Analog-to-Digital Converters For Secure and Emerging AIoT Applications

by

Ruicong Chen

B.S., Peking University (2019)
S.M., Massachusetts Institute of Technology (2021)

SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

IN ELECTRICAL ENGINEERING AND COMPUTER SCIENCE
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2023

Authored by: Ruicong Chen
Department of Electrical Engineering and Computer Science
May 16, 2023

Certified by: Anantha P. Chandrakasan
Vannevar Bush Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Certified by: Hae-Seung Lee
ATSP Professor of Electrical Engineering
Thesis Supervisor

Accepted by: Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

Analog-to-Digital Converters
For Secure and Emerging AIoT Applications

by

Ruicong Chen

Submitted to the Department of Electrical Engineering and Computer Science
on May 15, 2023, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

AI algorithms based on convolutional neural networks (CNNs), coupled with their high computational requirements, have stimulated the development of novel energy-efficient hardware. Analog neural networks (ANNs) with in-memory computing (IMC) using resistive random-access memory (RRAM) are promising architectures to reduce latency and increase energy efficiency for IoT devices. However, interface circuitry, including analog-to-digital converters (ADCs) between RRAM and digital components, is becoming the bottleneck of the RRAM-based ANNs. To address this challenge, a direct hybrid encoding for signed expressions (HESE) SAR is proposed to increase the sparsity of ADC output.

In addition to the performance requirements, the security of IoT devices is of paramount importance. An attacker can perform an ADC power side-channel attack (PSA) to expose confidential information by tapping into the power supply of the ADC. This attack exploits the strong correlation between the ADC digital output codes and the ADC power supply using neural networks based PSA. Previous works have implemented current equalizers or noise injections to protect ADCs from PSAs. However, the current equalizer introduces a large area and energy overhead for the ADC, which is not ideal for IoT applications. Additionally, the previous work with noise injection only protects the probing of CDAC supply. To overcome these limitations, two secure ADCs are proposed to improve both energy efficiency and security, making them more suitable for real-world applications.

Thesis Supervisor: Anantha P. Chandrakasan
Title: Vannevar Bush Professor of Electrical Engineering and Computer Science

Thesis Supervisor: Hae-Seung Lee
Title: ATSP Professor of Electrical Engineering

Acknowledgments

There are no words to adequately express my gratitude for the support and help I received from everyone around me throughout my graduate studies at MIT. The Covid-19 pandemic has profoundly impacted every aspect of our lives, but despite these challenges, I have been fortunate to have an incredibly fulfilling experience as a graduate student at MIT. I have had the opportunity to work on cutting-edge research alongside some of the most intelligent and talented individuals in the field. It is with great appreciation and admiration that I mention those who have supported and guided me during this journey.

First and foremost, I would like to extend my sincerest gratitude to my two advisors, Professor Anantha Chandrakasan, and Professor Hae-Seung Lee. They have consistently provided invaluable guidance and advice, not only in terms of research and academic matters but also in navigating the complexities of life at MIT. They have shown me how to think critically, develop a deep understanding of my field, and conduct myself with the professionalism and integrity that is expected of a mature researcher and scientist. Furthermore, they have demonstrated the qualities of a good mentor, both in how they have guided me and in how they have encouraged me to be a supportive mentor to others. They have offered thoughtful suggestions regarding my career development and have always been eager to share their extensive knowledge and experience. I am also grateful to Professor Song Han, who served on my thesis committee and provided invaluable expertise in the area of neural networks.

Additionally, I am deeply thankful to all my colleagues in the HSLee Group and Anantha Group for their unwavering support and camaraderie. Our discussions about technical challenges and personal matters have been both enlightening and encouraging. It has been a pleasure to share this journey with them, and I will always cherish the memories we have created together. In particular, I would like to acknowledge Mohamed Radwan Abdelhamid, who has been a fantastic mentor, guiding me through my early years at MIT and assisting me with chip design issues. I have also enjoyed engaging in conversations about campus life with Preetinder Garcha and

Sirma Orguc. Rishabh Mittal and Maitreyi Ashok have been incredibly helpful with analog design, while Miaorong Wang has generously shared her knowledge of digital synthesis.

Finally, I must express my heartfelt gratitude to my family and friends for their unconditional love and unwavering support. Their presence has been a constant source of strength during the pandemic, and I cannot imagine what life would have been like without them. Their encouragement has lifted me up during difficult times and has motivated me to persevere in my studies. ChatGPT also helps with the grammar of the thesis.

I feel incredibly fortunate to have been surrounded by such a supportive network of mentors, colleagues, friends, and family during my time at MIT. Each one of them has played a vital role in my journey, and I am immensely grateful for the opportunity to learn from and grow with them.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

1.1.1 AIoT Applications

Following Moore's Law [3], contemporary microprocessors accommodate billions of transistors within just a few square millimeters of silicon area, offering remarkable computing power, high processing performance, and accommodating a diverse array of functionalities. This advancement has facilitated numerous applications, such as portable electronics (e.g., laptops, smartphones), wearable health-monitoring systems, and wireless sensor nodes. An increasing number of these devices possess computing capabilities and are progressively interconnected, forming the "Internet of Things" (IoT). The remarkable surge in computing power, alongside advancements in algorithms and access to vast amounts of data, has also empowered modern computing systems to execute various "Artificial Intelligence" (AI) tasks, including image recognition [4], speech recognition [5], and natural language understanding [6]. IoT devices with AI processing are called AIoT applications.

One particularly intriguing aspect of these recognition systems is the "always-ON" sensing capability, enabling continuous data collection and environmental monitoring. Upon detecting an event of interest (e.g., a human face) through an "always-ON" component (e.g., face detection), more complex systems (e.g., a face recognition system typically in "sleep" mode to conserve energy) can be activated for further data processing. However, the growing number of "smart" and "always-ON" systems often face limited energy budgets, as they typically rely on batteries or ambient energy harvesting. Consequently, energy efficiency is a paramount concern for the sustainable implementation of these computing systems in real-world applications.

This thesis proposes a direct hybrid encoding for signed expressions (HESE) SAR to increase the sparsity of ADC output, which will enable improved energy efficiency of RRAM-based ANNs for AIoT applications.

1.1.2 ADC Attack Scenario

Analog-to-digital converters (ADCs) serve as essential components within electronic systems, responsible for transforming analog signals into their digital counterparts. With the increasing prominence of digital signal processing, there has been a concerted effort by researchers to enhance ADC performance, optimize power consumption, and minimize the physical area. Despite these advancements, recent studies [7][8] have identified a significant hardware security vulnerability arising from ADC side-channel attacks. This vulnerability is situated at the interface between analog circuits and digital processors, presenting a novel research direction for ADC studies. This thesis also aims to investigate and address the challenges posed by ADC side-channel attacks, contributing to the development of more secure and robust electronic devices.

In the field of cryptographic hardware research, side-channel attacks refer to those that capitalize on security vulnerabilities stemming from the physical implementation of a cryptographic algorithm, rather than any inherent weaknesses in the algorithm itself. For instance, the power side-channel attack (PSA) or electromagnetic side-channel attack (EMSA) occurs when an attacker exploits the power supply current waveforms or electromagnetic signals associated with an encryption engine to extract the secret key of a given cryptographic algorithm. Various classes of side-channel attacks exist, contingent upon the type of side-channel leakage exploited by the attacker.

Sensing hardware like sensors and ADCs serves as a crucial link between real-

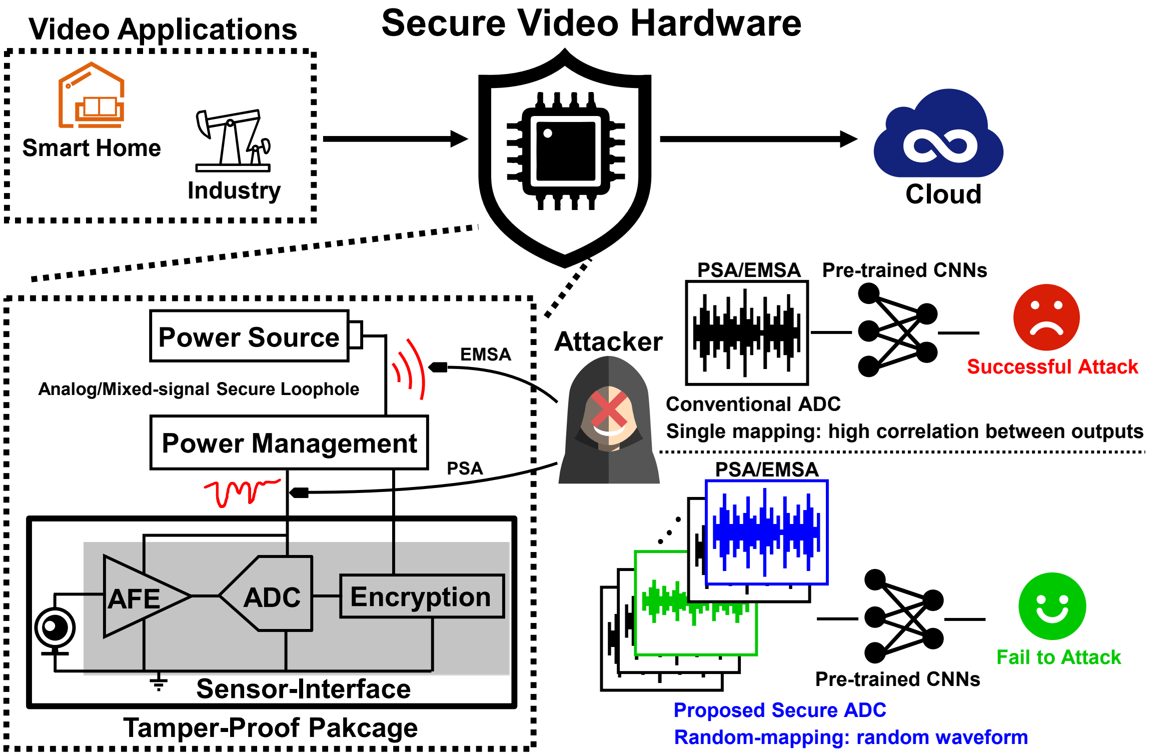

Figure 1-1: Potential security loopholes in video hardware

world physical phenomena and electronic systems, making it prevalent across various application domains. As depicted in Figure 1-1, typical sensing hardware comprises a sensor and a sensor interface circuit. The sensor transforms a physical quantity (e.g., temperature, pressure, light, or position) into an electrical signal, which is subsequently processed by the sensor interface circuit. An analog frontend (AFE) circuit interfaces with the sensor to condition the analog signal it produces. Following this, an analog-to-digital converter (ADC) digitizes the conditioned analog signal. Prior to providing the digitized sensor output value to the intended user, the sensor interface circuit may perform digital signal processing (DSP) and encryption for data reduction and security enhancement, respectively.

