit elements introduces additional spurs due to delay element mismatch. These spurs can be pushed to higher frequencies when t_{samp}/t_{inv} is large, at the expense of increasing input-referred jitter. Finally, the VCO phase can be decoded by identifying which of the oscillator elements is transitioning when the output voltage is sampled. Therefore, incorporating many delay elements requires a higher degree of decoding complexity. In this oscillator, a total of 15 high-speed inverters are used to provide a balance between low-frequency noise averaging and decoding simplicity.

In the physical design, each VCO is placed in a separate guard ring to provide isolation from the adjacent complementary VCO and prevent phase locking. While the inverter layout is naturally simple, careful layout practices are used to ensure uniformity between each inverter unit cell and reduce systematic mismatch. The output buffer is placed adjacent to each element in the ring oscillator to reduce capacitive parasitics and improve matching.

Figure 4.35 summarizes the simulated VCO gain (K_{VCO}), delay element noise (σ_{inv}/t_{inv}), and inverter delay (t_{inv}) as a function of bias point. Both high K_{VCO} and large t_{inv} can be obtained if the common mode input voltage (V_{CM}) is high enough to lower the supply voltage of the VCO. Similarly, lowering the VCO supply voltage can increase t_{inv} without significantly reducing K_{VCO}. While this increases σ_{inv}/t_{inv}, it lowers the overall input-referred voltage noise of the design as shown in Fig. 4.36 and discussed in the next section.

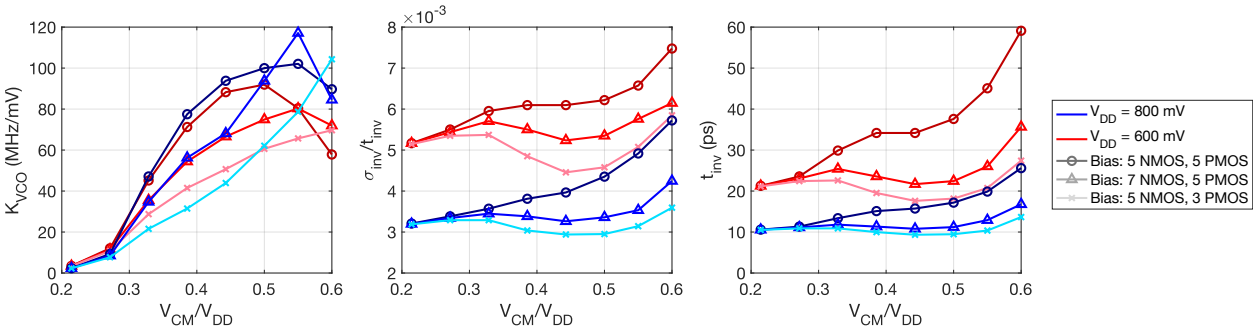

Figure 4.35. VCO gain (K_{VCO}, left), relative delay standard deviation (σ_{inv}/t_{inv}, center), and inverter delay (t_{inv}, right) as a function of VCO bias configuration.

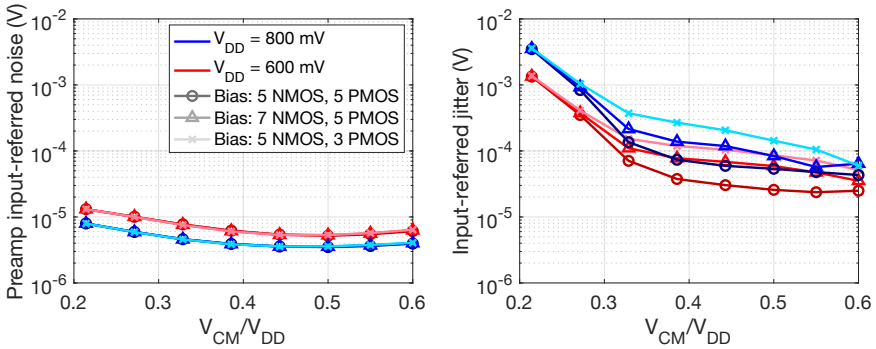

Figure 4.36. Input-referred preamplifier thermal noise (left) and ring oscillator jitter (right) at $f_{LO} = 1.7$ GHz.

4.3.4.3 *Preamplifier + oscillator performance*

Figure 4.36 compares the simulated input-referred thermal noise of the preamplifier and the oscillator as a function of bias point for an LO frequency of 1.7 GHz. The preamplifier contributes significantly less noise than the oscillator. Both the preamplifier noise and the oscillator noise improve with large G_m, as this boosts K_{VCO}. The G_m is maximized when the amplifier's input common mode voltage (V_{CM}) is near mid-rail. Increasing V_{CM} also lowers the preamplifier's output voltage, which reduces t_{inv} to improve the input-referred oscillator jitter. The preamplifier output voltage can also be reduced by lowering the supply voltage, which correspondingly lowers the VCO power consumption. At $f_{LO} = 1.7$ GHz, the core VCO (preamplifier, oscillator, and buffer) consumes 400 μW for $V_{DD} = 600$ mV and 1.1 mW for $V_{DD} = 800$ mV. The left plot of Fig. 4.36 illustrates that lowering V_{DD} raises the preamplifier's noise; however, at $V_{DD} = 600$ mV it remains below 25% of the oscillator's noise contribution. While lowering V_{DD} reduces both the VCO power consumption and total input-referred noise, higher V_{DD} is required for high mixer switch linearity.

4.3.4.4 *Signal decoding*

The VCO phase detector must provide a high degree of resolution to ensure that the converter's quantization noise is lower than the target noise floor of the receiver. Coarse phase

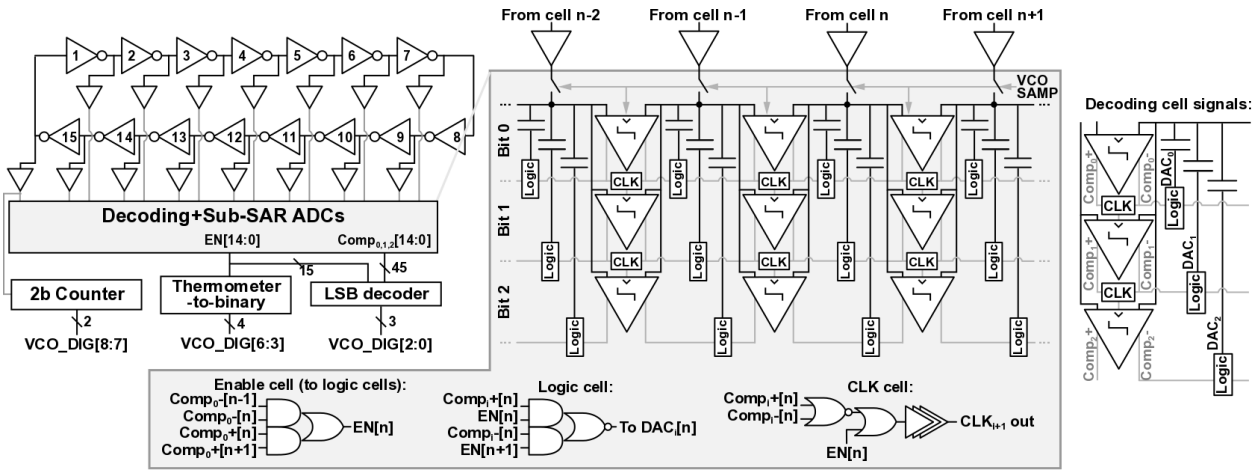

Figure 4.37. VCO decoding scheme. The inset shows how the buffered VCO output signals (labeled $n-2, n-1, n, n+1$ at top) connect to the output of adjacent signals

quantization can be performed using a digital counter to integrate the VCO frequency. Additional resolution can be obtained by using the output of each VCO element to identify the oscillator phase, as the sampled code will indicate which delay stage transitions during a particular sampling instant. To obtain a higher degree of precision, the analog voltage of the transitioning element can be sampled and digitized with a coarse ADC. As shown in Fig. 4.37, these three techniques are combined in this prototype to provide a wide input range. Two MSBs of the VCO are measured using a conventional counter-based approach, and three additional bits are obtained by identifying which of the inverter elements is transitioning during the sampling instant. The output of each element is sampled onto a 3-bit capacitive DAC that can be used as a coarse sub-SAR ADC to provide additional quantization resolution.

While digital counters are typically used in VCO-based ADC decoding [89, 101], most prior designs use flip-flops instead of comparators to decode the VCO phase [96, 99], and do not boost SQNR using an additional coarse converter. The decoding logic in this work uses comparators as a simple mechanism for translating between the low supply voltage of the ring oscillator output buffers and the higher digital supply voltage. The comparators can also be reused for the coarse-resolution sub-SAR ADC. As illustrated in Fig. 4.37, all 15 initial comparator results are first used to identify which voltage contains the transition residue. Subsequently, the relevant comparators are triggered asynchronously to provide three additional bits of resolution. Because quantization noise will be first-order shaped, only a few bits of sub-SAR resolution are required and the comparators and sampling DAC can be made minimum-sized to reduce power consumption. On the rising edge of the sampling clock (generated from the LO), the buffer output voltages are sampled onto a 2 fF, 2-bit capacitive MOM DAC.

In simulation, the decoding procedure completes in 400 ps, making it compatible with LO frequencies up to 2.5 GHz. The 3-bit sub-SAR operation completes in 250 ps, so faster LO frequencies can be supported when higher quantization noise is tolerable. At 1.7 GHz and $V_{DD} = 0.8$ V, the decoding scheme consumes 1.4 mW per sub-ADC, divided between

0.09 mW for the digital counter, 0.04 mW for the decoding logic, 0.43 mW to support the sub-SAR operation, and 0.88 mW to distribute the sampling clock. The buffers used to drive the DAC capacitance consume 0.33 mW. Because this power is largely dynamic, it can be reduced by lowering the sample rate when possible.

The on-chip signal decoding logic does not differentiate the output samples to generate a sequence of noise-shaped output values. This was done to simplify the prototype design and provide a higher degree of debugging flexibility, but an on-chip differentiator could be integrated into future designs. The main challenge in differentiating the VCO phase samples is the potential for overflow in the MSB counter. However, the continuously-integrating nature of the VCO-based ADC indicates that the VCO phase code should grow between successive samples. By guaranteeing that no viable input amplitude exceeds the detectable VCO phase range, any negative phase code difference can be corrected using the known counter range.

4.3.5 Signal recombination

An iterative optimization algorithm is implemented off-chip to estimate the gain from the VCO code to the SAR code and account for static non-idealities such as DAC capacitor mismatch. Each independent sub-DAC contains different capacitor weights, and routing mismatches will vary the offset of each DAC by changing the parasitic capacitance from the DAC top plate voltage to ground. Optimizing the converter's performance requires calculating these weights. To do this, an iterative gradient ascent algorithm is used to find coefficients that maximize the SNDR of constant-wave input training data. In this algorithm, a vector of parameters $\boldsymbol{p}$ is calculated by iteratively updating each individual parameter p_n according to the derivative of the function to be maximized (SNDR in this case). Specifically, a new estimate of p_n is found at iteration i of the algorithm by modifying the previous estimate (p_n^{i-1}) by an amount proportional to the partial derivative of SNDR with respect to p_n:

$$p_n^i = p_n^{i-1} + \alpha \frac{\partial SNDR(\boldsymbol{p}^{i-1})}{\partial p_n}$$

In this case, the partial derivative can be estimated by evaluating the SNDR twice for slightly different values of p_n, while keeping the other parameters in $\boldsymbol{p}$ fixed. Relative to a least mean squares (LMS) algorithm, as described in [98], this approach does not require fitting an ideal signal to the measurements. It requires FFT computations and comparison logic, but this functionality can be integrated into the digital baseband processor of a conventional receiver. Most digital baseband processors will already incorporate an FFT for decoding, as orthogonal frequency division multiplexing (OFDM) encodes wireless data in the frequency domain. Moreover, a one-time calibration is sufficient for measuring DAC mismatch, so the total relative the power overhead of this computation can be kept low. Alternatively, the conventional LMS weight-fitting algorithm can be modified to calculate the VCO residue gain by treating the unknown SAR to VCO gain as an additional weight in the unknown coefficient matrix that scales the measured VCO code.

4.3.6 Top-level design

Figure 4.38 shows both the layout of a single sub-ADC and a die micrograph of the flip-chip 16nm CMOS prototype. The active area of the four sub-ADCs is 0.26μm. As seen in the micrograph, the four sub-ADCs are arranged to minimize the length of RF signal routes. The differential RF input and high-speed $4f_{LO}$ clock are fed from two pairs of flip-chip solder bumps at the die edge to the right and left of the chip, respectively, using the top aluminum routing layer. The receiver placement relative to the RF input bumps was dictated by packaging constraints; in a customized design, the receiver would be placed as close as possible to the RF input to minimize signal losses. The 8192-element digital memory used to store the ADC results is not shown, but located to the left of the ADC. The remainder of the die is used to test separate IP.