Frequently, sensing hardware generates sensitive sensor data that should be exclusively accessible to authorized users like implantable devices. For instance, healthcare sensing hardware monitoring ECG signals may inadvertently expose confidential personal health information to potential adversaries. Similarly, data gathered by temperature sensors and utility meters within a smart home can reveal an individual's lifestyle patterns [9].

The security of ADCs is also a significant concern. An attacker can perform an ADC power side-channel attack (PSA) to steal confidential and sensitive signals by tapping into the power supply of the ADC. Such attacks can exploit the strong correlation between the ADC digital output codes and the ADC power supply with neural networks based PSA.

Potential security threats within sensing hardware are also shown in Figure 1-1. First, an attacker may attempt to intercept wireline or wireless communication between the sensing hardware and the user. This risk can be mitigated by encrypting sensor data [10] and adhering to established security protocols [11]. Second, an attacker may try to breach secure communication by executing a power side-channel attack (PSA) on the encryption engine. If successful in extracting the secret key of the encryption algorithm (e.g., AES) [12][13][14], the attacker could decode the encrypted data to access the original sensor information.

Third, an attacker might attempt to measure the analog sensor output signal directly. Given the typically high output impedance of most sensors, an additional sensor interface circuit introduced by an attacker is likely to disrupt the sensitive analog signal chain within the sensing hardware. This direct measurement can be thwarted by enclosing both the sensor and the sensor interface circuit within a tamper-proof package [15]. Lastly, if prior security vulnerabilities are addressed, an attacker may still perform a PSA or EMSA on analog/mixed-signal circuits, particularly ADCs. As ADC operations heavily rely on the sampled analog input voltage (or the equivalent digital output bits), the power supply current waveforms or EM leakage of the ADC may correlate with private sensor data. An attacker could exploit this correlation to reconstruct the digitized private sensor signal from the ADC power supply current waveforms.

Unlike the direct measurement of the analog sensor output signal, the PSA or EMSA of analog/mixed-signal circuits does not disturb the sensitive analog signal chain, as analog circuits are typically designed to withstand minor perturbations in their power supply voltage. Unfortunately, due to constraints like battery replacement requirements or physical size limitations of sensing hardware, tamper-proof packaging may not always extend to encompass the power source and power management circuit. In such cases, tamper-proof packaging is not a viable countermeasure against analog/mixed-signal domain PSAs and EMSAs [7][8][16].

The power or EM side-channel leakage of an ADC poses a significant security threat, as it exposes the private signal chain data before encryption. Considering the numerous applications of ADCs, the ADC PSA or EMSA scenario in sensing hardware can be extrapolated to other types of hardware that process sensitive information (e.g., communication systems).

1.2 ADC Implementation

Analog-to-digital converters (ADCs) are critical components in various systems and devices used in the world today. They enable the conversion of analog signals into digital signals that can be processed by digital circuits, allowing for the processing, storage, and transmission of data. ADCs are essential in numerous applications, including communications, healthcare, transportation, aerospace, and defense. For example, in medical devices, ADCs play a crucial role in the accurate measurement and monitoring of vital signs. In the transportation sector, ADCs are used in sensors to monitor and control various parameters, such as speed, acceleration, and temperature. ADCs have significant importance in enabling the digital transformation of various industries, which can have a positive impact on the world's economy and quality of life. Researchers push the performance of ADCs with various techniques [17].

1.2.1 Direct HESE SAR ADC: The first sparsity-aware ADC for analog neural networks

The proposed work introduces a novel bit-level sparsity-aware successive approximation register (SAR) ADC that directly produces Hybrid Encoding for Signed Expressions (HESE). This 12-bit resolution ADC is designed to support large artificial neural networks (ANNs) with good accuracy. The proposed HESE ADC incorporates two thresholds for 2-bits look-ahead (LA), and noise averaging (NA) is performed in the last two bit cycles. The proposed HESE SAR achieves an impressive figure of merit (FoM) of 15.2 fJ/conv.-step at 45MS/s, with a core area of only 0.072mm^2.

1.2.2 RaM-SAR: The first secure ADC for high-speed applications

Designing secure analog-to-digital converters (ADCs) is a significant challenge in integrated circuit design, especially for low-power, high-speed, and small form-factor systems used in the Internet of Things (IoT) and wearable devices. Such systems are vulnerable to side-channel attacks that can extract sensitive information from their power consumption or electromagnetic emissions. To address this challenge, we propose RaM-SAR, a secure random-mapping SAR ADC that provides resistance against both power and electromagnetic side-channel attacks. This 12-bit, 25 MS/s ADC achieves an energy consumption of 11.3 fJ per conversion step by using a novel random-mapping technique that randomly maps each conversion to one of 4096 conversion sequences. This randomization helps to protect against neural network-based power and electromagnetic side-channel attacks.

1.2.3 Sniff-SAR: The first detection-driven and un-trainable secure ADC

Because of the need for more secure ADCs capable of detecting and protecting against power and EM side-channel attacks, Sniff-SAR is developed. This 9.8fJ/c.-s 12b secure ADC incorporates detection-driven protection against side-channel attacks through the use of EMSA and PSA detectors. While the ADC core performs the analog-to-digital conversion, it normally operates in the unprotected SAR mode which is faster and more energy efficient. Periodically, the EMSA and PSA detectors check for side-channel attacks, and if detected, the ADC activates the secure SAR mode against both EMSA and PSA. The secure mode provides 3.6x10^{16} different switching traces, making it impractical to train the ADC with neural networks for PSA or EMSA. These innovative features make Sniff-SAR an effective solution for secure ADCs and have the potential to greatly enhance the security of a wide range of electronic systems.

1.3 Thesis Organization

Background knowledge is introduced in chapter 2. Analog neural networks and sparsity encoding are discussed. Common side-channel attacks are introduced and the basic of successive-and-approximation registers (SAR) ADCs is mentioned. In chapter 3, detailed implementation and measurement results are shown for the direct HESE SAR ADC [18]. In chapter 4, detailed implementation and measurement results are shown for the RaM-SAR [7]. In chapter 5, detailed implementation and measurement results are shown for the Sniff-SAR. Conclusions and future work are discussed in chapter 6.

Chapter 2

Background

2.1 Analog Neural Networks

2.1.1 Memory-wall

The term "memory-wall" [19] refers to the growing disparity between CPU clock speeds and memory access times. Although CMOS scaling has led to smaller and faster transistors, which boost CPU speeds, overall processing times remain constrained by slow memory access times. To sustain Moore's Law, multi-core processor designs began gaining traction around 2005. However, the parallel operation of multiple cores made on-chip memory bandwidth and energy consumption increasingly dominant concerns.

A primary factor contributing to memory being a bottleneck in modern computing systems is the traditional "von-Neumann" architecture, which physically separates memory and processor, with data flowing between them. This configuration results in limited data transfer bandwidth, dictated by the memory input/output (IO) capacity. Consequently, in contemporary computing systems, a significant portion of time and energy is expended moving data between memory and processing elements [20].

2.1.2 Analog Neural Networks

Analog neural networks (ANNs) [21] are a type of artificial neural network that uses analog electronic circuits to perform computations. They mimic the behavior of biological neurons by using continuous voltage signals instead of digital signals. ANNs are considered to be highly energy-efficient and can process signals in parallel, making them well-suited for tasks such as image and speech recognition. In-memory computing (IMC) using resistive random-access memory (RRAM) is a promising architecture for ANNs as it reduces latency and increases energy efficiency for IoT devices. However, the interface circuitry between RRAM and digital components, including analog-to-digital converters (ADCs), is becoming a bottleneck for the development of ANNs. Additionally, the security of ADCs is a significant concern as they are vulnerable to power and electromagnetic side-channel attacks. Therefore, there is a growing interest in developing secure and energy-efficient ADCs for ANNs that can enable the development of more efficient and secure AI algorithms for IoT devices.

With the increasing popularity of generative AI applications like chatGPT [22], the amount of storage needed is increasing tremendously for Large Language Models (LLMs). Hardware implementations of LLMs face major challenges in dealing with the huge amount of data movement. The memory access and data movement energies are the dominant ones, compared to the computation energy [23].

Various methods are possible for tackling the memory bottleneck in modern computing systems. Dynamic Voltage Scaling (DVS) is an effective way to reduce the energy consumption of the memory subsystem [24]. However, memory like SRAM does not scale easily as in logic circuits, due to the reduction of operating margins.

One promising approach is the concept of analog neural networks, which blurs the line between conventional memories and compute elements, by computing inside the memory.

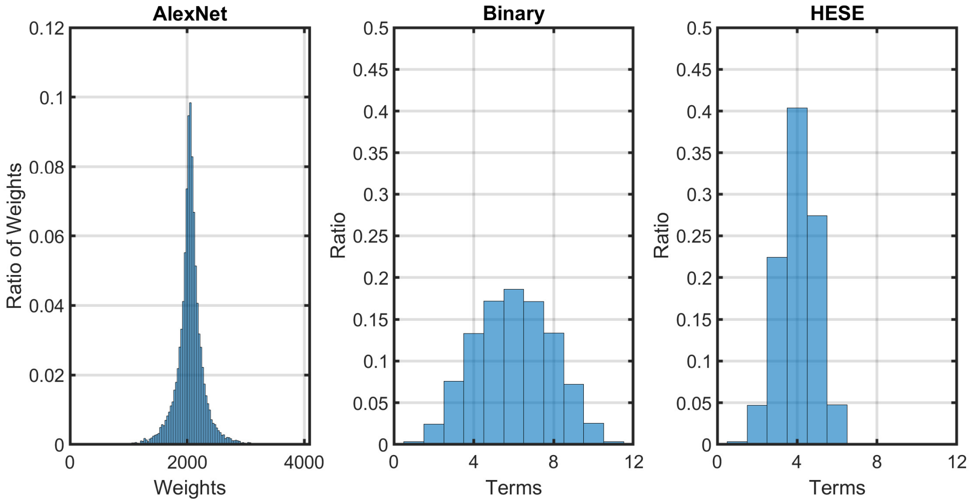

Figure 2-1: The distributions of weights in AlexNet shape the distribution of the number of terms in a binary encoding and the efficient HESE encoding.

2.1.3 Sparsity Encoding

Sparsity encoding is a powerful technique in machine learning that aims to enhance the efficiency and performance of algorithms by capitalizing on the inherent sparsity found in various machine learning models. In the context of machine learning, sparsity refers to the notion that not all features within a dataset hold equal significance in predicting the outcome of a specific task [25]. As a result, by identifying and either removing or diminishing the impact of less critical features, algorithms can be streamlined and made more accurate. To accomplish this, sparsity encoding techniques, such as Hybrid Encoding for Signed Expressions (HESE) and term quantization (TQ) [1], have been developed and refined over time.

These techniques facilitate the effective encoding of sparse matrices, which frequently emerge in machine learning tasks. By minimizing the number of non-zero elements in these matrices, sparsity encoding can dramatically decrease the computational complexity of machine learning algorithms, leading to faster processing times and more efficient use of resources. This is particularly beneficial in applications with constrained computational capabilities, such as those found in mobile devices, Internet of Things (IoT) systems, and other embedded systems.

To implement sparsity encoding effectively, several techniques have been proposed.

Hybrid Encoding for Signed Expressions (HESE) is one such method that aims to compress sparse matrices by identifying patterns in the data and encoding them in a more compact form [1]. This can lead to significant reductions in memory usage and computational overhead, making it particularly well-suited for applications with limited resources.

Another sparsity encoding technique, term quantization (TQ), focuses on quantizing the terms in a group of weights rather than truncating all the lower bits of the weights [1]. By quantizing less important terms, TQ can minimize its influence on the model's predictions, allowing it to concentrate on the most significant aspects of the data. This results in a more efficient and accurate model that is better equipped to handle real-world problems.

As the demand for machine learning applications in various domains continues to grow, the need for efficient and effective algorithms becomes increasingly important. Sparsity encoding techniques, such as HESE and TQ, play a crucial role in addressing this need by enabling algorithms to focus on the most relevant features and discard the rest. By reducing computational complexity and memory requirements, these techniques pave the way for more advanced machine learning applications in resource-constrained environments, unlocking new possibilities for mobile devices, IoT systems, and other embedded platforms.

2.2 Side-channel Attacks

2.2.1 Digital Hardware PSA

Following the initial demonstration of power side-channel attacks (PSAs) on digital cryptographic hardware [12], researchers have investigated both attack and countermeasure strategies for the digital encryption engine PSAs. Given the similarities between cryptographic hardware PSAs and ADC PSAs, ADC PSA and EMSA researchers can benefit from the concepts and techniques previously developed in cryptographic hardware PSA research.

Cryptographic hardware power side-channel attacks (PSAs) can be broadly categorized into two classes based on the assumptions made to construct a power analysis model that correlates the secret key with the power traces: profiled attacks and non-profiled attacks.

A profiled attack is based on the assumption that an attacker has the opportunity to experiment with a training device before gaining access to the target device [26]. The training device is a piece of digital cryptographic hardware that shares the same part number and algorithm implementation as the target device. Profiled attacks comprise two phases: profiling and attacking.

During the profiling phase, the attacker has full control over the training device and collects numerous power traces for all conditions intended to be modeled. Subsequently, the attacker constructs the power analysis model using the gathered power traces. In the attacking phase, the attacker measures the power traces of the target device and performs a Maximum Likelihood Estimation (MLE) with the prepared power analysis model. This approach enables the attacker to uncover the secret key of the target cryptographic hardware. Template attacks and stochastic attacks, based on multivariate Gaussian distribution, have been introduced as profiled attacks [26]. Recently, neural networks have emerged as a new tool for implementing power analysis models in profiled attacks [27].

A non-profiled attack assumes an attacker performs a PSA directly on the target cryptographic hardware without utilizing a training device. Since the attacker has no prior knowledge of the target cryptographic hardware, the power analysis model for a non-profiled attack is based on a hypothesis that connects the secret key to the hardware power traces. For all potential secret key guesses, the attacker evaluates the relationship between the key guess and the measured power traces of the target device based on the hypothesis.

The attacker quantifies this evaluation using a distinguisher and identifies the secret key with the highest distinguisher output value. Non-profiled attacks typically require a significantly larger number of power traces than profiled attacks to correctly identify the secret key using the distinguisher. In line with profiled attack research, a recent non-profiled attack study utilized neural networks to construct its distinguisher [14].

Countermeasures have also been studied for the digital hardware PSA: algorithm-level countermeasure and hardware-level countermeasure.

Algorithm-level countermeasures are implemented in software and involve modifying the internal computations of an encryption algorithm in a manner that reduces the correlation between the secret key and the hardware power traces without compromising the security of the encryption algorithm. Masking is a widely used algorithmic-level countermeasure that accomplishes this by concealing the intermediate data of an encryption algorithm using random values known as masks [28].

For hardware-level countermeasures, circuit implementation is modified to increase the PSA-resistance of digital hardware. Equalization and randomization are two methods. For equalization, [29] uses differential logic gates to equalize the power traces and [30] equalizes the power traces with an on-chip current equalizer. For randomization, [31] introduces random dithering in the control loop of a buck regulator.

2.2.2 ADC PSA and EMSA

ADC PSA was first introduced in [15]. For slope-based ADCs, attackers could steal the private conversion output by monitoring the duration of the A/D conversion. Slope-based ADCs, however, are less common than SAR ADCs in many high-speed low-power applications.

Noise injection technique was reported in [32]. It injects the same dither to the CDAC as well as the comparator to randomize the conversion sequence. While the technique showed promise against correlation based template-matching attacks, the resiliency against more sophisticated neural net based attacks has not been demonstrated.

Switched capacitor equalizer protection was reported in [16]. The current equalizer obscures the power traces by supplying the power from an on-chip capacitor. The power supply traces are nearly identical regardless of the ADC outputs. The switched capacitor current equalizer has 3 phases. They are charge, supply, and purge phases.

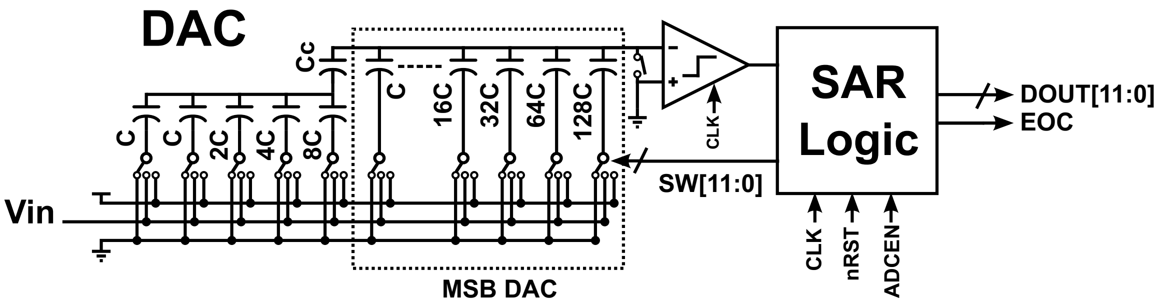

Figure 2-2: A single-ended 12 bits SAR ADC with 8-4 segmented capacitor array

In the charge phase, the supply capacitor is charged to VDD. In the supply phase, the supply capacitor supply power to the ADC. The current supplied by the supply capacitor depends on the DAC capacitor switching, but it is isolated from the power supply pin of the chip. Thus, the ADC is protected from PSA. In the purge phase, the supply capacitor is discharged to a fixed voltage before being recharged.

Random switching scheme was reported in [8]. The first few MSB capacitors of the CDAC are split into unit capacitors. These capacitors are switched at random times. The decision timing is also random. This makes training for power or EM side-channel attacks difficult.

ADC PSA and EMSA share the same objectives with digital cryptographic hardware PSA. Both attackers need to find the correlation between the side-channel information and the private data. Defense against both attacks aims to decouple the private data and the side-channel leakage.

Digital cryptographic hardware PSA, however, tries to extract the exact key of the cryptographic algorithm. However, every bit should be correct to break the algorithm. It's fine for ADC PSA or EMSA to get time-series data with limited resolution. For example, 8 bits are enough to estimate the heart rate from an ECG signal [33]. The effectiveness of the ADC PSA or EMSA is evaluated with relative error to the real ADC outputs.

2.3 Successive-Approximation Registers (SAR) ADCs

Successive approximation register (SAR) analog-to-digital converters (ADCs) [34] are a type of ADC that use a binary search algorithm to convert analog signals into digital signals. SAR ADCs are widely used in various applications due to their high accuracy and low power consumption. They work by first comparing the input signal to the midpoint of the ADC's range and then using the result to determine which half of the range to further search. This process continues until the resolution is achieved. SAR ADCs are particularly well-suited for low-power and low-speed applications where high resolution is required. They are also used in systems where accuracy and linearity are critical, such as in medical devices and sensors. As the demand for low-power and high-resolution ADCs continues to grow in various industries, there is a need for continued development of SAR ADC technology to meet these demands.

Figure 2-2 shows the basics of single-ended SAR ADC with 8 MSBs and 4 LSBs segmented capacitor array and bottom-plate sampling. The SAR ADC typically consists of the sample-and-hold circuitry, a comparator, a feedback DAC, and the SAR logic. In the common capacitor DAC implementation, the sample-and-hold function is incorporated in the DAC.

The components and the connections of the SAR are described as follows. The bottom plate of the capacitors is connected to one of the three voltages (Vin, GND, and VDD). SW signal controls the switches of the capacitors. The top plate of the capacitor array is connected to the negative input of the comparator. The positive input of the comparator is connected to the ground. In a single-supply system, the positive input of the comparator is connected to a common mode voltage. The comparator output serves as an input for the SAR Logic block. CLK is the global clock that controls the comparator strobe and the SAR Logic. The nRST signal resets the state machine inside the SAR Logic. The ADCEN signal is the enable signal for the SAR. EOC is the end-of-conversion signal for a later phase. The DOUT signal contains 12 bits of converted binary digital code. The DAC array is separated into MSB DAC and LSB DAC to save energy. The LSB DAC is not labeled for simplicity.

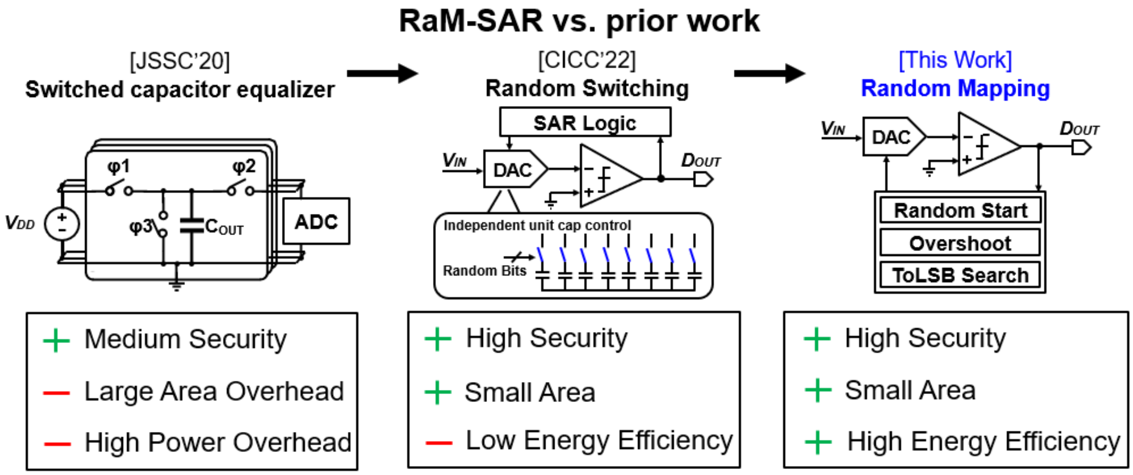

Figure 2-3: Potential security loopholes in video hardware

The bridge capacitor C_c helps to interpolate the LSB in MSB DAC into 4 bits.