The layout of each sub-ADC is shown at the left of Fig. 4.38. While the LO clock receiver is located to the left of the four sub-ADCs, the LO phase generator is placed centrally to balance the length of traces from the phase generator to the mixer switches and reduce systematic phase shift. High levels of harmonic rejection are still achievable in the presence of mixer mismatch, however, because the recombination weights can be tuned digitally. After the input signal is sampled onto the baseband capacitor via mixer switches near the LO phase generator, the low-frequency baseband signal is routed to the input of the four differential DAC banks. As shown in Fig. 4.25, the FIR control signal generation logic is placed to the right of the DAC banks. Due to the large capacitance required to keep sampled thermal noise low and the area overhead of sampling logic within each unit cell, the DAC banks consume the largest portion of the active die area. Power switches to alternate between DAC reference voltages are located at the top and bottom of the DAC bank. The SAR comparator/logic and VCO are placed near the baseband capacitor. During VCO-only operation, the down-converted baseband signal will be close to the VCO preamplifier input. Each sub-ADC also contains a set of 16-way deserializers to interface with the digital memory, which is sampled at $f_{LO}/16$. The remainder of the receiver area is used for decoupling capacitance and scan signal routing.

4.4 Measurement results

This section presents measurements of the architecture-configurable receiver prototype detailed in the previous section. First, Section 4.4.1 describes the measurement and test configuration. The measured receiver performance in the VCO-only, SAR-only, and VCO+SAR configurations is then compared in Sections 4.4.2, 4.4.3, and 4.4.4. Section 4.4.5 then concludes by discussing the advantages and disadvantages of these three configurations and comparing the results to previously published RF-to-digital receivers.

4.4.1 Measurement setup

Figure 4.39 illustrates the hardware setup used to characterize the test chip. High-speed signal generators provide the RF input signal (centered around f_{LO}) and RF sampling clock (operating at $4f_{LO}$). All DC supplies and bias voltages are generated using off-chip low-dropout regulators (LDOs) on a separate PCB. The on-chip scan signal interface and digital memory are accessed with an Opal Kelly FPGA, which is controlled on a test laptop via Python. This FPGA also programs the LDO voltages and reads the current draw of each

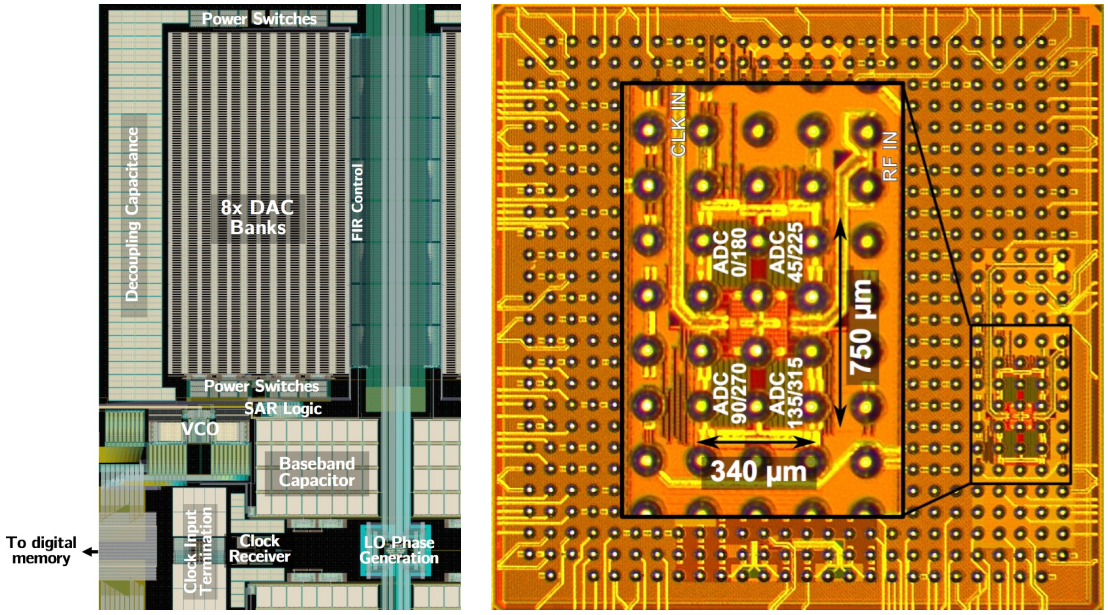

(a) Single sub-ADC layout. (b) 16nm CMOS die micrograph.

Figure 4.38. Top-level chip implementation.

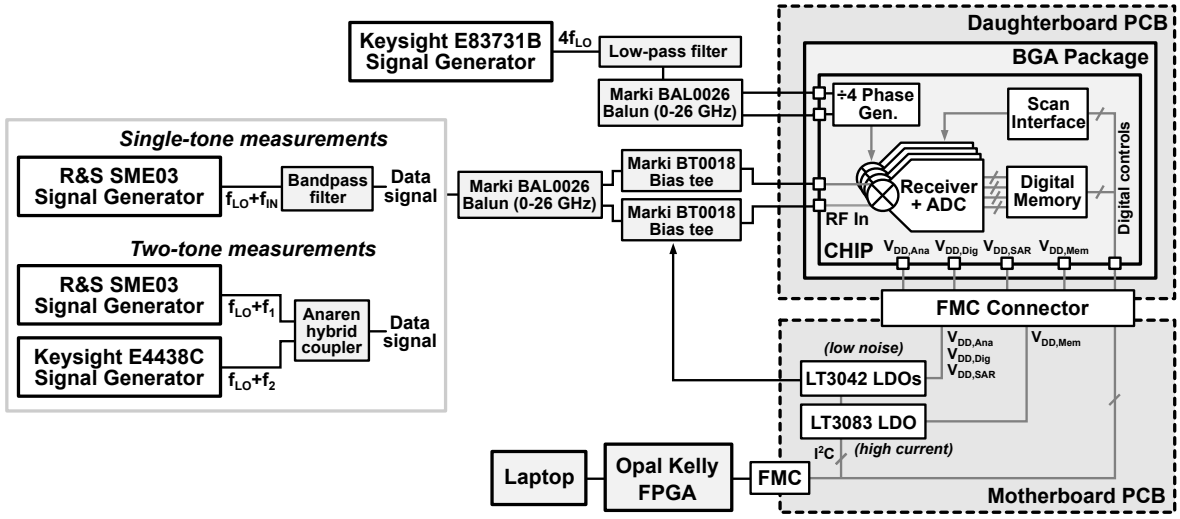

Figure 4.39. Receiver measurement setup.

supply to measure the power breakdown of the chip.

As shown in Fig. 4.39, the 5 mm × 5 mm die is bonded to a ball grid array (BGA) package. To simplify testing multiple chips, this package is soldered to a daughterboard PCB containing SMA connectors for high-speed signal inputs and a high pin count (HPC) FPGA mezzanine card (FMC) for connecting low-frequency signals to a motherboard PCB. The RF clock and signal traces on the daughterboard are designed to provide 100 Ω differential impedance. The daughterboard also contains test structures that can be used to characterize the attenuation of the SMA connectors and PCB trace. The FMC connector feeds DC voltages from LDOs on the motherboard to the test chip, and connects digital signals from the test chip to an FPGA connector on the motherboard. The motherboard contains I^2C-programmable LDOs and current sensors; the low-noise LT3042 regulator generates the analog supplies, while the high-power LT3083 regulator provides the chip-level V_{DD}. The digitally programmable AD5170 and AD5274 potentiometers are used to set the bias voltages, while current is measured using the MAX9611 amplifier and ADC.

High-frequency baluns are used to generate the differential RF clock and input signals from the single-ended signal generator output. Two bias tees set the common mode input voltage of the RF data; the DC bias is generated by an LDO on the motherboard PCB. The RF clock is AC coupled and does not require DC biasing. Losses through the SMA cables, filters, and PCB routing are measured with a spectrum analyzer to accurately estimate the signal power at the input of the test chip. The following sections describe in further detail how the measurement setup is configured for both single-tone and multi-tone tests.

4.4.1.1 *Single-tone tests*

Single-tone tests are used to characterize the sensitivity, maximum resolution, and frequency response of the receiver. Both the sensitivity and resolution of the ADC are measured from the signal to noise and distortion ratio (SNDR) of the ADC. A 5000-point FFT of the ADC output data stored in the on-chip memory is used to find the SNDR; the SAR and VCO results are scaled and recombined off-chip in MATLAB. The sensitivity is characterized by finding the input power that provides 0 dB SNDR, while the peak SNDR is obtained by sweeping the maximum input power to find the optimal balance between distortion and thermal noise. The frequency response of the receiver is measured directly by sweeping the frequency of the input signal. For each input tone, the RMS amplitude of the received digital sequence is measured to obtain the signal strength. Large input amplitudes are used to characterize the frequency response to boost the SNR of each sample.

To prevent FFT spectral leakage, the input tone frequency (offset from f_{LO}) is chosen from the ADC sampling frequency (f_{samp}), number of FFT points (N_{pts}), and a prime number p as follows:

$$f_{sig} = f_{samp}\left(\frac{p}{N_{pts}}\right)$$

Due to the high flicker noise corner in deeply scaled CMOS process nodes, $f_{sig} \approx 8$ MHz. This allows the 10-20 MHz signal bandwidth to exclude low-frequency noise, and includes the effects of second harmonic distortion. Noise and distortion from the signal generators are mitigated using a low-pass filter on the clock source and a narrow band-pass filter on the RF input signal, as shown in Fig. 4.39.

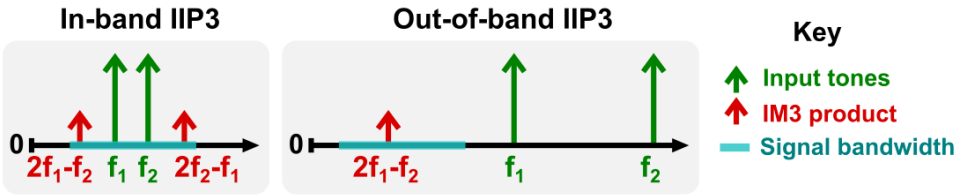

Figure 4.40. In-band and out-of-band IIP3 measurement configuration. Frequencies listed are relative to the LO frequency.

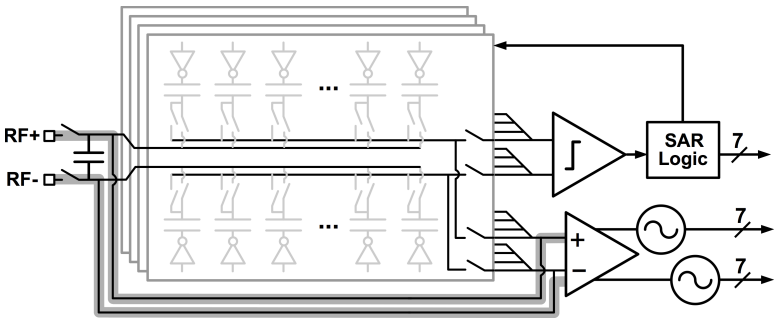

Figure 4.41. Receiver configuration in VCO-only mode. All cells in the DAC are disabled and the baseband signal is fed directly to the VCO.

4.4.1.2 *Multi-tone measurements*

Two-tone tests are used to measure the IIP3 of the receiver. As shown in Fig. 4.39, signals from multiple generators are combined using a hybrid coupler. The in-band IIP3 is measured by placing both signals within the 10-20 MHz signal bandwidth, and measuring the power of their third-order intermodulation products (IM3) at $2f_1 - f_2$ and $2f_2 - f_1$. For out-of-band IIP3 measurements, tones are offset > 40 MHz from the signal band at frequencies that will generate an in-band IM3, as illustrated in Fig. 4.40. The IIP3 is measured from the known input tone power P_{in} and relative power of the IM3 products ΔP as:

$$IIP3 = P_{in} + \frac{\Delta P}{2}$$

The power of the input tones must be strong enough to generate detectable IM3 products. Because this resolution is limited by the SNDR of the ADC and the number of FFT samples, the input power must approach the compression point of the receiver, which may attenuate the input signal and underestimate the IIP3 of the receiver.

4.4.2 VCO mode

The VCO-only mode can provide improved sensitivity and low power consumption when high linearity is not required. In this setting, the SAR ADC bypass switches are closed and each unit cell in the sampling DAC is disabled, as shown in Fig. 4.58. As a result, the down-converted baseband signal stored on the sampling capacitor is fed directly to the VCO input. To bias the VCO preamplifier near mid-rail for improved G_m and lower input-referred noise, a moderately high common mode input voltage (V_{CM}) must be applied to

the RF input signal. While this reduces the linearity of the NMOS sampling switches in the mixer, the nonlinearity of the VCO dominates the receiver's linearity. This relaxed linearity requirement also enables the receiver to operate with a low supply voltage (600 mV) to reduce power without degrading IIP3. As shown in Fig. 4.35, operating at a lower supply voltage also reduces the input-referred jitter of the oscillator.