Differential operations are popular in the SAR ADC implementation to reduce the sensitive common mode and power supply disturbance. All ADCs in this thesis are implemented in differential mode.

Conventional ADCs have a high correlation between the side-channel information and the digital outputs. The attack can use pre-trained CNNs to decode the information and steal the sensitive information. RaM-SAR uses the random-mapping (RaM) technique that let the ADC has a random waveform during each conversion. This makes it very difficult for the attacker to the PSA or EMSA.

2.4 Previous Work on Secure ADCs

[16] (Figure. 2-3, left) implemented a current equalizer to decorrelate the supply current of all blocks from digital outputs. The idea is to use the current equalizer as the voltage supply so that the power pattern is similar for each conversion. The current equalizer has three switches and a large capacitor. The capacitor is first charged to a fixed value. The capacitor then uses charge to supply the ADC core. In the end, the capacitor is discharged to a fixed value. This scheme provides medium security against EMSA as the attacker can get the ADC power trace using an EM probe. The area and power have room for improvement due to the large capacitor.

The technique is also not secured against EMSA as the attacker can directly probe the power trace of the ADC core.

[8] (Figure. 2-3, middle) proposed to control the unit capacitors for the MSB bits independently and switch them randomly to randomize the power traces. The technique can protect against both PSA and EMSA. The energy efficiency has room for improvement due to extra wiring and control logic.high-speed

Recently, integrated regulators with control loop randomizers [31][35] and shunt linear regulators [36][37] have been proposed for general PSA and EMSA protection with the trade-off of the performance of the core circuit. [31] present the enhanced PSA resistance provided by an on-die all-digital high-frequency integrated inductive voltage regulators (IVRs), fabricated in 130nm CMOS technology, for a standard (unprotected) 128-bit Advanced Encryption Standard (AES) core designed using static CMOS logic. The IVR is equipped with a configurable digital proportional-integral-derivative (PID) controller, a digital discontinuous conduction mode (DCM) controller, and a loop randomization (LR) block.

[35] presents enhanced power and EM SCA resistance for standard (unprotected) 128-bit AES engines with parallel (P-AES, 128-bit) and serial (S-AES, 8-bit) datapaths using an on-die security-aware all-digital series low-dropout (DLDO) regulator, commonly employed for fine-grain SoC power management. The security-aware DLDO bolsters SCA resistance by introducing control-loop induced perturbations in a baseline DLDO, further enhanced by a random switching noise injector (SNI) through power stage control, and a randomized reference voltage (R-VREF) generator combined with all-digital clock modulation (ADCM).

[36] builds upon the concept of signature attenuation in the current domain but introduces a fully-synthesizable design featuring digital current sources, control loop, and bleed. This approach increases the minimum traces to disclosure (MTD) from 10 million to 250 million (a 25x improvement) using a single synthesizable countermeasure. Furthermore, by combining the digital signature attenuation circuit (DSAC) with a second synthesizable generic technique in the form of a time-varying transfer function (TVTF), this work achieves an MTD of over 1.25 billion for both EM and

power SCA, demonstrating a robust and scalable solution for enhancing cryptographic algorithm security.

Considering that the correlated current is the source of both power (at the supply pin) and EM leakage (radiation throughout the current path), [37] adopts current-domain 'signature attenuation' (CDSA) as a low-overhead generic countermeasure against both EM and power side-channel attacks. This approach aims to achieve the highest MTD reported to date, providing a more efficient and secure solution for protecting cryptographic algorithms against side-channel attacks.

Chapter 3

Direct HESE SAR ADC for AIoT Applications

3.1 Introduction

The application space of AIoT is huge, ranging from fundamental research to personal daily life. AIoT has great potential in the future. By 2030, ~350 billion AIoT devices are expected to be in operation, reaching $16 trillion, or 14% of total GDP [38][39]. With the increasing need for edge computing and long battery life, AIoT devices with low standby power and high efficiency for neural network inference are in great demand. The applications include microphones, industry monitoring, vital-sign monitoring devices, etc. End devices with speech interfaces can benefit greatly from ultra-low power AIoT devices, such as Voice Activity Detection (VAD) and Keyword Spotting (KWS) [40] systems. Both VAD and KWS systems have to be always-on and highly efficient in inference. AIoT devices are also ubiquitous in machine health monitoring products that minimize downtime with sensor signals [41]. Moreover, the AIoT system with a reconfigurable rectenna opens up opportunities for wireless and battery-less in-body vital-sign monitoring [42].

Conventional AIoT systems, however, still need to improve their energy efficiency. For battery-constrained AIoT systems, energy-efficient implementations would bring longer battery life and better user experience. For AIoT systems with connected power sources, low-power designs are still preferred as they are more environmentally friendly. Some features in conventional AIoT systems can be explored to lower power consumption and improve the energy-efficiency. For sensing, the input signals usually have low activity. For computing, the data in embedded neural networks are highly sparse. For memory access, data reuse can improve energy efficiency associated with the memory wall in the von Neumann architecture.

AI algorithms based on CNNs, coupled with their high computational requirements, have stimulated the development of novel energy-efficient hardware, such as Eyeriss [43]. Previous work applies term quantization to the weights [44] and uses binary encoding [21][45] in ANNs. Hybrid encoding for signed expressions (HESE) and term quantization (TQ) to the outputs of each layer further reduces the non-zero terms and increase sparsity (Figure 3-1). The HESE signed digit representation (SDR) is directly generated during the analog-to-digital conversion. The HESE SDR has both positive and negative terms to reduce the non-zero terms. The TQ prunes out small terms in a group basis.

In this work, we propose the first bit-level sparsity-aware successive approximation register (SAR) ADC which directly produces HESE. The 12-bit resolution can support large ANNs with good accuracy. The proposed HESE ADC has two thresholds for 2-bits look-ahead (LA) and noise averaging (NA) is performed in the last two bit cycles. The proposed HESE SAR achieves a FoM of 15.2 fJ/conv.-step at 45MS/s. The core area of the SAR ADC is 0.072mm^2.

The main contributions of the work are:

- The first direct HESE SAR ADC to provide sparse encoding during the analog-to-digital conversion with 2-bit look-ahead and the noise averaging at the last couple of cycles.

- The use of direct HESE SAR ADC along with term quantization (TQ) to increase the sparsity in analog neural networks (ANNs).

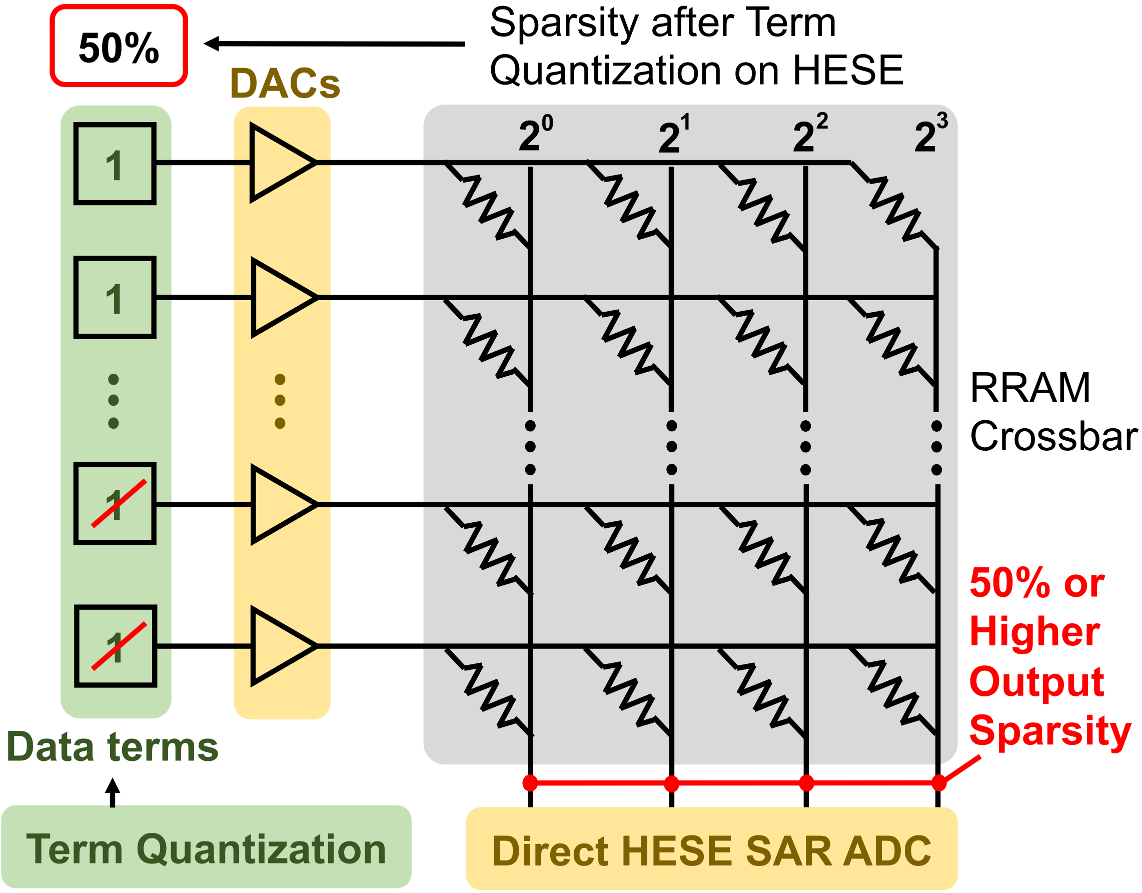

Figure 3-1: An example crossbar of RRAM with 1-bit RRAM cells and 1-bit input values for computing dot products. Both data and weights are bit-sliced, with each weight term occupying a separate RRAM column. The proposed SAR ADC directly provides the hybrid encoding for signed expressions (HESE) signed-digit representation (SDR) to minimize the number of non-zero terms. Also, term quantization [1] sets low-order power-of-two terms to 0, indicated by red slashes, to satisfy a group term budget. The Figure illustrates a case when 50% of input terms are zeros. In forming the dot product of the input and an RRAM column of weight terms, products on the column to be output for accumulation have 50% or higher sparsity, given that some weight terms may be zeros. The direct HESE SAR ADC introduces extra sparsity other than TQ and reduces the energy of computation.