In the VCO-only test configuration, the cutoff frequency of the first-order filter in the sampling network is programmed by setting the baseband capacitor size. The input matching resistors are fixed at $\approx 100\Omega$ using three of the parallel unit cells described in Section 4.3.2. The VCO preamplifier bias is tuned to optimize the sensitivity of the VCO. No offset is applied by the offset cancellation DAC, as the large preamplifier input devices have a relatively small input-referred offset.

4.4.2.1 *Output spectrum*

To demonstrate the achievable resolution of the ADC, Fig. 4.42 shows the measured ADC output spectrum in VCO-only mode for a peak SNDR of 52 dB. In this 5000-point FFT, the input power is -28 dBm and $f_{LO} = 1700$ MHz. The first-order noise shaping of the VCO is evident, with the shaped noise falling below the noise floor at frequencies under 40 MHz. This fixed noise floor is dictated by the input-referred thermal noise of the preamplifier and ring oscillator; flicker noise is also present at low frequencies. As a result, a 7-17 MHz IF signal bandwidth is considered in order to eliminate flicker noise. A longer FFT could more clearly illustrate low-frequency noise, but the length of the FFT is restricted by the size of the on-chip memory bank.

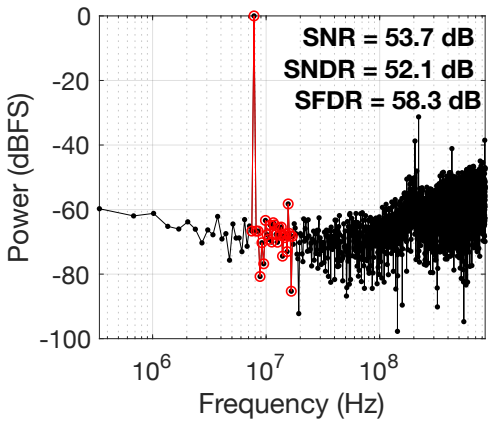

Figure 4.42. Frequency spectrum of receiver in VCO-only mode ($f_{LO} = 1700$ MHz).

The frequency spectrum in Fig. 4.42 also shows second-order harmonics due to mismatch between the two VCOs in the pseudo-differential structure, and a small amount of third-order harmonic distortion due to the nonlinear characteristic of the VCO. Spurs at multiples the SAR ADC sampling frequency of 212 MHz ($f_{LO}/8$) are caused by the switching activity of the offset cancellation DAC.

The 45° phase offset between the three sub-ADCs causes the SNDR improvement from averaging to be lower than the SNDR improvement that could be obtained without this phase shift. Figure 4.43 shows the FFTs of all three sub-ADCs, illustrating that the averaging

provides roughly 3 dB SNDR benefit relative to the best stand-alone design instead of the 4.8 dB benefit expected from averaging three in-phase signals with uncorrelated noise. While this phase shift degrades the resolution of the converter, it is required to support harmonic rejection. Additionally, Fig. 4.43 demonstrates that a third harmonic tone is present in the individual sub-ADCs that is eliminated after signal averaging. Calibration techniques could be used to compensate for the nonlinearity of the VCO to further improve performance.

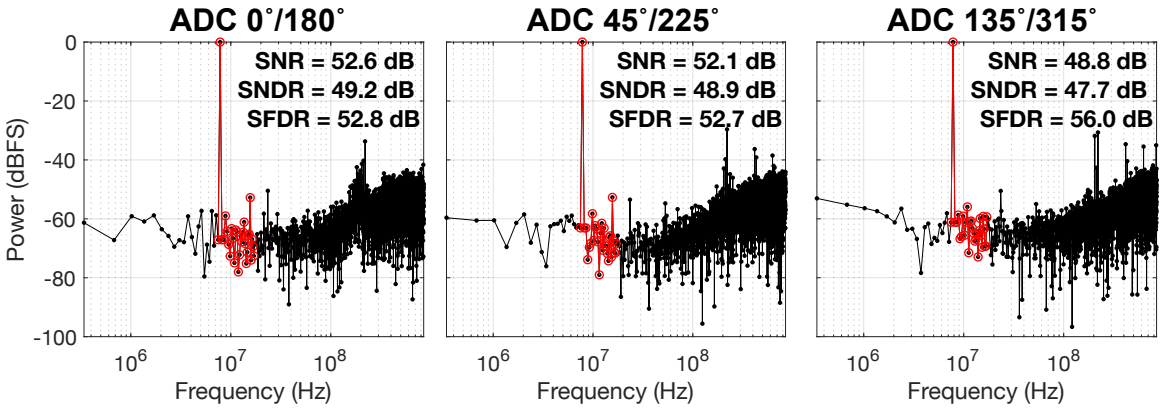

Figure 4.43. Frequency spectrum of receiver in VCO-only mode for all sub-ADCs (f_{LO} = 1700 MHz).

Finally, to characterize the linearity of the VCO, Fig. 4.44 shows an FFT of the receiver output when driven by two in-band or two out-of-band tones. In these measurements, the LO frequency is fixed at 1700 MHz, the analog V_{DD} = 0.6 V and the digital V_{DD} = 0.7 V. Figure 4.44(a) demonstrates an in-band IIP3 of -4.9 dBm, while Fig. 4.44(b) demonstrates an out-of-band IIP3 of +11 dBm, measured using tones offset 50 MHz from the carrier frequency. The IM2 measurements illustrate an in-band IIP2 of +16.6 dBm and out-of-band IIP2 of +35.6 dBm. The linearity could be improved further digitally, since each of the two oscillators in the pseudo-differential VCO are measured separately to compensate for mismatch between the two oscillators.

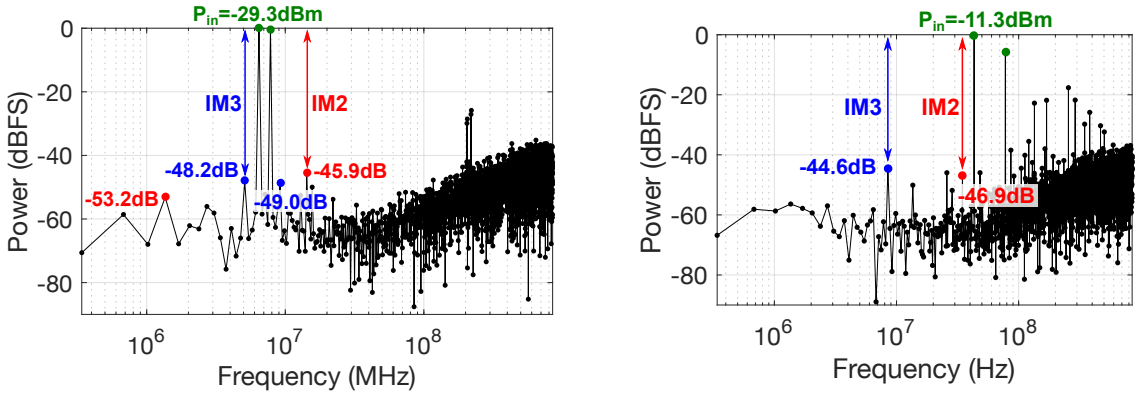

(a) In-band IIP3 measurement.

(b) Out-of-band IIP3 measurement.

Figure 4.44. Linearity measurements in VCO-only mode.

4.4.2.2 *Sensitivity & resolution*

Figure 4.58 shows the measured SNDR in 10 MHz, 20 MHz, and 40 MHz bandwidths as a function of input power for $f_{LO} = 1700$ MHz. At this carrier frequency, the sensitivity of the receiver (approximated as the 0 dB SNDR input power) is -85 dBm, -83 and -80 dBm in 10, 20, and 40 MHz bandwidths, respectively. The peak SNDR is 52, 50, and 48 dB for these bandwidths. The SNDR scales linearly with P_{in} because the noise power is independent of input signal strength. As P_{in} approaches -30 dBm, the distortion introduced by the VCO receiver dominates the SNDR. For P_{in} above -28 dBm, the harmonic distortion grows sharply and degrades the SNDR.

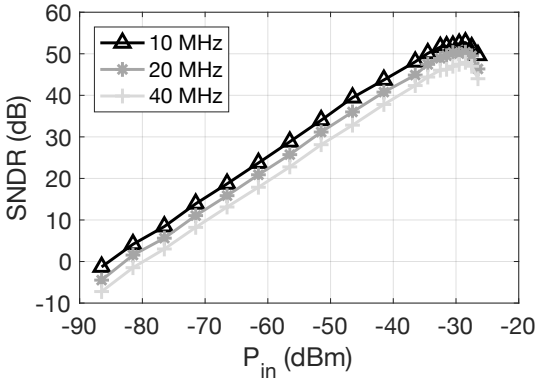

Figure 4.45. P_{in} vs. SNDR in VCO-only mode in 10 MHz, 20 MHz, and 40 MHz bandwidths.

Only bandwidths up to 40 MHz are characterized in Fig. 4.58 because above this range, shaped quantization noise will enter the signal bandwidth. This is set by the thermal noise floor and the shaped noise amplitude, as shown in Fig. 4.42. Up to 40 MHz, the receiver maintains the expected 3 dB degradation in SNDR as the bandwidth doubles, increasing the total in-band noise. Supporting higher bandwidths would require an alternative loop filter topology. Moreover, the mixer's baseband capacitance would have to scale to smaller values to move the cutoff frequency of the mixer, as discussed next.

4.4.2.3 *Filter response*

Figure 4.46 shows the measured frequency response of the receiver in VCO-only mode for both the maximum-bandwidth (smallest C_{BB}) and minimum-bandwidth (largest C_{BB}) settings. The measurements are compared to the calculated response of single-pole filters with varying bandwidths to demonstrate that the 3 dB bandwidth scales from roughly 14 MHz to 36 MHz. The sinc filter provides a slight amount of additional filtering at higher input frequencies.

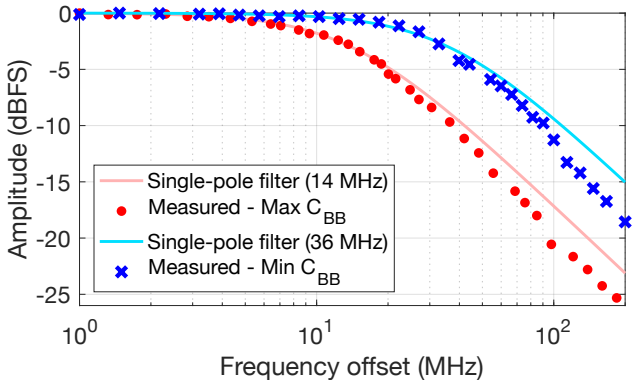

Figure 4.46. Frequency response of receiver in VCO-only mode ($f_{LO} = 1700$ MHz).

Because the VCO-only mode bypasses the configurable FIR filter in the SAR ADC, it experiences predominantly only first-order filtering from the matching resistor and baseband sampling capacitor. As demonstrated, the integrating nature of the VCO contributes a slight amount of sinc filtering, though the short integration period of the filter ($1/f_{LO}$) limits the rejection it provides. While the modest amount of filtering restricts the ability of the receiver to handle large blockers, the VCO-only mode is designed to handle low-power input signals. Moreover, a high-order anti-alias filter is not required in this setting because the VCO sampling frequency of f_{LO} is much larger than the signal bandwidth.

4.4.2.4 *Wideband performance*

To illustrate the wideband nature of the receiver, Fig. 4.47 shows the peak SNDR, sensitivity, and IIP3 of the receiver in VCO-only mode as a function of f_{LO}. Because lowering f_{LO} reduces the ADC sampling frequency, it increases the total in-band thermal noise by decreasing the oversampling ratio of the ADC. However, reducing f_{LO} does not degrade the SQNR of the converter, because the VCO phase is subsequently integrated over a longer sample period to increase the effective quantizer resolution (as discussed in Section 4.2.2). While the ADC is sampled at f_{LO} in this prototype because packaging constraints supported only a single input clock, the ADC sampling frequency could be de-coupled from the LO clock to maintain a constant oversampling ratio and digital power consumption with RF frequency.

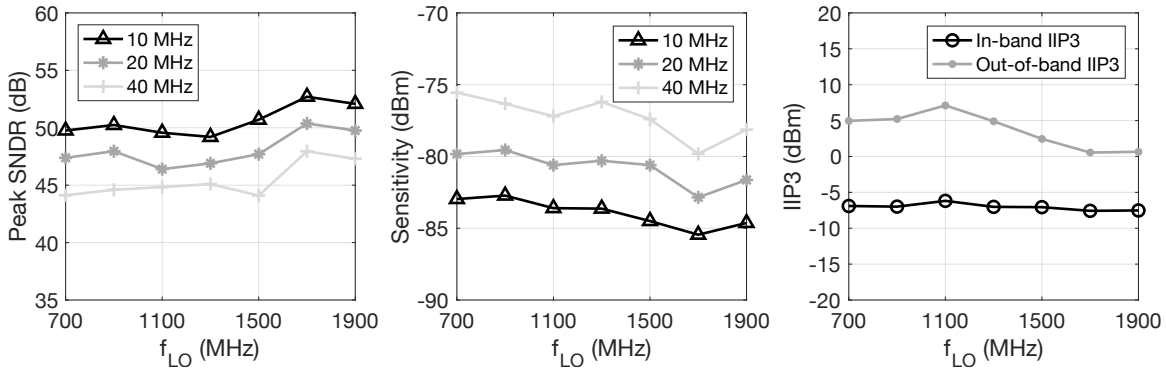

Figure 4.47. Peak SNDR, sensitivity, and linearity of the receiver in VCO-only mode.

Figure 4.47 shows that the peak SNDR is consistently kept above 49, 46, and 44 dB in 10, 20, and 40 MHz bandwidths (respectively) across the RF frequency range. Correspondingly, the receiver sensitivity is kept below −83/−80/−75 dBm in 10/20/40 MHz bandwidths. The in-band IIP3 remains relatively constant at -7 dBm for all LO frequencies, while the out-of-band IIP3 is kept above +0 dBm. All measurements were taken for an analog V_{DD} of 600 mV and a digital V_{DD} of 700 mV. As shown in Fig. 4.36, raising V_{DD} to 800 mV can actually lead to higher input-referred jitter, and the total input-referred noise is heavily bias-dependent. In turn, the sensitivity measurements are susceptible to variation, as bias parameters were tuned for each LO frequency. The peak SNDR measurements, however, are largely distortion-limited and therefore less susceptible to bias point.

4.4.2.5 *Power consumption*

Finally, the power consumption of the receiver in VCO-only mode is presented in Fig. 4.48; this figure shows both a breakdown of the receiver power at 1700 MHz and a plot of the power consumption as a function of LO frequency. Because nearly 60% of the power is digital, a significant portion of the power scales directly with f_{LO} (which sets the sampling frequency of the ADC). A fixed amount of static power is consumed by the active preamplifiers in the VCO and the LO buffer. Because less than a quarter of the receiver power is active power from the VCO, additional power could be allocated to the VCO to reduce input-referred noise and improve sensitivity without substantially impacting the total power consumption. Alternatively, simplifying the digital decoding circuitry could greatly reduce the total receiver power.

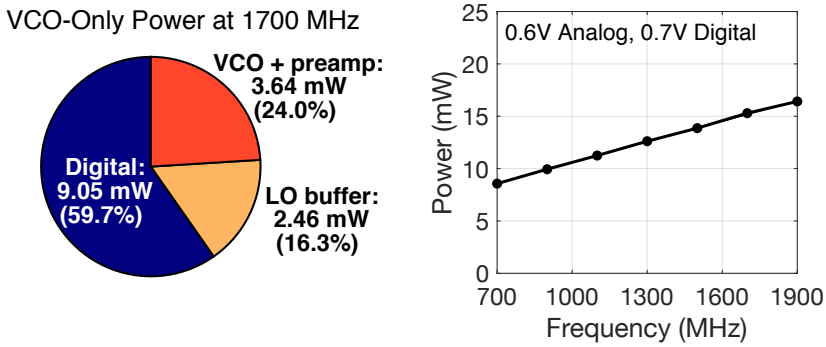

Figure 4.48. Power consumption of the receiver in VCO-only mode.

4.4.3 SAR mode

The SAR-only mode provides high linearity due to the capacitive input sampling but poor SNDR due to the large quantization noise from the 7-bit DAC. Because the quantization noise exceeds the DAC sampling noise, cells in the DAC can be disabled as shown in Fig. 4.51 to reduce the total capacitance and lower the power consumption by reducing the number of switched cells. To maintain 7 effective bits of resolution using fewer than 129 DAC unit cells, the DAC weights can be configured with radix <2 with two redundant bits. The FIR filter integrated into the SAR sampling provides an additional degree of anti-alias filtering. Reducing the number of unit cells simply limits the configurability of this filter.

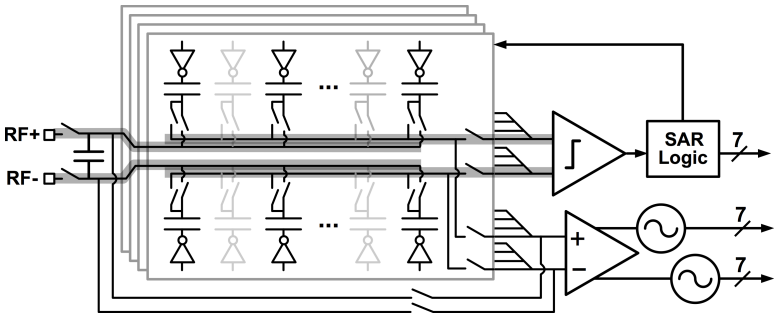

Figure 4.49. Receiver configuration in SAR-only mode. Some DAC cells can be disabled to lower power consumption.

Because the quantization noise is large, more distortion from the input sampling network is tolerable, which allows the ADC to operate at a lower supply voltage of 600 mV. This enables a reduction in power consumption without significantly degrading performance. Moreover, a low input V_{CM} can be used in this configuration to improve linearity because a PMOS-input comparator is used in the SAR ADC. While the input-referred noise of the comparator will be highest when V_{CM} is low, the SAR switching algorithm can be configured to bias the comparator at higher V_{CM} in the final ADC decisions to lower input-referred noise. As discussed in the previous chapter, however, the switching algorithm with variable V_{CM} will contribute missing codes that degrade the differential nonlinearity of the ADC.

4.4.3.1 *Output spectrum*

To characterize the performance of the receiver in SAR-only mode, Fig. 4.50 presents a FFT of the measured receiver output for an input power of -5 dBm with an LO frequency (f_{LO}) of 1700 MHz. Because the SAR ADC sampling frequency is $f_{LO}/8$ but the digital output is sampled at f_{LO}, the effective FFT size is limited to 625 points by the 5000-point digital memory. This results in a high noise floor that limits the detectable SFDR.

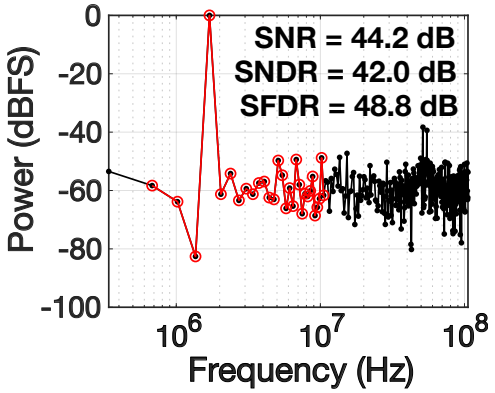

Figure 4.50. Frequency spectrum of receiver in SAR-only mode (f_{LO} = 1700 MHz) for a non-binary radix.

From Fig. 4.50, we can see that the ADC has an SNDR of 42 dB in a 10 MHz bandwidth. As expected, the frequency spectrum is relatively flat because quantization noise dominates.

Out-of-band spurs at $f_{LO}/32$ are caused by mismatch between the four sub-DACs in each ADC. Smaller in-band tones may be caused by missing codes in the transfer function of the ADC, possibly due to the variable common mode SAR switching algorithm. As previously demonstrated in Fig. 3.46, such nonlinearity can introduce spurious tones.

4.4.3.2 *Sensitivity & resolution*

To illustrate the achievable dynamic range and sensitivity of the receiver, Fig. 4.51 shows the receiver SNDR as a function of input power for both binary-weighted and non-binary DAC configurations in a variety of signal bandwidths. The LO frequency was fixed at 1.7 GHz, giving a SAR ADC sample rate of 212.5 MS/s. In these measurements, the analog $V_{DD} = 600$ mV and the digital $V_{DD} = 800$ mV.

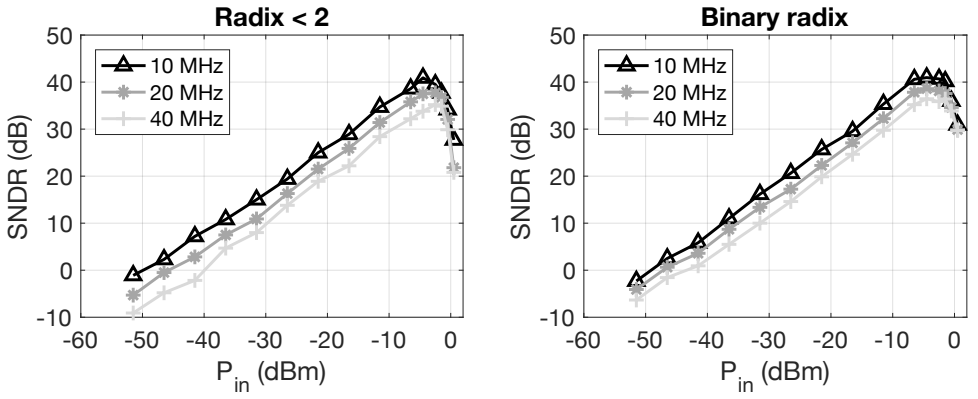

Figure 4.51. P_{in} vs. SNDR in SAR-only mode, measured in 10 MHz, 20 MHz, and 40 MHz bandwidths. The left plot shows the response when half of the unit cells are disabled to implement a sub-binary radix DAC, and the right plot shows the conventional binary radix DAC results with all unit cells enabled.

For the non-binary radix configuration, the peak SNDR is 40, 38, and 35 dB in 10, 20, and 40 MHz bandwidths. The 0 dB SNDR point (sensitivity level) is measured at input powers of -50, -46, and -40 dBm for these same bandwidths. In the 40 MHz bandwidth, spurious tones due to the converter's differential nonlinearity are more likely to fall in-band and reduce the SNDR. This leads to over a 3 dB difference between SNDRs in the 20 MHz and 40 MHz bandwidth configurations. If the noise floor were completely flat with frequency, a 3 dB SNDR difference would be expected. The main advantage of the binary radix configuration is improved linearity. The peak SNDR in binary-weighted mode is 40/38/35 dB in the 10/20/40 MHz bandwidths, and the sensitivity is -50, -48, and -42 dBm. The performance is nearly identical, illustrating that using fewer DAC sells to lower power consumption does not significantly degrade performance.

4.4.3.3 *Filter response*

The frequency response of the FIR filter is presented in the next section with the SAR+VCO mode characterization (Fig. 4.59). In SAR+VCO mode, the VCO does not provide additional sinc filtering because the VCO-based ADC operates on the sampled ADC conversion residue instead of the full input signal. As a result, enabling the VCO provides higher dynamic range to improve the measurement resolution but does not affect the frequency

response. Figure 4.59 shows that the receiver can provide over 50 dB of passive anti-alias filtering.

4.4.3.4 *Wideband performance*

To characterize the wideband performance, Figs. 4.52 and 4.53 present the measured peak SNDR and sensitivity in SAR-only mode as the LO frequency is swept from 700 MHz to 1.9 GHz. Figures 4.52 and 4.53 show measurements for the SAR ADC with binary-weighted (radix = 2) and non binary-weighted (radix < 2) DAC configurations at analog supply voltages of 600 mV and 800 mV.

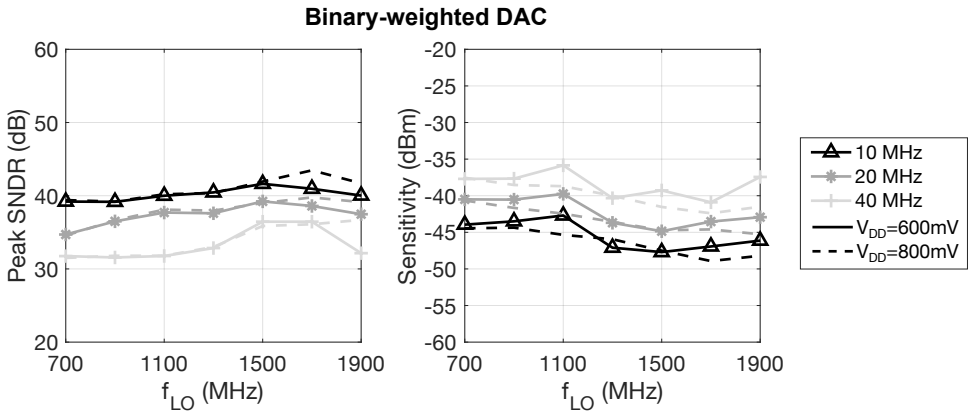

Figure 4.52. Peak SNDR and sensitivity of the receiver in SAR-only mode with binary-weighted DAC for $V_{DD} = 600$ mV and 800 mV.