3.2 Specific Background

3.2.1 Analog Neural Networks for CNNs

Analog Neural Networks (ANNs) typically use RRAM crossbars to both store CNN weights and perform matrix multiplication in-memory in an analog fashion [21][45].

An example crossbar of RRAM with 1-bit RRAM cells and 1-bit input values for computing dot products is shown in Figure 3-1. Both data and weights are bit-sliced, with each weight term occupying a separate RRAM column. The proposed SAR ADC directly provides the hybrid encoding for signed expressions (HESE) signed-digit representation (SDR) to minimize the number of non-zero terms. Also, term quantization [1][46][47] sets low order power of two terms to 0 with red slashes to

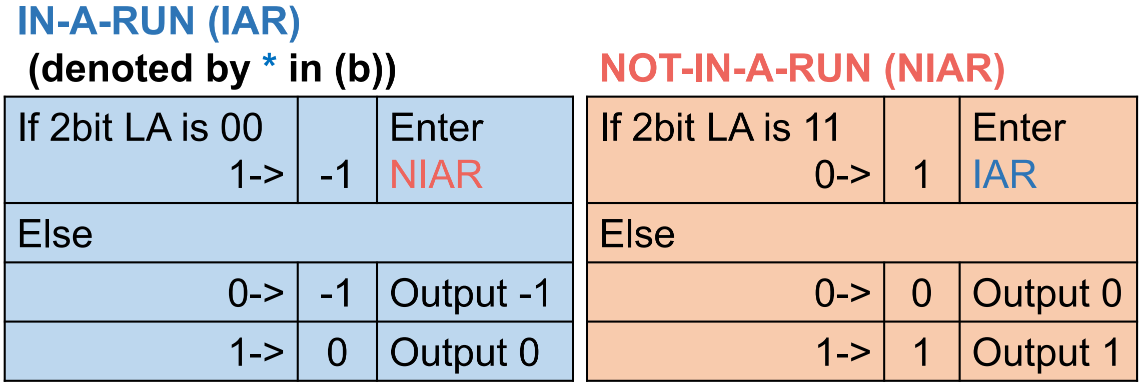

IN-A-RUN (IAR)
(denoted by * in (b))

If 2bit LA is 00 1->	-1	Enter NIAR
Else		
0->	-1	Output -1
1->	0	Output 0

NOT-IN-A-RUN (NIAR)

If 2bit LA is 11 0->	1	Enter IAR
Else		
0->	0	Output 0
1->	1	Output 1

Figure 3-2: An illustration of term quantization (TQ), which keeps the largest non-zero 10 terms across a group of 5 data.

As shown in Figure 3-2, HESE SDR [1][48] is an efficient one-pass encoding scheme to produce minimum-length signed expression representations. Along with the term quantization (TQ) in [44], HESE SDR can further reduce the non-zero terms. This L2R HESE SDR with 2-bit LA is co-designed between the encoding algorithm and the circuit design, which can enable a direct HESE SAR ADC design. The rules for

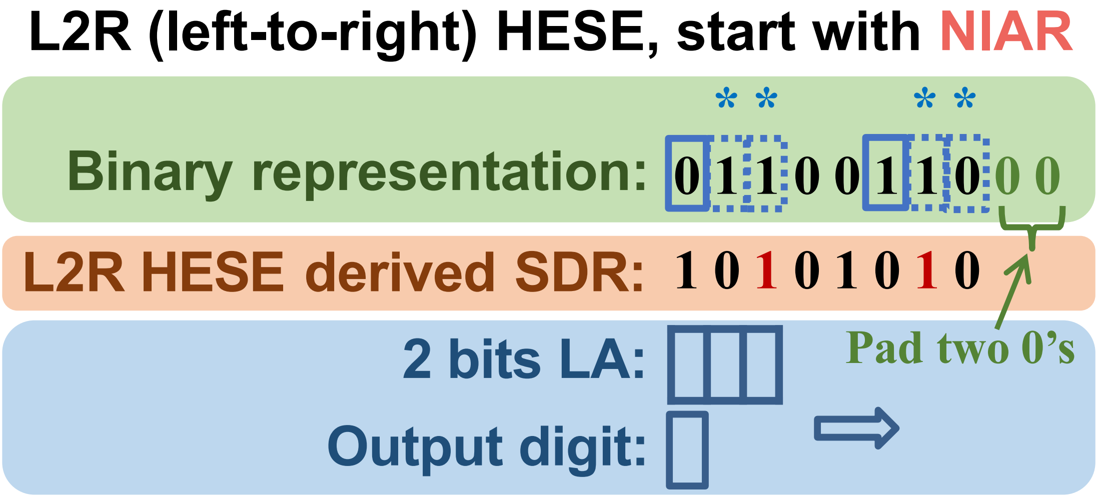
L2R (left-to-right) HESE, start with NIAR
* *
* *
Binary representation: 0 1 1 0 0 1 1 0 0 0
L2R HESE derived SDR: 1 0 1 0 1 0 1 0
Pad two 0's
2 bits LA:
Output digit:

Figure 3-3: The encoding looks at the current bit and the next 2 bits to decide the encoded term. LA stands for look-ahead. Left to right (L2R) HESE SDR finds a minimum-length SDR and red 1 stands for -1

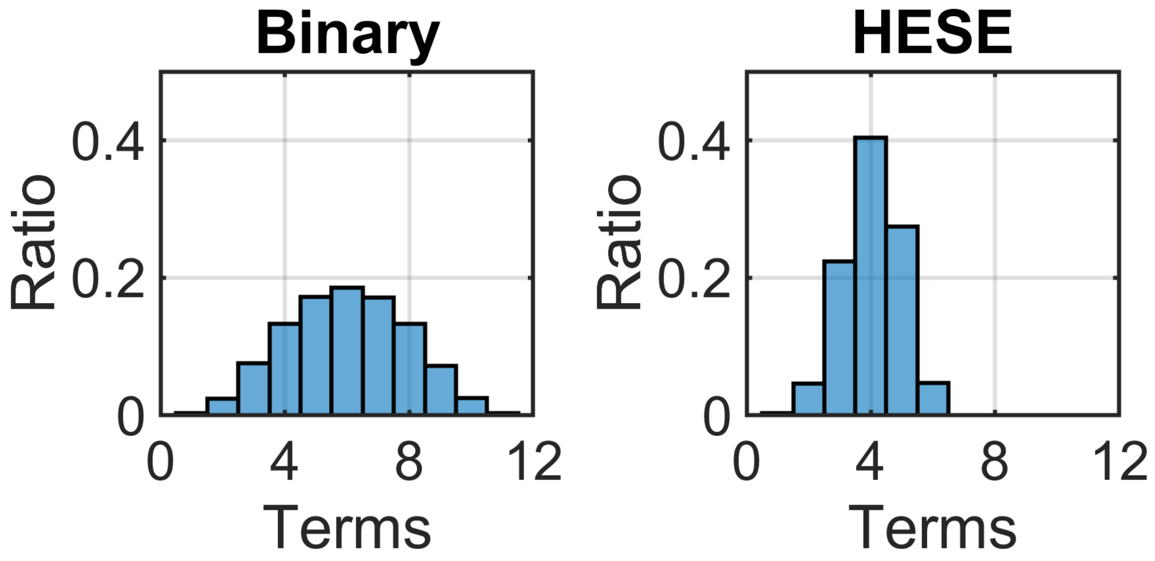

Figure 3-4: Over 95% of the weights in a pre-trained AlexNet [2] can be represented by only half of the HESE SDR terms due to bit-level sparsity

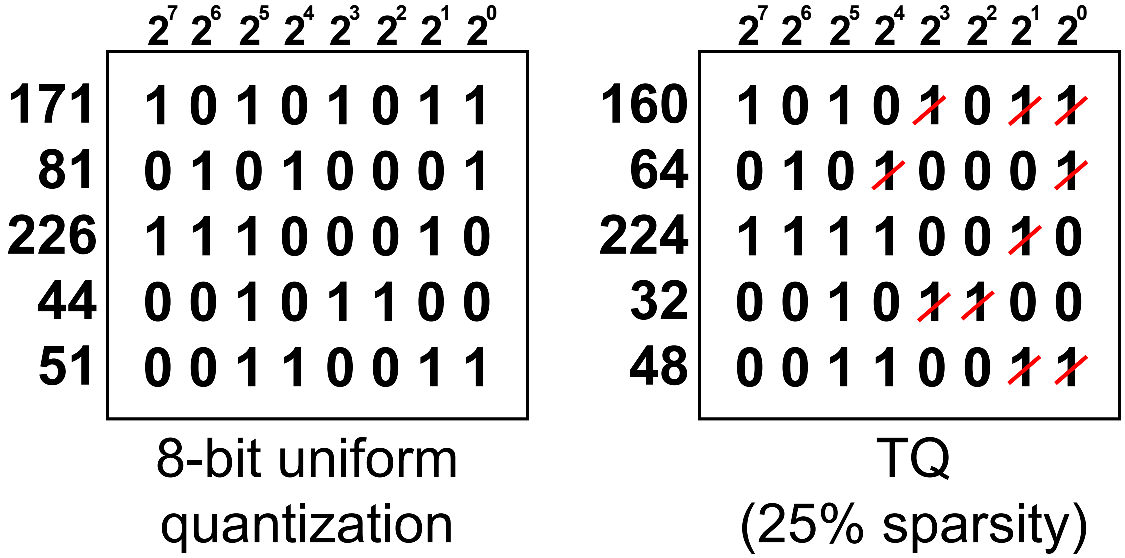

Figure 3-5: An illustration of term quantization (TQ), which keeps the largest non-zero 10 terms across a group of 5 data.

HESE SDR are shown in Figure 3-2. Figure 3-3 shows the illustration of the one-pass HESE SDR encoding. Figure 3-4 shows the weight distribution of a pre-trained AlexNet and terms distribution. HESE SDR can greatly reduce the non-zero terms and thus increase sparsity.

3.2.3 Term Quantization

Term quantization (TQ) [1] prunes out small terms in a group of data, as shown in Figure 3-5. TQ can increase the bit-level sparsity with a trivial impact on the classification accuracy. For uniform quantization, all values are truncated uniformly. TQ truncates the data in a group of data, which can keep more information

3.3 Contributions

The research introduced in this work is situated within the vast application space of Artificial Intelligence of Things (AIoT) which has a significant potential for growth in the future, particularly within the context of edge computing and long battery life. The research aims to improve energy efficiency in AIoT devices such as Voice

Activity Detection (VAD) and Keyword Spotting (KWS) systems which have a need to be always-on and highly efficient in inference.

In order to address the energy efficiency issues in conventional AIoT systems, the research explores certain features in these systems such as data sparsity in embedded neural networks to improve energy efficiency.