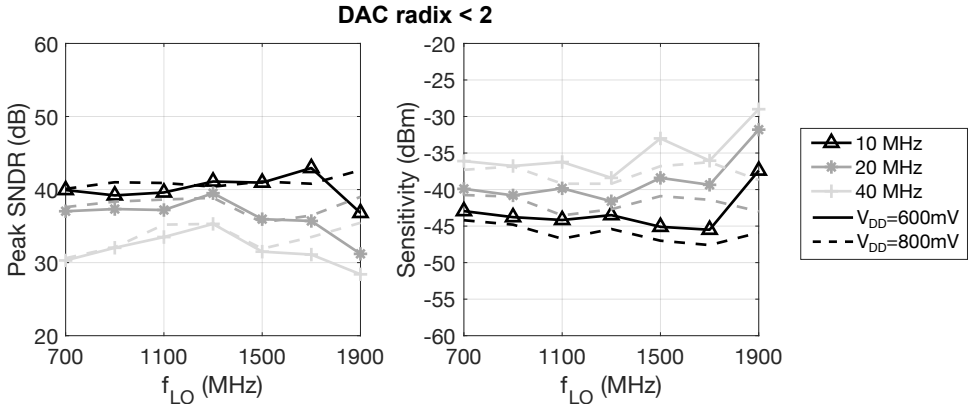

Figure 4.53. Peak SNDR and sensitivity of the receiver in SAR-only mode with sub-radix 2 DAC for $V_{DD} = 600$ mV and 800 mV.

Using the conventional binary-weighted DAC, the maximum achievable SNDR across all frequencies is 40/35/31 dB in 10/20/40 MHz bandwidths. This result is independent of supply voltage. With the alternative DAC radix, the peak SNDR is 35/31/28 dB on a 600 mV supply voltage and 40/35/30 on an 800 mV supply voltage. While the SNDR is relatively independent of supply voltage at most frequencies, as the LO frequency increases, a higher V_{DD} becomes necessary to improve resolution. This is because the higher voltage

becomes necessary to support faster sampling rates associated with higher LO frequencies. As the sample rate increases, less timing margin is available to allow for complete DAC settling, and comparator metastability will become increasingly likely. This dependence on LO frequency can also be observed in the sensitivity measurements. While the maximum sensitivity is -30/-32/-37 dBm on a 600 mV supply using the sub-binary radix DAC, it improves to -36/-40/-45 dBm using the 800 mV supply. Using the binary-weighted DAC, the sensitivity is -36/-40/-43 dBm on a 600 mV supply and -37/-40/-45 dBm on an 800 mV supply. As with the peak SNDR measurements, the receiver performance is more consistent across supply voltage using the binary-weighted DAC. Next, Fig. 4.54 shows how the IIP3 varies with LO frequency.

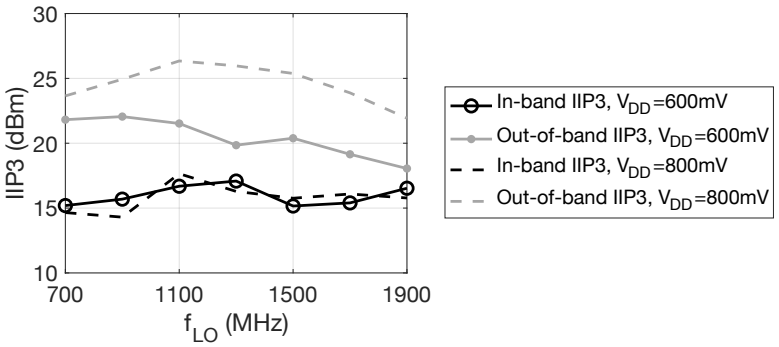

Figure 4.54. Linearity of the receiver in SAR-only mode with binary-weighted DAC for $V_{DD} = 600$ mV and 800 mV.

The results in Fig. 4.54 are obtained using a conventional binary-weighted DAC. It illustrates relatively flat and V_{DD}-independent performance across LO frequency for the in-band IIP3, but frequency and supply dependent performance for the out-of-band IIP3. The in-band IIP3 measurement accuracy is limited by the size of the FFT. To generate detectable IM3 products, the input power must be high, which can underestimate IIP3. The out-of-band IIP3 measurements, however, depend on the degree of filtering provided by the receiver. Because the nulls of the discrete-time FIR filter occur frequencies set by the FIR sampling frequency, changing LO frequency affects the rejection, in turn altering the IIP3. As expected, increasing the supply voltage improves linearity.

4.4.3.5 *Power consumption*

To conclude the SAR ADC measurements, Fig. 4.55 summarizes the active power consumption of the receiver in this mode. Because the SAR ADC and FIR filter sampling is fully digital, the only analog component of the receiver power is in the LO buffer. Therefore, a significant fraction of the total power consumption is frequency-dependent.

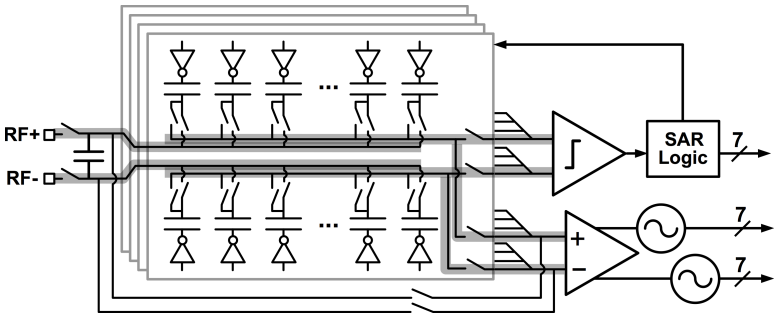

Figure 4.56. Receiver configuration in SAR+VCO mode. All features are enabled.

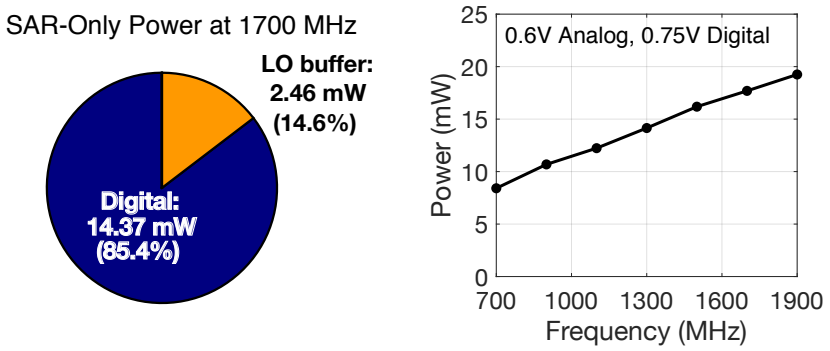

Figure 4.55. Power consumption of the receiver in SAR-only mode.

Most of the digital power is due to the FIR filter sampling, given that each DAC cell can be sampled during an independently programmable index. This requires digital logic gates in each of the 128 cells that operates at the LO sampling frequency, driving up the digital power consumption. Moreover, because these cells must be distributed across a large area of the chip to construct a large DAC with low thermal sampling noise, many routing parasitics exist in the design to drive up the digital power. In future work, fixing the anti-alias filter sampling procedure could help save power at the expense of configurability.

4.4.4 SAR+VCO mode

The combined SAR+VCO mode provides both high resolution and high linearity at the expense of power consumption. As shown in Fig. 4.56, all cells in the DAC must be enabled to minimize sampled thermal noise. Because the VCO-based ADC is used to quantize the SAR residue, the SAR resolution does not need to be high — this configuration is resilient to the SAR quantization noise and comparator decision errors. The redundant bits in the SAR ADC are only needed if the amplitude of the conversion residue saturates the input range of the VCO.

To maximize the linearity of the NMOS sampling switches, the analog supply voltage is set to 800 mV. The common mode voltage of the RF input signal is kept low ($\approx$ 150 mV) to improve linearity, given that the SAR switching algorithm can be configured to bias the VCO-based ADC near mid-rail for optimal sensitivity. The digital supply voltage may be lowered when f_{LO} is small, but must be high enough to operate the FIR filter sampling, SAR logic, and VCO decoding at a higher rate when f_{LO} increases. It is kept to 800 mV for

these measurements.

4.4.4.1 *Output spectrum*

To first illustrate the performance of this converter, Fig. 4.57 presents an FFT of the receiver output measured for an input power of -6 dBm at an LO frequency of 1700 MHz. While the VCO-only receiver was linearity-limited at input powers of -28 dBm, the addition of the SAR ADC maintains a relatively flat frequency response even for -6 dBm input signals. The high-frequency noise shaping afforded by the VCO can also be observed, though to a smaller extent than the VCO-only converter because only the SAR quantization residue is noise-shaped. Additionally, spurious tones are observed at multiples of the SAR ADC sampling frequency (212.5 MHz) and the sub-DAC switching frequency of 53 MHz. By oversampling the ADC, these tones remain outside of the desired 10 MHz signal bandwidth.

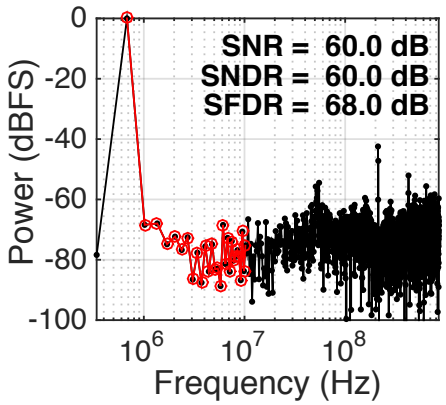

Figure 4.57. Frequency spectrum of receiver in SAR+VCO mode (f_{LO} = 1700 MHz).

4.4.4.2 *Sensitivity & resolution*

To summarize the achievable resolution and sensitivity of the receiver at 1700 MHz operation, Fig. 4.58 shows the measured SNDR as a function of P_{in} for the SAR+VCO receiver in 10, 20, and 40 MHz bandwidths. This demonstrates a peak SNDR of 60, 57, and 50 dB in these bandwidths. The 40 MHz performance begins to be limited by the tones generated by mismatch between sub-DACs within each ADC. The corresponding sensitivity (0 dB SNDR point) is limited to -75/-72/-69 dBm for the bandwidths considered. Thermal noise introduced during the input sampling process reduces the sensitivity of the converter. Attenuation due to charge-sharing between the baseband capacitor and the sampling capacitance results in further degradation of the receiver sensitivity in this mode of operation.

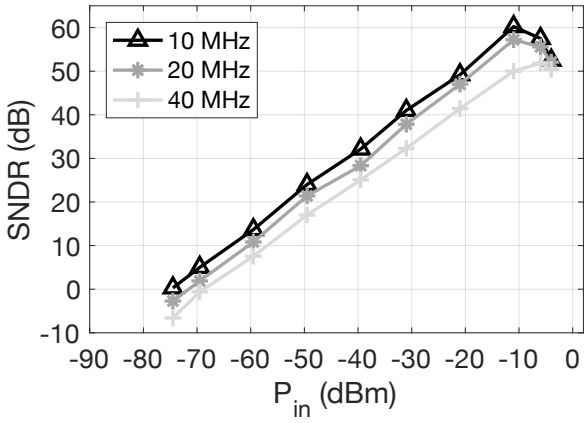

Figure 4.58. P_{in} vs. SNDR in SAR+VCO mode, measured in 10 MHz, 20 MHz, and 40 MHz bandwidths.

4.4.4.3 *Filter response*

Figure 4.59 compares the measured and calculated frequency response of the receiver in SAR+VCO mode for two different filter configurations with $f_{LO} = 1700$ MHz. Measurements are presented from a single sub-ADC; averaging between sub-ADCs will contribute further rejection at multiples of the LO frequency.

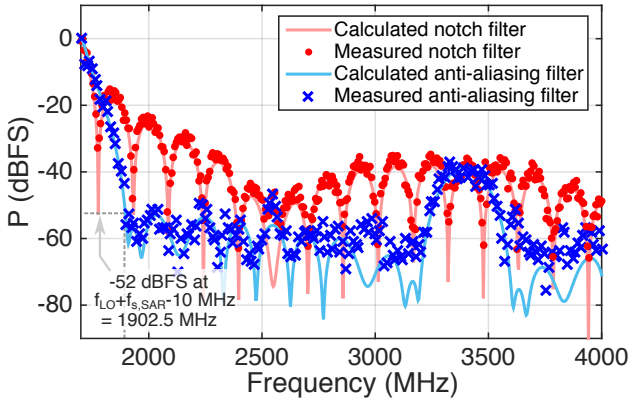

Figure 4.59. Frequency response of receiver in SAR+VCO mode ($f_{LO} = 1700$ MHz).