A novel bit-level sparsity-aware successive approximation register (SAR) Analog-to-Digital Converter (ADC) is proposed that directly produces Hybrid Encoding for Signed Expressions (HESE). Coupled with term quantization, this approach increases sparsity in Analog Neural Networks (ANNs). This direct HESE SAR ADC is the first of its kind and can provide sparse encoding during the analog-to-digital conversion process.

The research offers a unique approach to increase energy efficiency and sparsity in ANNs. It holds practical implications for real-world scenarios where AIoT devices are required to be always-on and efficient, making it a promising solution for future AIoT developments.

3.4 Proposed Direct HESE SAR ADC

The global architecture of the proposed SAR ADC is shown in Figure 3-6. The SAR ADC is split into two half DACs and two half-sized comparators to provide two thresholds. The SAR ADC has two thresholds for each bit-cycling to perform the 2-bit look-ahead (LA). Noise averaging between the two halves is performed in the last couple of cycles to eliminate noise penalty due to half size DACs and comparators. Bottom-plate sampling is used and the sampling switches are bootstrapped to enhance the linearity. The comparators are fully dynamic with no static power. Two foreground calibration schemes are implemented to improve the linearity. One is the bridge capacitor calibration [49]. The other is the 4 largest MSB capacitors calibration [50].

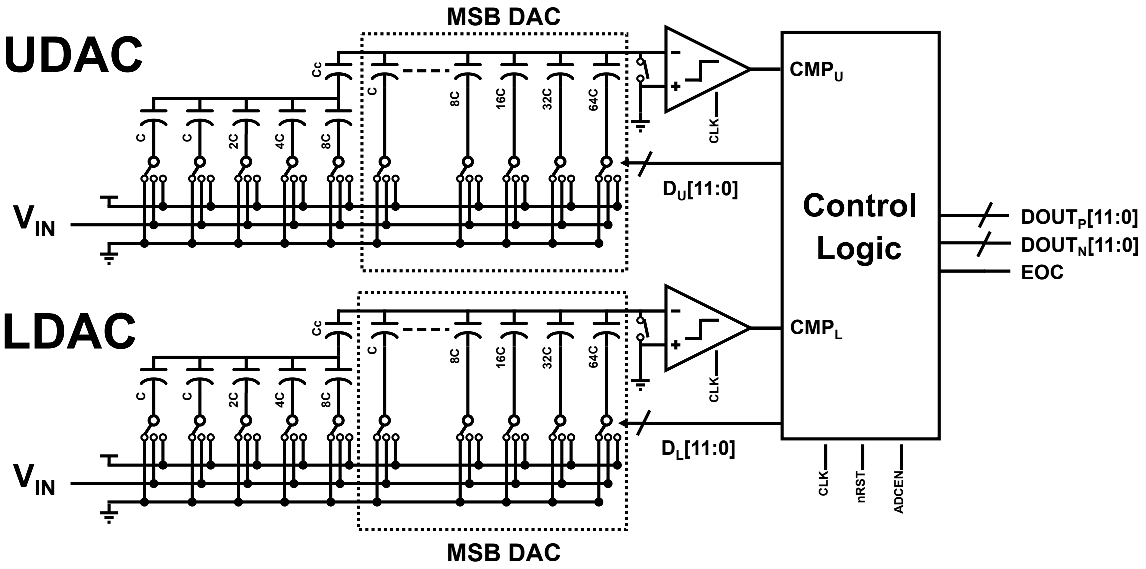

Figure 3-6: Global architecture of the proposed SAR is shown. Two DACs and comparators are implemented for the 2-bit look-ahead (LA) of hybrid encoding for signed expressions (HESE) SDR. Noise averaging (NA) is used to reduce the capacitor size. The accumulated current is converted to the voltage by the sample and hold circuitry.

3.4.1 Conversion Plan of the HESE SAR

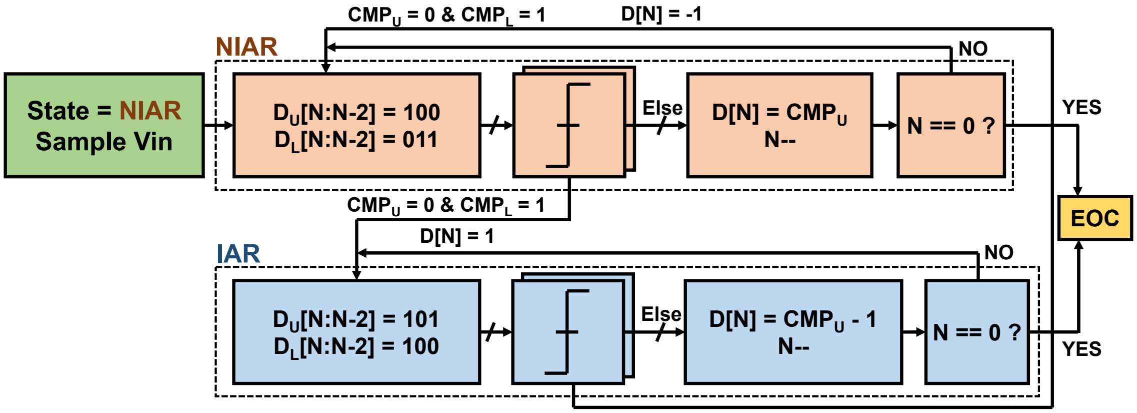

Figure 3-7: Flowchart of the conversion plan

The HESE SAR switches between the IN-A-RUN (IAR) and the NOT-IN-A-RUN (NIAR) state when the input lies between two thresholds. The extra threshold provides the analog 2-bit look-ahead (LA) for direct HESE. V_{IN} stands for the sampled

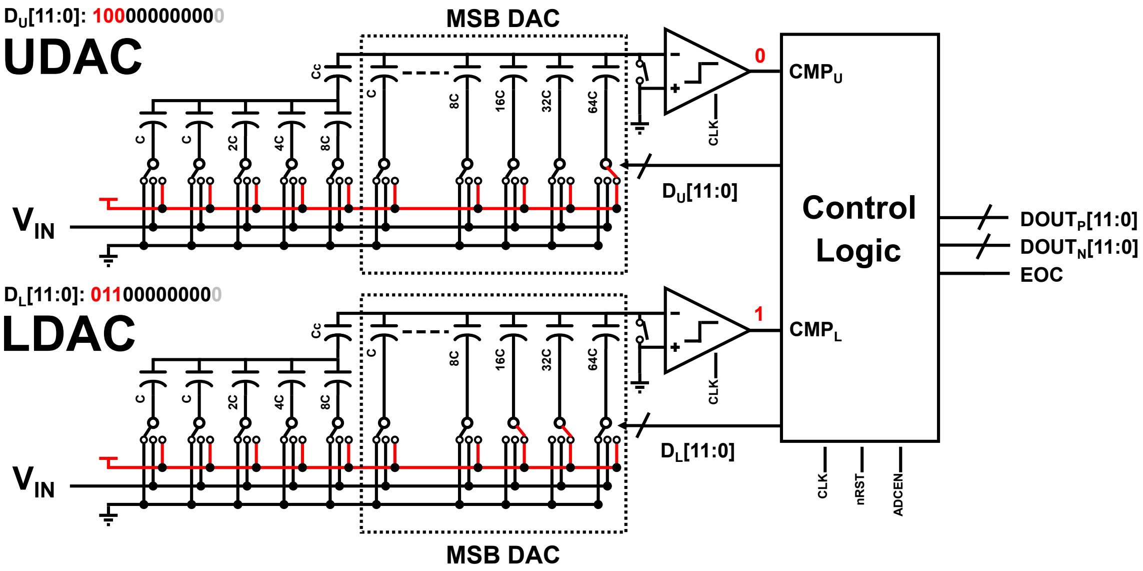

Figure 3-8: The HESE SAR starts in the NIAR state. D_U[N:N-2] and D_L[N:N-2] are set to 100 and 011, respectively. The SAR can look for 2bit LA of two consecutive 1's with this configuration. When CMP_U is 0 and CMP_L is 1, D[N] is encoded to 1 and the SAR switches to the IAR state.

The SAR starts with the NIAR state. As shown in Figure 3-8, D_U and D_L are the digital inputs to the upper and lower DAC, respectively. In the NIAR state, D_U[N:N-2] and D_L[N:N-2] are set to 100 and 011, respectively. The SAR can look for 2bit LA of two consecutive 1's with this configuration. When CMP_U is 0 and CMP_L is 1, D[N] is encoded to 1 and the SAR switches to the IAR state. Otherwise, D[N] = CMP_U.

In the IAR state, D_U[N:N-2] and D_L[N:N-2] are set to 101 and 100, respectively. The SAR looks for 2bit LA of two consecutive 0's. When CMP_U is 0 and CMP_L is 1, D[N] is encoded to -1 and the SAR switches to the NIAR state. Otherwise, D[N] = CMP_U - 1.

When N = 1 or 0, the SAR enters the ending states. The ending states are slightly different to handle the padded zeros. If N = 1 or 0 and the SAR is in the NIAR state, the LA is not necessary and 2bit LA cannot be two consecutive 1's because only zeros

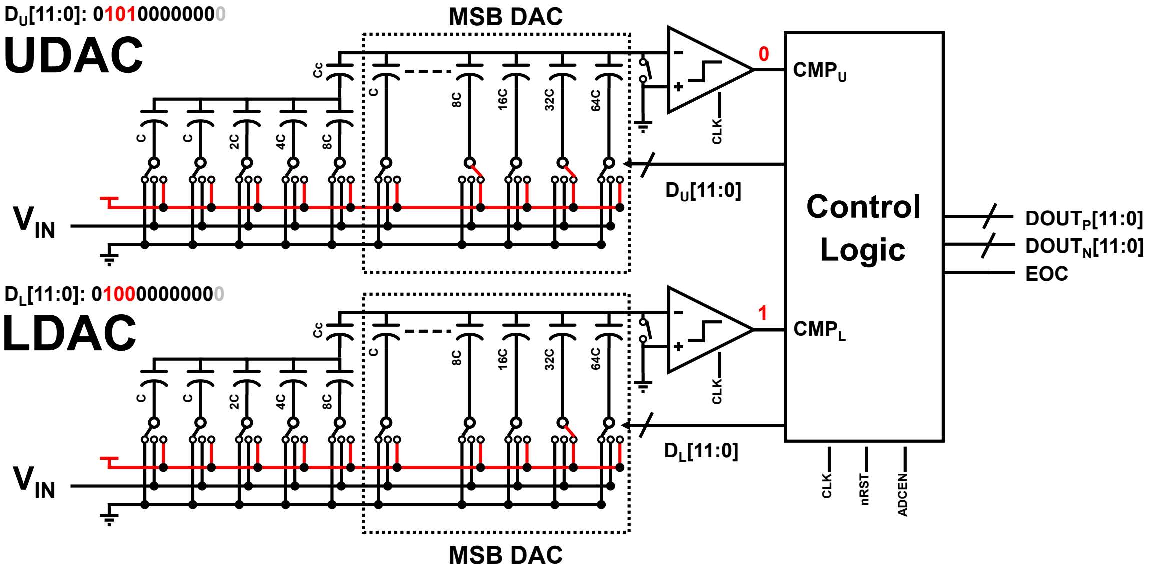

Figure 3-9: In the IAR state, D_U[N:N-2] and D_L[N:N-2] are set to 101 and 100, respectively. The SAR can look for 2bit LA of two consecutive 0's. When CMP_U is 0 and CMP_L is 1, D[N] is encoded to -1 and the SAR switches to the NIAR state.

are padded. If N = 1 and the SAR is in the IAR state, the current bit and next bit of D_U and D_L are set to 11 and 10, respectively, to look for 2bit LA of two consecutive 0's. If N = 0 and the SAR is in the IAR state, the encoded output is -1 regardless of the outputs of the comparators due to the padded two zeros. The NIAR and the IAR switch when CMP_U is 0 and CMP_L is 1. When LA is not necessary, the LDAC and UDAC are connected in parallel to perform noise averaging.

Figure 3-10 shows an example conversion waveform of the HESE SAR. Only 6 MSBs are shown for simplicity. The conversion starts with the NIAR state. In the first cycle, D_U[11:9] is set to 100 and D_L[11:9] is set to 011. If V_{IN} is larger than $VDAC_L$ and smaller than $VDAC_H$, the 2bit LA is 11. The HESE SAR enters the IAR state which would provide negative ones to increase sparsity. In the second cycle, D_U[10:8] is set to 101 and D_L[10:8] is set to 100. V_{IN} is larger than $VDAC_U$ and the HESE SAR stays in the IAR state. In the third cycle, V_{IN} is within two thresholds and the HESE SAR enters the NIAR state.

To reduce the sampled kT/C noise and the comparator noise, both the UDAC and the LDAC are connected in parallel except for the dummy LSB capacitors in the last two bit-cycles where LA is not necessary. Since both the UDAC and LDAC have 11-

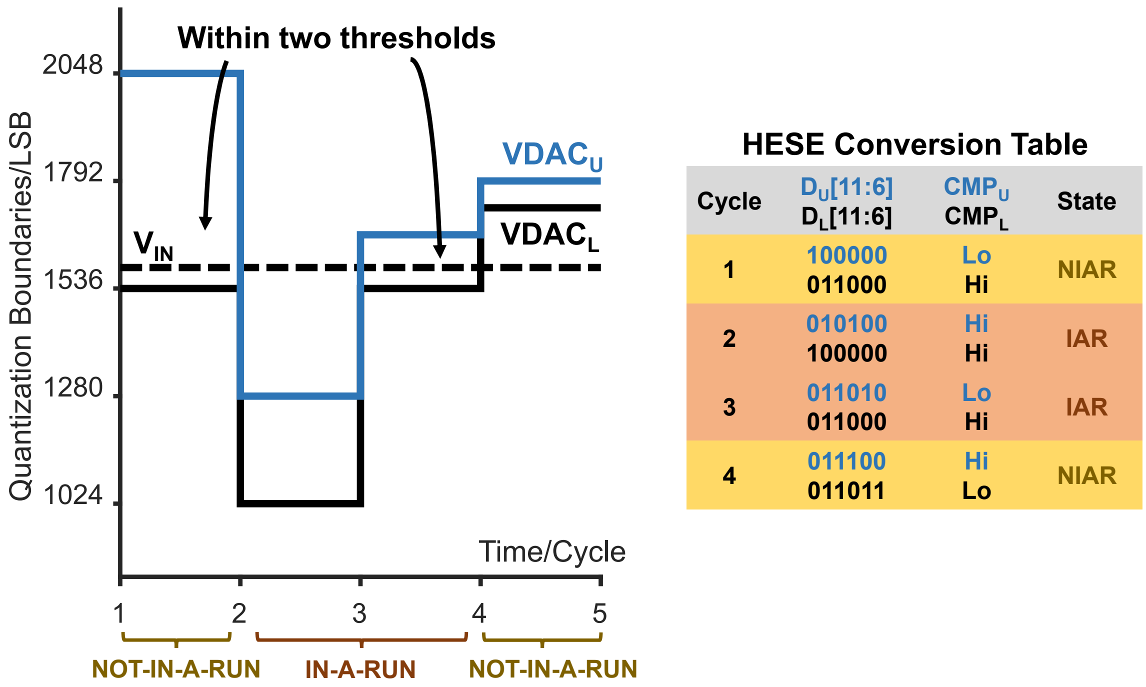

HESE Conversion Table

Cycle	D_U[11:6] D_L[11:6]	CMP_U CMP_L	State
1	100000 011000	Lo Hi	NIAR
2	010100 100000	Hi Hi	IAR
3	011010 011000	Lo Hi	IAR
4	011100 011011	Hi Lo	NIAR

Figure 3-10: An example conversion waveform of the HESE SAR. $VDAC_U$ stands for the output of upper DAC. V_{IN} stands for the output of lower DAC. CMP_U is the output of the upper comparator. CMP_L is the output of the lower comparator. Comparator output is 1 when $V_{IN} > V_{DAC}$.

bit resolution, one more bit decision is required to provide an 12-bit result. The LSB capacitors are separately actuated to provide 1 extra bit decision. Compared with conventional SAR ADCs, no power/area/noise penalty is incurred in the proposed direct HESE SAR ADC.

3.4.2 Sampling Network

A top-plate sampling network has several advantages over a bottom-plate sampling network in a SAR ADC design [51]. The top-plate sampling network removes the need for MSB capacitor in the CDAC. This brings significant area and energy saving.

The direct HESE SAR uses a bottom-plate sampling network to avoid the disadvantages of a top-plate sampling network to achieve 12-bit linearity. Firstly, the sampling switch has a charge injection effect on the top-plate nodes of the CDAC when ADC turns off its sampling switch [52]. This introduces a gain error since the injected charge is a function of the ADC input voltage. The gain error can be calibrated using background calibration technique [53]. The channel charge can also result in nonlinearity since the threshold voltage of the sampling switch is a nonlinear function of the ADC input voltage. This is caused by the back-gate effect. The charge injection is constant because the sampling switch always sees the same voltage. The charge redistribution is not affected on the top-plate nodes.

Secondly, the bottom-plate sampling removes the effect of charge sharing with the parasitic capacitor on the top-plate node. For the top-plate sampling, the input is charging the parasitic capacitor and the CDAC during sampling. However, during the conversion the CDAC charge is shared with the parasitic capacitance. This causes nonlinearity if the parasitic capacitance is nonlinear. This is caused by the nonlinear sampling switch junction capacitor and the comparator input capacitor. For the bottom-plate sampling, when the conversion is complete, the voltage across the parasitic capacitance returns to its initial condition reversing any charge sharing between the CDAC and the parasitic capacitance. Therefore, the bottom-plate sampling does not suffer from gain error or nonlinearity due to nonlinear parasitic capacitance.

As shown in Figure 3-11, bootstrapped switches are also used to achieve a 12-bit

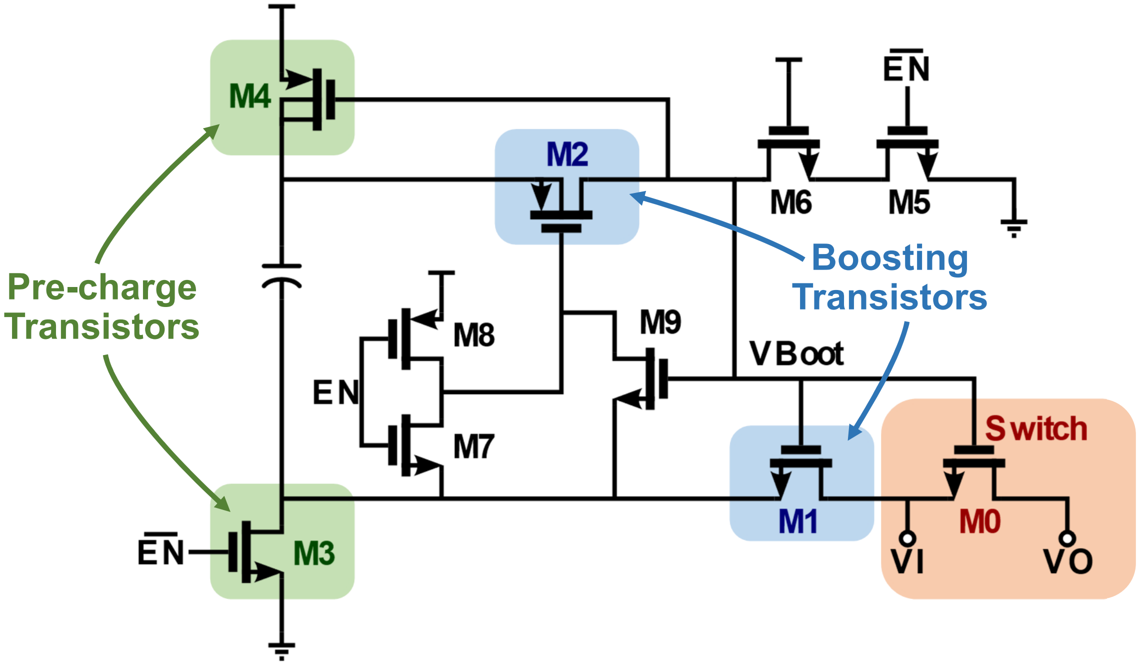

Figure 3-11: The schematic of the bootstrapped switch

resolution. The on-resistance of the sampling switch depends on the input voltage level and this brings nonlinearity in high accuracy ADCs [54]. The RaM SAR bootstraps the V_{GS} of its bottom-plate tracking switches with a charge-pump circuit [55]. The V_{GS} remains constant for all ADC inputs. The switches of the sampling network are sized to make the largest input settling error lower than 0.01LSB. To reduce the on-resistance of the sampling switches, the width of the sampling switches can be increased.

M3 and M4 are for pre-charing the capacitor. During the pre-charge phase, the capacitor is connected to the ground via M3 and VDD via M4. M1 and M2 are for boosting the gate voltage. During the sampling phase, the source of M0 is connected to the bottom plate of the capacitor via M1 and the gate of M0 is connected to the top plate of the capacitor via M2. The other transistors are for reliability and control signal generation. Note that M2 is PMOS and M0, M1, M5 and M6 are NMOS transistors. M6 and M8 are needed to improve the circuit reliability.