The experimental and theoretical frequency responses align well, though the maximum achievable SNDR of the receiver and maximum output power of the signal generator limits the measurable filter rejection. Random mismatch between DAC unit elements will also introduce variation in the filter response, as discussed in Section 4.2.1. The anti-alias filter response shown in blue illustrates that up to 52 dB anti-alias filtering is achievable. At f_{LO}, the FIR filter does not provide additional anti-aliasing and only first-order filtering from the input network provides signal rejection. However, harmonic rejection from the 8-phase mixer provides additional filtering of signals centered around f_{LO}. The notch filter shown in red illustrates that the filter response can be configured to place nulls closer to the target signal bandwidth.

4.4.4.4 *Wideband performance*

To illustrate the wideband performance of the receiver, Fig. 4.60 shows the sensitivity (P_{in} corresponding to 0 dB SNDR) as a function of LO frequency in both a 10 MHz and 20 MHz bandwidth. In a 40 MHz bandwidth, tones at the sub-DAC switching frequency ($f_{LO}/8/4$) fall in-band for a low LO frequency of 700 MHz, so this bandwidth is not presented in the measurements.

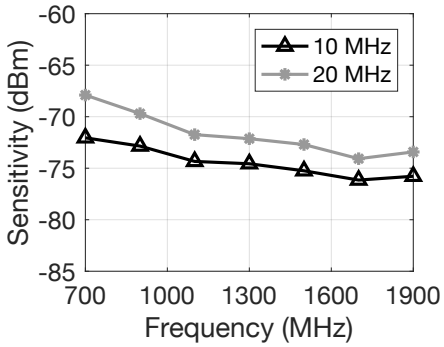

Figure 4.60. Sensitivity of the receiver in SAR+VCO mode.

The sensitivity of the receiver is maintained below -72/-68 dBm in the 10/20 MHz bandwidths considered. Because the ADC sample rate is tied to the LO frequency but the signal bandwidth is fixed, the sensitivity of the receiver degrades with LO frequency due to the lower oversampling rate that results in a larger fraction of in-band noise.

4.4.4.5 *Power consumption*

The receiver's power consumption in SAR+VCO mode is presented in Fig. 4.61. To support a higher dynamic range and improve linearity, the receiver operates on an 800 mV analog and digital V_{DD} in this mode. Similar to the VCO-only mode, nearly 60% of the power is digital, and therefore frequency-dependent. Increasing the analog V_{DD} to 800 mV results in a larger fraction of power consumed by the LO buffer. Because the clock is generated with digital logic gates, the power consumption will scale with V_{DD}^2.

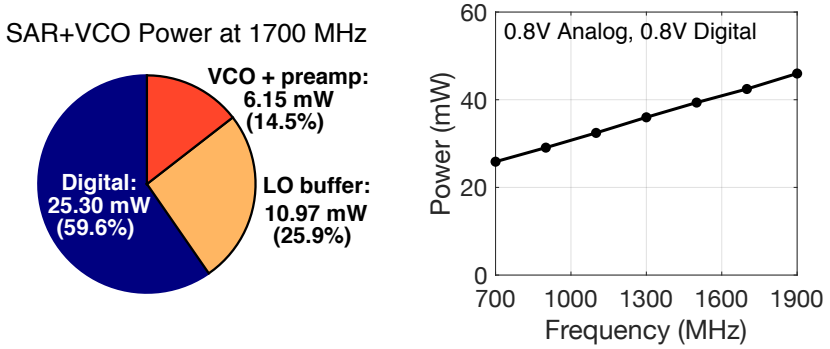

Figure 4.61. Power consumption of the receiver in SAR+VCO mode.

In this mode of operation, only 14.5% of the total less than a quarter of the receiver power is active power from the VCO, additional power could be allocated to the VCO to

	Englund, ISSCC '14	Wu, JSSC '14	Kim, ISSCC '11	This work		
Arch.	ΣΔ	ΣΔ	VCO	VCO	SAR	SAR+VCO
RF freq. (GHz)	0.7-2.7	0.4-4	0.2-1.8	0.7-1.9		
BW (MHz)	15	10	10	10		
SNDR (dB)	43	65-52	50	47	42	60
IB-IIP3 (dBm)	-	+10	-6.8	-7.6	+15.8	+17.2
OB-IIP3 (dBm)	-2	+13.5	-	+5.3	+16.3	+19.5
Sensitivity (dBm)	-89	-88	-86	-82	-44	-72
Noise figure (dB)	5.9-8.8	16	18	20-22	55-60	28-32
Uses LNA?	Yes	Yes	Yes	No	No	No
Power (mW)	90	17-70.5	42-49	8.6-16	8.4-21	26-46
Supply (V)	1.1	1.1 / 1.5	1.35	0.6 / 0.7	0.6 / 0.75	0.8
Area (mm^2)	1.0	0.56	0.4	0.07	0.26	0.26
Tech. (nm)	40	65	90	16		

Figure 4.62. Comparison to state-of-the-art RF-to-digital receiver designs [2].

reduce input-referred noise and improve sensitivity without substantially impacting the total power consumption. Alternatively, simplifying the digital decoding circuitry could greatly reduce the total receiver power.

4.4.5 Architecture comparison

To understand the overall capabilities of this receiver, Fig. 4.62 presents a table comparing the measured performance of this RF-to-digital receiver to previously published RF-to-digital receiver designs. This illustrates that the achievable SNDR is among state of the art, and the achievable linearity is the best among prior works due to the capacitive sampling nature of the SAR ADC. However, the corresponding sensitivity of the receiver is low relative to previous work because this design has no active low-noise amplification stage. The SAR noise floor could be improved by increasing the sampling capacitance size or improving the oversampling rate at the expense of power. Because SAR ADC thermal noise is exponentially related to capacitance size, the power required to substantially lower the SAR-only mode noise will be significant. An active LNA stage would relax the required size of the capacitance, at the expense of additional power, design complexity, and likely lower linearity. To improve the minimum achievable noise figure in VCO-only mode, the VCO could simply be redesigned VCO for lower input-referred jitter. Given that the power of this design is among the lowest of previously published works, additional headroom exists to lower input-referred noise at the expense of power consumption.

4.5 Conclusions

Overall, this work shows that a digital-intensive, scalable architecture using combined SAR and VCO-based techniques can result in performance competitive with prior work, while requiring only a single self-biased amplification stage that operates down to 0.6V. The SAR sampling stage enables the design to achieve state-of-the-art linearity, while the configurability can offer over 2.5X power reduction going from SAR+VCO to VCO-only or SAR-only

mode. The minimum power consumption in VCO-only mode is the lowest of previously published RF-to-digital converters, though additional power savings could be obtained by simplifying the digital logic. The additional bits of VCO phase resolution provided by the 3-bit SAR ADC improve the in-band VCO SNDR by < 0.1 dB; this stage could be eliminated. Additional power savings in the SAR stage could be obtained by replacing the programmable FIR filter with a fixed anti-alias filter. As described in the previous chapter, configurability translates to a fixed power overhead that limits performance efficiency.

The major limitation of the existing prototype is its achievable noise figure, determined by the total input-referred noise of the converter. Without an explicit LNA stage, the sampling capacitance must be particularly large to maintain competitive noise figure. Because of the large area and power overhead of the sampling capacitance, introducing an additional amplification stage is the most practical means of improving the sensitivity in SAR+VCO mode. The noise performance of the VCO-based ADC, however, could be improved simply by redesigning the preamplifier and oscillator and moderately increasing power consumption. Future work could explore ways to reduce noise in the digital domain utilizing information from all four sub-ADCs; however, this would require detecting the desired signal phase using a reference tone. Efforts to redesign the digitally-intensive receiver architecture could look at alternative methods to optimize the linearity of the VCO-based ADC without sampling the signal onto a separate DAC.

Chapter 5

Conclusion

Configurability is a promising means of both improving the energy efficiency of wireless systems and allowing these systems to fully leverage the enhanced digital processing capabilities enabled by CMOS technology scaling. In a reconfigurable wireless system, the performance of a radio receiver could be tuned to provide high linearity and good noise figure only when required by a user's environment, and not simply to meet the worst-case conditions outlined in a wireless standard. To that end, this dissertation has described two approaches to constructing data converters for digitally-adaptive radio. In this conclusion, Section 5.1 first reviews the two techniques considered and highlights their utility in the broader context of future wireless systems. Finally, Section 5.2 details potential extensions of this research for further study.

5.1 Main contributions

The first main contribution of this work is the design and implementation of a resolution-configurable SAR ADC to serve as a building block for energy-efficient massive MIMO radio arrays with digital beamforming. To obtain adequate spatial selectivity to support multiple users, these arrays require many elements. While oversampled ADCs are often employed in single-element receivers to meet stringent dynamic range requirements, the SAR architecture is uniquely suited to the moderate to low resolution requirements of array-based receivers. SAR ADCs are also easy to implement in scaled CMOS because they can be constructed without any active gain stages. While previous work developing scalable SAR ADCs emphasizes bandwidth tuning to support multiple wireless standards and often requires supply voltage scaling, this ADC enables resolution configurability at a fixed bandwidth and supply voltage specifically for use as a common module in array-based wireless systems with a programmable number of elements. Using a tunable DAC size, comparator size, and switching algorithm, the 80 MS/s prototype ADC scales power by a factor of two from 0.4-0.8 mW while achieving 7.0-9.1 effective bits of resolution in a 10 MHz signal bandwidth. It is integrated with a scalable mixer-first receiver and implemented in 65nm CMOS. In the context of beamforming radio systems, the prototype design illustrates how overhead power consumption must be carefully minimized to maintain energy efficiency across a broad range of array sizes.

The second core contribution of this dissertation is the proposed architecture-configurable

RF-to-digital converter suitable for fully-integrated adaptive radio systems. To build complex radio systems in a monolithic IC, receiver architectures must first be easy to implement in modern CMOS process nodes without significant performance degradation. The proposed receiver design therefore uses a hybrid SAR and VCO-based ADC architecture, combining the high linearity of the capacitive sampling process in a SAR ADC with the sensitivity of a VCO-based ADC. The chosen digital-intensive design is highly process-scalable, as the only active analog component in the receiver is a self-biased preamplifier in the VCO. A single pole of passive RC filtering and a digitally-configurable passive FIR filter provides over 50 dB anti-alias filtering without any active stages. Next, by using a digitally-configurable filter and by allowing the VCO and SAR component of the receiver to be disabled, the proposed design can adapt bandwidth, noise figure and linearity to a meet a system's needs. Complex adaptive wireless systems require highly configurable receivers. The 0.7-1.9 GHz, 10-20 MHz bandwidth receiver prototype implemented in 16nm CMOS demonstrates +16 dBm in-band IIP3 and 60 dB dynamic range when the SAR ADC is enabled. Without the noise contribution of the SAR ADC, the VCO-based ADC can detect signals down to -82 dBm on a 0.6V analog supply. Through architecture configurability, the receiver power consumption can be scaled by over a factor of 2, from 8.6-16 mW in VCO-only mode to 26-46 mW in the high-performance SAR+VCO mode. Overall, this work demonstrates the potential power savings afforded by configurability and proposes a receiver architecture that is uniquely suited to implementation in scaled nodes optimized for digital circuits.

The digital-intensive receiver designs proposed in this work are increasingly valuable as CMOS technology scaling progresses. Despite predictions that Moore's law scaling will slow down, foundries are continuing to introduce CMOS technologies down to 7nm and below. While scaling offers many improvements in digital processing capabilities, reduced supply voltage, higher levels of mismatch, and complex design rules all present challenges to designing custom analog components in these nodes. Both the SAR and VCO-based ADC architectures are promising means of building interface circuits using effectively digital building blocks. Moreover, as scaling reduces the cost of on-chip digital computation, digitally-configurable receivers can use this processing power to enable smart wireless systems. This dissertation provides deeper insight into when configurability is a useful design tool for various applications.

5.2 Future work

This thesis introduces techniques for constructing digitally-configurable receivers, but further work can help refine these approaches and apply them to real-world wireless systems. In addition to the implementation-level design improvements proposed in Chapters 3 and 4 to enhance performance, continued research can demonstrate these designs in a full system-on-a-chip (SoC) with integrated calibration. These challenges could involve studying the following:

- *Rapid design implementation:* To integrate the proposed digital-intensive design with digital systems that can be rapidly generated in hardware description languages, the proposed designs could be implemented using design automation tools such as the Berkeley Analog Generator [102].

- *Learning & adaptation algorithms:* Digital algorithms are needed to detect properties of a user's environment and then modify the performance of wireless receivers accordingly to minimize power consumption. This could include tasks such as optimizing the filter topology using the architecture-configurable design in Chapter 4, or determining the proper allocation of noise between the receiver and ADC using the scalable design in Chapter 3.

- *Noise reduction algorithms:* Using many parallel ADCs (Chapter 4) or receivers (Chapter 3) has the potential to lower noise via averaging. However, any phase offset between signals added together can lower the signal-to-noise ratio improvement offered by this approach, as it lowers the effective signal strength. Alternative digital calibration techniques and training sequences may be able to recover some of this SNR degradation.

- *Quantifying overhead power:* Any digital-assist techniques used to enhance performance will consume additional power. As discussed in the previous chapter, overhead power fundamentally limits the efficiency of configurable designs. This cost must be quantified to fully evaluate the digital-intensive approach.

The work in this dissertation provides a useful framework for understanding the benefits and challenges of constructing digital-intensive receivers. To fully benefit from the system-level improvements offered by configurability, however, additional studies can be considered to demonstrate its use in a fully integrated SoC.

Bibliography

[1] K. Trotskovsky *et al.*, "A 0.25-1.7GHz, 3.9-13.7mW power-scalable, -10dBm harmonic blocker-tolerant mixer-first RF-to-digital receiver for massive MIMO applications," *IEEE Solid-State Circuits Letters*, vol. 1, February 2018.

[2] A. Whitcombe, F. Sheikh, E. Alpman, A. Ravi, and B. Nikolić, "A dual-mode reconfigurable RF-to-digital receiver in 16nm FinFET," in *IEEE Symposium on VLSI Circuits*, pp. 23–24, June 2018.

[3] A. Smith, "Record shares of americans now own smartphones, have home broadband." http://www.pewresearch.org/fact-tank/2017/01/12/evolution-of-technology/, January 2017.

[4] B. Chen, "What's the right age for a child to get a smartphone?." The New York Times, 2016.

[5] "Cisco visual networking index: Global mobile data traffic forecast update, 2016-2021," white paper, Cisco, 2017.

[6] G. Moore, "Cramming more components onto integrated circuits," *Electronics*, vol. 38, April 1965.

[7] Qualcomm, "Qualcomm Snapdragon product information." https://www.qualcomm.com/snapdragon/, 2018.

[8] A. Whitcombe, S. Taylor, M. Denham, V. Milovanović, and B. Nikolić, "On-chip random telegraph noise and i-v characterization in deeply scaled cmos," in *European Solid-State Device Research Conference*, September 2016.

[9] P. Mak and R. P. Martins, "High-/mixed-voltage RF and analog CMOS circuits come of age," *IEEE Circuits and Systems Magazine*, pp. 27–39, 2010.

[10] M. White, "Are you really ready for your next node?." https://blogs.mentor.com/calibre/blog/2017/01/11/are-you-really-ready-for-your-next-node/.

[11] N. Ding, D. Wagner, X. Chen, A. Pathak, Y. Hu, and A. Rice, "Characterizing and modeling the impact of wireless signal strength on smartphone battery drain," in *SIGMETRICS*, pp. 29–40, 2013.

[12] C. Andrews and A. Molnar, "A passive mixer-first receiver with digitally controlled and widely tunable RF interface," *IEEE Journal of Solid-State Circuits*, vol. 45, pp. 2696–2708, December 2010.

[13] J. Zhu and P. Kinget, "Frequency-translational quadrature-hybrid receivers for very-low-noise, frequency-agile, scalable inter-band carrier aggregation," *IEEE Journal of Solid-State Circuits*, vol. 51, pp. 3137–3151, December 2016.

[14] D. Murphy *et al.*, "A blocker-tolerant, noise-canceling receiver suitable for wideband wireless applications," *IEEE Journal of Solid-State Circuits*, vol. 47, pp. 2943–2963, December 2012.

[15] Y. Xu and P. Kinget, "A chopping switched-capacitor RF receiver with integrated blocker detection," *IEEE Journal of Solid-State Circuits*, vol. 53, pp. 1607–1617, June 2018.

[16] V. Giannini *et al.*, "A 2-mm^2 0.1–5 GHz software-defined radio receiver in 45-nm digital CMOS," *IEEE Journal of Solid-State Circuits*, vol. 44, pp. 3486–3489, December 2009.

[17] A. Geis, J. Ryckaert, L. Bos, G. Vandersteen, Y. Rolain, and J. Craninckx, "A 0.5 mm^2 power-scalable 0.5–3.8-GHz CMOS DT-SDR receiver with second-order RF band-pass sampler," *IEEE Journal of Solid-State Circuits*, vol. 45, pp. 2375–2387, November 2010.

[18] M. Kitsunezuka *et al.*, "A 30-MHz–2.4-GHz CMOS receiver with integrated rf filter and dynamic-range-scalable energy detector for cognitive radio systems," *IEEE Journal of Solid-State Circuits*, vol. 47, pp. 1084–1093, May 2012.

[19] X. Huang, P. Harpe, G. Dolmans, H. de Groot, and J. Long, "A 780–950 MHz, 64–146 μW power-scalable synchronized-switching OOK receiver for wireless event-driven applications," *IEEE Journal of Solid-State Circuits*, vol. 49, pp. 1135–1147, May 2014.

[20] T. Haque *et al.*, "A reconfigurable architecture using a flexible LO modulator to unify high-sensitivity signal reception and compressed-sampling wideband signal detection," *IEEE Journal of Solid-State Circuits*, vol. 53, pp. 1577–1591, June 2018.

[21] S. Li, J. Li, X. Gu, H. Wang, M. Tang, and Z. Zhuang, "A continuously and widely tunable 5 dB-NF 89.5 dB-gain 85.5 dB-DR CMOS TV receiver with digitally-assisted calibration for multi-standard DBS applications," *IEEE Journal of Solid-State Circuits*, vol. 48, pp. 2762–2774, November 2013.

[22] D. Chamla, A. Kaiser, A. Cathelin, and D. Belot, "A switchable-order G_m-C baseband filter with wide digital tuning for configurable radio receivers," *IEEE Journal of Solid-State Circuits*, vol. 42, pp. 1513–1521, July 2007.

[23] F. Behbahani, W. Tan, A. Karimi-Sanjaani, A. Roithmeier, and A. Abidi, "A broad-band tunable CMOS channel-select filter for a low-if wireless receiver," *IEEE Journal of Solid-State Circuits*, vol. 35, pp. 476–489, April 2000.

[24] D. Lin, L. Li, S. Farahani, and M. Flynn, "A flexible 500 MHz to 3.6 GHz wireless receiver with configurable DT FIR and IIR filter embedded in a 7b 21 MS/s SAR ADC," *IEEE Transactions on Circuits and Systems — I: Regular Papers*, vol. 59, pp. 2846–2857, December 2012.

[25] D. Lin, H. Chae, L. Li, and M. Flynn, "A low-power adaptive receiver utilizing discrete-time spectrum-sensing," *IEEE Transactions on Microwave Theory and Techniques*, vol. 61, pp. 1338–1346, March 2013.

[26] C. Luo, P. Gudem, and J. Buckwalter, "A 0.2–3.6-ghz 10-dbm b1db 29-dbm iip3 tunable filter for transmit leakage suppression in saw-less 3g/4g fdd receivers," *IEEE Transactions on Microwave Theory and Techniques*, vol. 63, pp. 3514–3524, October 2015.

[27] F. Rivet, Y. Deval, D. Begueret, J. andn Dallet, P. Cathelin, and D. Belot, "A disruptive receiver architecture dedicated to software-defined radio," *IEEE Transactions on Circuits and Systems — II: Express Briefs*, vol. 55, pp. 344–348, April 2008.

[28] A. Mirzaei, H. Darabi, and D. Murphy, "A low-power process-scalable superheterodyne receiver with integrated high-Q filters," *IEEE Journal of Solid-State Circuits*, vol. 46, pp. 2920–2932, December 2011.

[29] D. Daly and A. Chandrakasan, "An energy-efficient OOK transceiver for wireless sensor networks," *IEEE Journal of Solid-State Circuits*, vol. 42, pp. 1003–1011, May 2007.

[30] H. Fuketa, S. O'uchi, and T. Matsukawa, "A 0.3-V 1-μW super-regenerative ultrasound wake-up receiver with power scalability," *IEEE Transactions on Circuits and Systems — II: Express Briefs*, vol. 64, pp. 1027–1031, September 2017.

[31] S. Sen, D. Banerjee, M. Verhelst, and A. Chatterjee, "A power-scalable channel-adaptive wireless receiver based on built-in orthogonally tunable LNA," *IEEE Transactions on Circuits and Systems — I: Regular Papers*, vol. 59, pp. 946–957, May 2012.

[32] Y. Xu, B. Chi, X. Yu, N. Qi, P. Chiang, and Z. Wang, "Power-scalable, complex bandpass/low-pass filter with I/Q imbalance calibration for a multimode GNSS receiver," *IEEE Transactions on Circuits and Systems — II: Express Briefs*, vol. 59, pp. 30–34, January 2012.

[33] J. Mitola, "The software radio architecture," *IEEE Communications Magazine*, pp. 26–38, May 1995.

[34] A. Abidi, "The path to the software-defined radio receiver," *IEEE Journal of Solid-State Circuits*, vol. 42, pp. 954–966, May 2007.

[35] M. Verhelst and W. Dehaene, "A flexible, ultra-low-energy 35 pJ/pulse digital back-end for a QAC IR-UWB receiver," *IEEE Journal of Solid-State Circuits*, vol. 43, pp. 1677–1687, July 2008.

[36] X. Wang *et al.*, "1.9-2.6GHz tuning range variable gain low-noise amplifier with digital assisted automatic tuning loop," in *IEEE Radio Frequency Circuits Symposium*, 2011.

[37] D. Stewart and C. Saavedra, "Extending the bandwidth of low-noise microwave amplifier through digital assist," *Electronics Letters*, vol. 50, pp. 528–530, March 2014.

[38] R. Gangarajaiah, M. Abdulaziz, H. Sjoland, P. Nilsson, and L. Liu, "A digitally assisted nonlinearity mitigation system for tunable channel select filters," *IEEE Transactions on Circuits and Systems — II: Express Briefs*, vol. 63, pp. 69–73, January 2016.

[39] L. Zhang and H. Krishnaswamy, "A 0.1-to-3.1 GHz 4-element MIMO receiver array supporting analog/RF arbitrary spatial filtering," in *IEEE International Solid-State Circuits Conference*, pp. 410–411, February 2017.

[40] "The LTE standard," tech. rep., Qualcomm Signals Research Group, https://www.qualcomm.com/ media/documents/files/the-lte-standard.pdf, April 2014.

[41] "Introduction to LTE device testing." http://download.ni.com/evaluation/rf/ Introduction_to_LTE_Device_Testing.pdf, National Instruments.

[42] Z. Feng, "Introduction of wi-fi standardization and interoperability certification test." Tutorial slides, October 2017.

[43] "Introduction to wireless LAN measurements." http://download.ni.com/evaluation/rf/ Introduction_to_WLAN_Testing.pdf, National Instruments.

[44] A. Puglielli *et al.*, "Design of energy- and cost-efficient massive MIMO arrays," *Proceedings of the IEEE*, vol. 104, pp. 586–606, March 2016.

[45] J. Wu *et al.*, "A 4GS/s 13b pipelined ADC with capacitor and amplifier sharing in 16nm CMOS," in *IEEE International Solid-State Circuits Conference*, pp. 466–467, February 2016.

[46] B. Vaz *et al.*, "A 13b 4GS/s digitally assisted dynamic 3-stage asynchronous pipelined-SAR ADC," in *IEEE International Solid-State Circuits Conference*, pp. 276–277, February 2017.

[47] A. Ali *et al.*, "A 14-bit 2.5GS/s and 5GS/s RF sampling ADC with background calibration and dither," in *IEEE Symposium on VLSI Circuits*, pp. 206–207, June 2016.

[48] M. Englund *et al.*, "A programmable 0.7-to-2.7GHz direct $\Delta\Sigma$ receiver in 40nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 50, pp. 644–655, March 2015.

[49] B. Xia, A. Valdes-Garcia, and E. Sanchez-Sinencio, "A 10-bit 44-MS/s 20-mW configurable time-interleaved pipeline ADC for a dual-mode 802.11b/bluetooth receiver," *IEEE Journal of Solid-State Circuits*, vol. 41, pp. 530–539, March 2006.

[50] M. Taherzadeh-Sani and A. Hamoui, "A reconfigurable and power-scalable 10-12 bit 0.4-44 MS/s pipelined ADC with 0.35-0.5 pJ/step in 1.2 V 90 nm digital CMOS," *IEEE Transactions on Circuits and Systems — I: Regular Papers*, vol. 60, pp. 74–82, January 2013.

[51] I. Ahmed and D. Johns, "A 50-MS/s (35 mW) to 1-kS/s (15 uW) power scaleable 10-bit pipelined ADC using rapid power-on opamps and minimal bias current variation," *IEEE Journal of Solid-State Circuits*, vol. 40, pp. 2446–2455, December 2005.

[52] I. Ahmed and D. Johns, "A high bandwidth power scalable sub-sampling 10-bit pipelined ADC with embedded sample and hold," *IEEE Journal of Solid-State Circuits*, vol. 43, pp. 1638–1647, July 2008.

[53] O. Nys and E. Dijkstra, "On configurable oversampled A/D converters," *IEEE Journal of Solid-State Circuits*, vol. 28, pp. 736–742, July 1993.

[54] J. E. Park, D. E. Lim, and D. K. Jeong, "A reconfigurable 40-to-67 dB SNR, 50-to-6400 Hz frame-rate, column-parallel readout IC for capacitive touch-screen panels," *IEEE Journal of Solid-State Circuits*, vol. 49, pp. 2305–2318, October 2014.

[55] H. Shibata *et al.*, "A DC-to-1 GHz tunable RF delta-sigma ADC achieving DR=74 dB and BW=150 MHz at $f_0 = 450$ MHz using 550 mW," *IEEE Journal of Solid-State Circuits*, vol. 47, pp. 2888–2897, December 2012.

[56] R. Ritter, J. Kauffman, J. Becker, and M. Ortmanns, "A 10 MHz bandwidth, 70 dB SNDR continuous time delta-sigma modulator with digitally improved reconfigurable blocker rejection," *IEEE Journal of Solid-State Circuits*, vol. 51, pp. 660–670, March 2017.

[57] Y. Ke, J. Craninckx, and G. Gielen, "A design approach for power-optimized fully reconfigurable delta-sigma A/D converter for 4g radios," *IEEE Transactions on Circuits and Systems — II: Express Briefs*, vol. 55, pp. 229–233, March 2008.

[58] A. Rusu, D. Rodriguez de Llera Gonzalez, and M. Ismail, "Reconfigurable ADCs enable smart radios for 4G wireless connectivity," *IEEE Circuits and Devices Magazine*, pp. 6–11, May/June 2006.

[59] G. Taylor and I. Galton, "A refonfigurable mostly-digital delta-sigma ADC with a worst-case FOM of 160 dB," *IEEE Journal of Solid-State Circuits*, vol. 48, pp. 983–995, April 2013.

[60] M. Yip and A. Chandrakasan, "A resolution-reconfigurable 5-to-10-bit 0.4-to-1 V power scalable SAR ADC for sensor applications," *IEEE Journal of Solid-State Circuits*, vol. 48, pp. 1453–1464, June 2013.

[61] M. Yip and A. Chandrakasan, "A resolution-reconfigurable 5-to-10-bit 0.4-to-1 V power scalable SAR ADC," in *IEEE International Solid-State Circuits Conference*, pp. 190–191, February 2011.

[62] N. Verma and A. Chandrakasan, "An ultra low energy 12-bit rate-resolution scalable SAR ADC for wireless sensor nodes," *IEEE Journal of Solid-State Circuits*, vol. 42, pp. 1196–1205, June 2007.

[63] P. Harpe, G. Dolmans, K. Philips, and H. de Groot, "A 0.7V 7-to-10b 0-to-2MS/s flexible SAR ADC for low-power wireless sensor nodes," in *European Solid-State Circuits Research Conference*, pp. 373–376, September 2012.

[64] P. Harpe, Y. Zhang, G. Dolmans, K. Philips, and H. De Groot, "A 7-to-10b 0-to-4MS/s flexible SAR ADC with 6.5-to-16fJ/conversion-step," in *IEEE International Solid-State Circuits Conference*, pp. 472–473, February 2012.

[65] H. Nakane *et al.*, "A fully integrated SAR ADC using digital correction technique for triple-mode mobile transceiver," *IEEE Journal of Solid-State Circuits*, vol. 49, pp. 2503–2514, November 2014.

[66] B. Malki, T. Yamamoto, B. Verbruggen, P. Wambacq, and J. Craninckx, "A 70 dB DR 10 b 0-to-80 MS/s current-integrating SAR ADC with adaptive dynamic range," *IEEE Journal of Solid-State Circuits*, vol. 49, pp. 1173–1183, May 2014.

[67] C. Y. Liou and C. C. Hsieh, "A 2.4-to-5.2fJ/conversion-step 10b 0.5-to-4MS/s SAR ADC with charge-average switching DAC in 90nm CMOS," in *IEEE International Solid-State Circuits Conference*, pp. 280–281, February 2013.

[68] K. Yoshioka, A. Shikata, R. Sekimoto, T. Kuroda, and H. Ishikuro, "An 8bit 0.35-0.8V 0.5-30MS/s 2bit/step SAR ADC with wide range threshold configuring comparator," in *European Solid-State Circuits Research Conference*, pp. 381–384, September 2012.

[69] W. Yu, S. Sen, and B. Leung, "Distortion analysis of MOS track-and-hold sampling mixers using time-varying volterra series," *IEEE Transactions on Circuits and Systems — II: Analog and Digital Signal Processing*, vol. 46, pp. 101–113, February 1999.

[70] S. Kuboki, N. Miyakawa, and K. Matsubara, "Nonlinearity analysis of resistor string A/D converters," *IEEE Transactions on Circuits and Systems*, vol. CAS-29, pp. 383–390, June 1982.

[71] C. C. Liu, S. J. Chang, G. Y. Huang, and Y. Z. Lin, "A 10-bit 50-MS/s SAR ADC with a monotonic capacitor switching procedure," *IEEE Journal of Solid-State Circuits*, vol. 45, pp. 731–740, April 2010.

[72] Y. K. Chang, C. S. Wang, and C. K. Wang, "A 8-bit 500-ks/s low power SAR ADC for biomedical applications," in *IEEE Asian Solid-State Circuits Conference*, pp. 228–231, November 2007.

[73] Y. Zhu, C. H. Chan, U. F. Chio, S. W. Sin, S. P. U, R. P. Martins, and F. Maloberti, "A 10-bit 100-MS/s reference-free SAR ADC in 90 nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 45, pp. 1111–1121, June 2010.

[74] A. Sanyal and N. Sun, "An energy-efficient low frequency-dependence switching technique for SAR ADCs.," *IEEE Transactions on Circuits and Systems — II: Express Briefs*, vol. 61, pp. 294–298, May 2014.

[75] P. Nuzzo, F. De Bernardinis, P. Terreni, and G. Van der Plas, "Noise analysis of regenerative comparators for reconfigurable ADC architectures," *IEEE Transactions on Circuits and Systems — I*, vol. 55, pp. 1441–1454, July 2008.

[76] P. Figueiredo and J. Vital, "Kickback noise reduction techniques for CMOS latched comparators," *IEEE Transactions on Circuits and Systems — II: Express Briefs*, vol. 53, pp. 541–545, July 2006.

[77] J. E. Eklund and C. Svensson, "Influence of metastability errors on SNR in successive-approximation A/D converters," *Analog Integrated Circuits and Signal Processing*, vol. 26, pp. 183–190, March 2001.

[78] A. Waters, J. Muhlestein, and U. Moon, "Analysis of metastability errors in conventional, LSB-first, and asyncrhonous SAR ADCs," *IEEE Transactions on Circuits and Systems — I: Regular Papers*, vol. 63, pp. 1898–1909, November 2016.

[79] B. Verbruggen, M. Iriguchi, and J. Craninckx, "A 1.7 mW 11b 250 MS/s 2-times interleaved fully dynamic pipelined SAR ADC in 40 nm digital CMOS," *IEEE Journal of Solid-State Circuits*, vol. 47, pp. 2880–2887, December 2012.

[80] M. Miyahara, Y. Asada, D. Paik, and A. Matsuzawa, "A low-noise self-calibrating dynamic comparator for high-speed ADCs," pp. 269–272, November 2008.

[81] B. Murmann, "ADC performance survey 1997-2018," May 2018.

[82] R. Winoto and B. Nikolić, "A highly reconfigurable 400-1700MHz receiver using a down-converting Sigma-Delta A/D with 59-dB SNR and 57-dB SFDR over 4-MHz bandwidth," in *IEEE Symposium on VLSI Circuits*, pp. 142–143, June 2009.

[83] C. Wu, E. Alon, and B. Nikolić, "A wideband 400 MHz-to-4 GHz direct RF-to-digital multimode $\delta\sigma$ receiver," *IEEE Journal of Solid-State Circuits*, vol. 49, pp. 1639–1652, July 2014.

[84] C. Wu and B. Nikolić, "A 0.4 GHz - 4 GHz direct RF-to-digital $\Sigma\Delta$ multi-mode receiver," in *IEEE European Solid-State Circuits Research Conference*, pp. 275–278, September 2013.

[85] M. Englund *et al.*, "A programmable 0.7-to-2.7GHz direct $\Delta\Sigma$ receiver in 40nm CMOS," in *IEEE International Solid-State Circuits Conference*, pp. 470–471, February 2014.

[86] F. Opteynde, "A maximally-digital radio receiver front-end," in *IEEE International Solid-State Circuits Conference*, pp. 450–451, February 2010.

[87] G. Taylor and I. Galton, "A mostly-digital variable-rate continuous-time delta-sigma modulator ADC," *IEEE Journal of Solid-State Circuits*, vol. 45, pp. 2634–2646, December 2010.

[88] F.-C. Huang, S.-C. Hsu, Y.-L. Tsai, Y.-Y. Lin, and T.-H. Lin, "LMS-based digital background linearization technique for VCO-based Delta-Sigma ADC," in *IEEE Midwest Symposium on Circuits and Systems*, pp. 753–756, 2014.

[89] J. Kim, W. Yu, H.-K. Yu, and S. Cho, "A digital-intensive receiver front-end using VCO-based ADC with an embedded 2nd-order anti-aliasing sinc filter in 90nm CMOS," in *IEEE International Solid-State Circuits Conference*, pp. 176–177, February 2011.

[90] S. Rao, A. Elshazly, W. Yin, N. Sasidhar, and P. Hanumolo, "A 71dB SFDR open loop VCO-based ADC using 2-level PWM modulation," in *IEEE Symposium on VLSI Circuits*, pp. 270–271, June 2011.

[91] M. Straayer and M. Perrott, "A 12-bit, 10-MHz bandwidth, continuous-time $\Sigma\Delta$ ADC with a 5-bit, 950-MS/s VCO-based quantizer," *IEEE Journal of Solid-State Circuits*, vol. 43, pp. 805–814, April 2008.

[92] B. Cook, A. Berny, A. Molnar, S. Lanzisera, and K. Pister, "Low-power 2.4-GHz transceiver with passive RX front-end and 400-mV supply," *IEEE Journal of Solid-State Circuits*, vol. 41, pp. 2757–2766, December 2006.

[93] C. Andrews and A. Molnar, "Implications of passive mixer transparency for impedance matching and noise figure in passive mixer-first receivers," *IEEE Transactions on Circuits and Systems — I: Regular Papers*, vol. 57, pp. 3092–3103, December 2010.

[94] C. Wu, Y. Wang, B. Nikolić, and C. Hull, "An interference-resilient wideband mixer-first receiver with LO leakage suppression and I/Q correlated orthogonal calibration," *IEEE Transactions on Microwave Theory and Techniques*, vol. 64, pp. 1088–1101, April 2016.

[95] D. Murphy, A. Mirzaei, H. Darabi, M.-C. Chang, and A. Abidi, "An LTV analysis of the frequency-translational noise-cancelling receiver," *IEEE Transactions on Circuits and Systems — I: Regular Papers*, vol. 61, pp. 266–279, January 2014.

[96] A. Sanyal and N. Sun, "A 18.5 fJ/step VCO-based 0-1 MASH $\delta\sigma$ ADC with digital background calibration," in *IEEE Symposium on VLSI Circuits*, pp. 26–27, June 2016.

[97] R. Winoto and B. Nikolić, *Multi-Mode/Multi-Band RF Transceivers for Wireless Communications*, ch. Discrete-Time Processing of RF Signals. Wiley and Sons, October 2010.

[98] H. Garvik, C. Wulff, and T. Ytterdal, "An 11.0 bit ENOB, 9.8 fJ/conv.-step noise-shaping SAR ADC calibrated by least squares estimation," in *IEEE Custom Integrated Circuits Conference*, pp. 1–4, 2017.

[99] J. Kim, T.-K. Jang, Y.-G. Yoon, and S. Cho, "Analysis and design of voltage-controlled oscillator based analog-to-digital converter," *IEEE Transactions on Circuits and Systems — I: Regular Papers*, vol. 57, pp. 18–30, January 2010.

[100] A. Abidi, "Phase noise and jitter in CMOS ring oscillators," *IEEE Journal of Solid-State Circuits*, vol. 41, pp. 1803–1816, August 2006.

[101] B. Kim, S. Kundu, S. Ko, and C. Kim, "A VCO-based ADC employing a multi-phase noise-shaping beat frequency quantizer for direct sampling of sub-1mV input signals," in *IEEE Custom Integrated Circuits Conference*, no. 1-4, September 2014.

[102] E. Chang *et al.*, "BAG2: A process-portable framework for generator-based AMS circuit design," in *IEEE Custom Integrated Circuits Conference*, April 2018.