The operation of the boosting switch is described below. EN is the enable signal. When EN is high, the bootstrapped switch is on. When EN is low, the bootstrapped

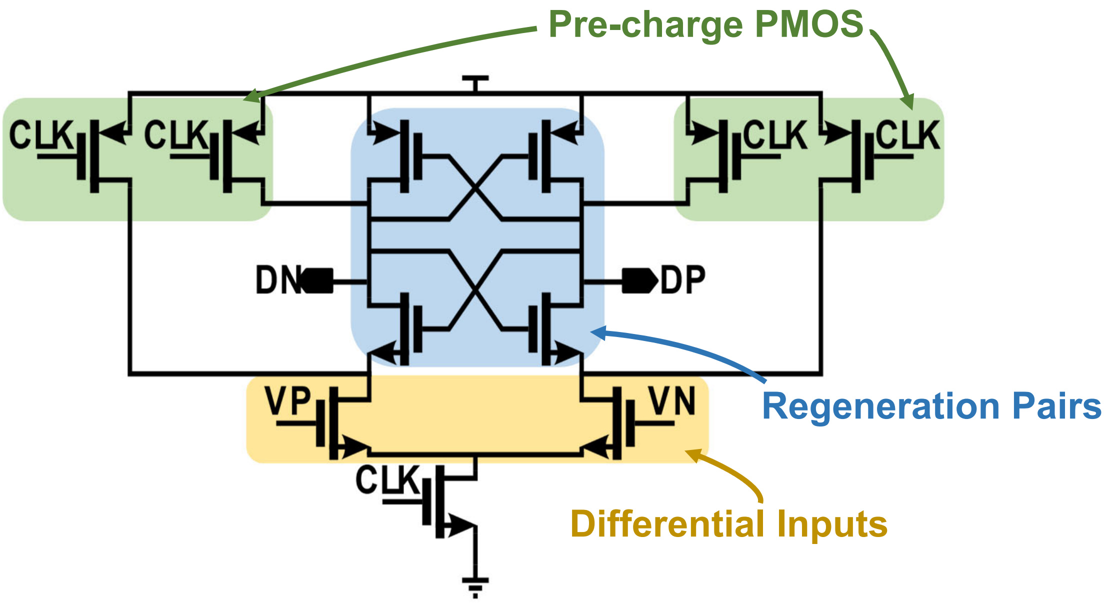

Figure 3-12: The schematic of the dynamic comparator

switch is off. When EN is low, M3 and M4 are turned on. The top and bottom plates of the capacitor are connected to supply voltage and ground, respectively. M7 is off and M8 is on. The gate voltage of M2 is high and M2 is off. VBoost is connected to the ground by M5 and M6. M0, M1, and M9 are off. When EN is high, M3 and M4 are off and the charge is stored in the capacitor. M5 is off and the VBoost is no longer connected to the ground. M2 is on and the VBoot is boosted by the capacitor voltage. Note that this circuit would generate a voltage level that's higher than the supply voltage. M9 and M6 are extra transistors to protect the M2 and M5, respectively. Since the top plate of the capacitor can be larger than the supply voltage, the bulk of M2 and M4 are connected to the top plate of the capacitor.

3.4.3 Comparator

Figure 3-12 shows the strong-arm latch of the SAR ADC [56]. The outputs are fed into an RS-latch for the comparator output. For every bit-cycling, the comparator determines the sign of the differential CDAC top plate voltage and passes the result to the RaM SAR logic. The strong-arm latch has two NMOS pairs and one PMOS pair. VP and VN are differential inputs from the top plate of CDACs. VP and Vn are connected to the differential NMOS input pair. The other NMOS pair and the PMOS pair form a regeneration circuit.

When CLK is low, the tail NMOS is closed and the internal nodes are pre-charged to VDD by PMOS. When the rising edge of CLK comes, the differential pair amplifies the difference between the input voltages and convert the difference to current. The regeneration pairs decide the outputs.

3.4.4 CDAC

Figure 4-1 shows the CDAC of the RaM-SAR ADC. The CDAC is split into two capacitors arrays that are bridged by a bridge capacitor [57] to reduce area.

There are two main considerations for choosing the unit capacitance for the CDAC. The total sampling capacitance should be large enough to suppress KT/C noise based on the SNR target. The unit capacitor of the RaM-SAR is chosen to be 11.3fF.

For the layout, the CDAC uses the common centroid layout technique to suppress the linear gradient effects [58]. The CDAC also uses the equal-edge-ratio layout technique to increase the matching [59]. A p-well ground guardring surrounds the entire CDAC to make the substrate below the CDAC quiet. Routing of the CDAC is done carefully to ensure the symmetry and all the metal density rules are met by manually placing ground-connected dummy metals. This avoids random dummy metal insertion of the foundry.

3.5 Results

A prototype chip is fabricated in low-power 65nm technology. The chip micrograph is shown in Figure 3-13. The prototype demonstrates direct sparse encoding along with A/D conversion.

Figure 3-14 shows that the HESE SDR minimizes the non-zero terms compared to the binary encoding. The HESE SDR can save up to 60% of the terms compared to binary encoding.

Figure 3-15 shows the measured spectrum of the direct HESE SAR. The effective

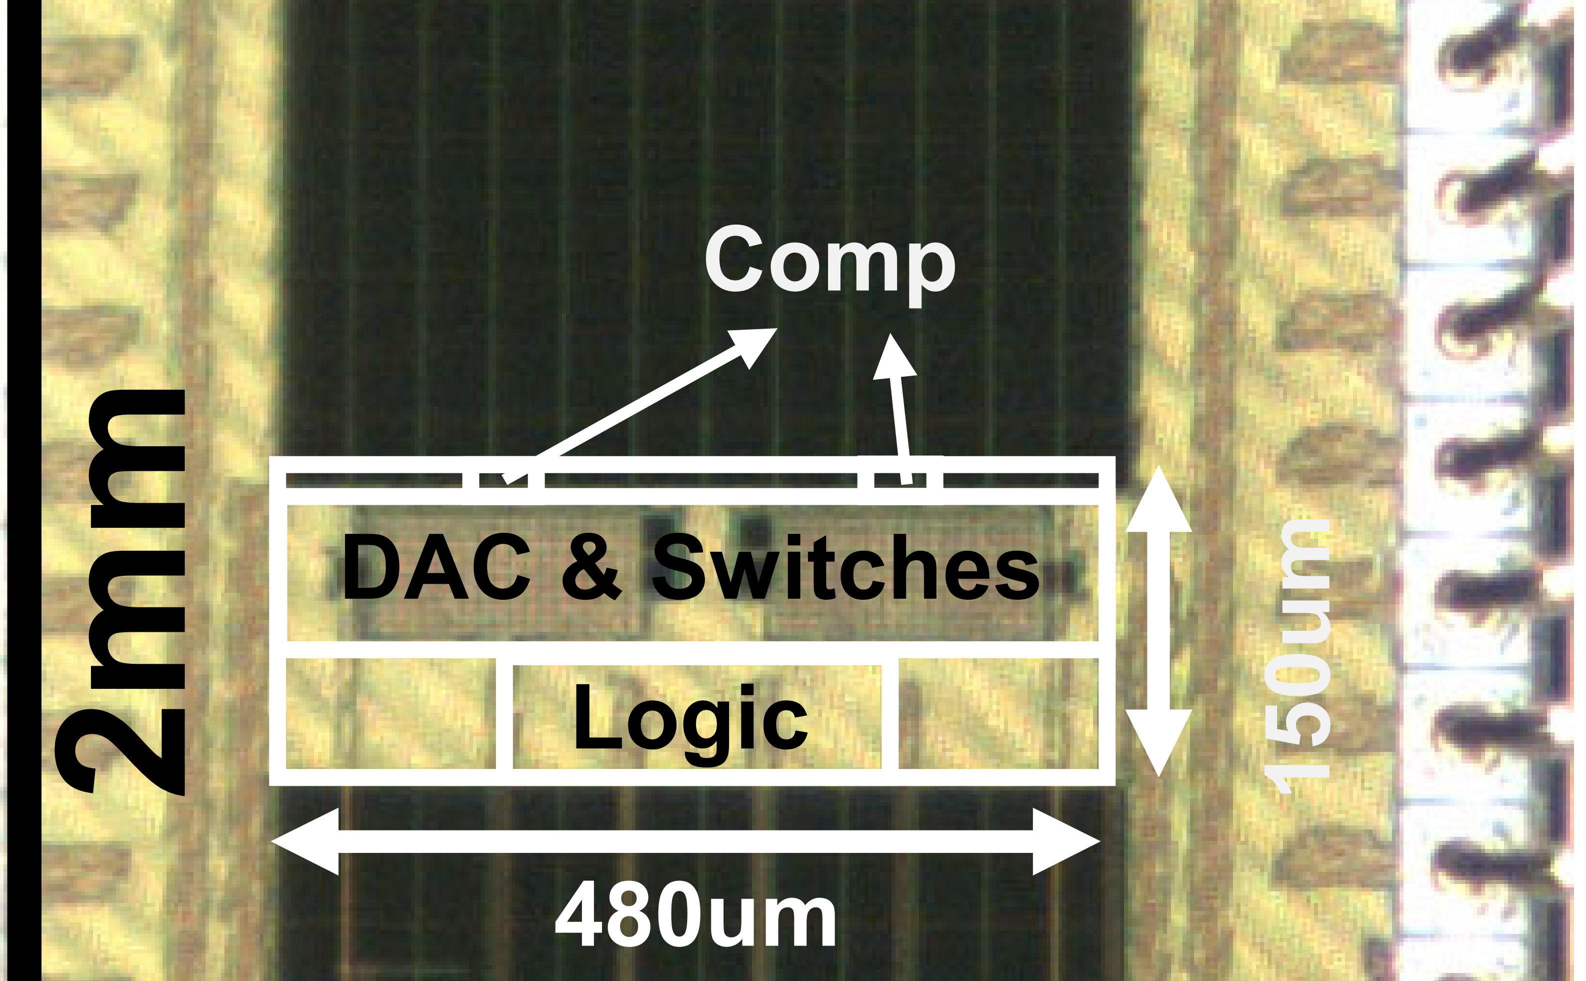

Figure 3-13: Chip micrograph

number of bits (ENOB) is 10.6b. Signal bins are highlighted. The sampling rate is 45MS/s. 16384 data points are used for FFT calculation. The input frequency is a 22.1MHz sine wave. The positive and negative outputs are read out separately and reconstructed to binary representations for FFT calculation.

The HESE SAR is the first sparsity-aware SAR ADC to demonstrate direct sparsity encoding with competitive energy efficiency, resolution, and area.

3.6 Conclusion

This work is the first bit-level sparsity-aware ADC in ANNs with direct hybrid encoding for signed expressions (HESE) leveraging algorithm-circuit co-design. ANN with HESE SAR minimizes the non-zero terms and enables a reduction in energy along with the term quantization (TQ). A prototype in 65nm low-power technology achieves Walden FoM of 15.2fJ/conv.-step at 45MS/s. The direct HESE SAR offers a

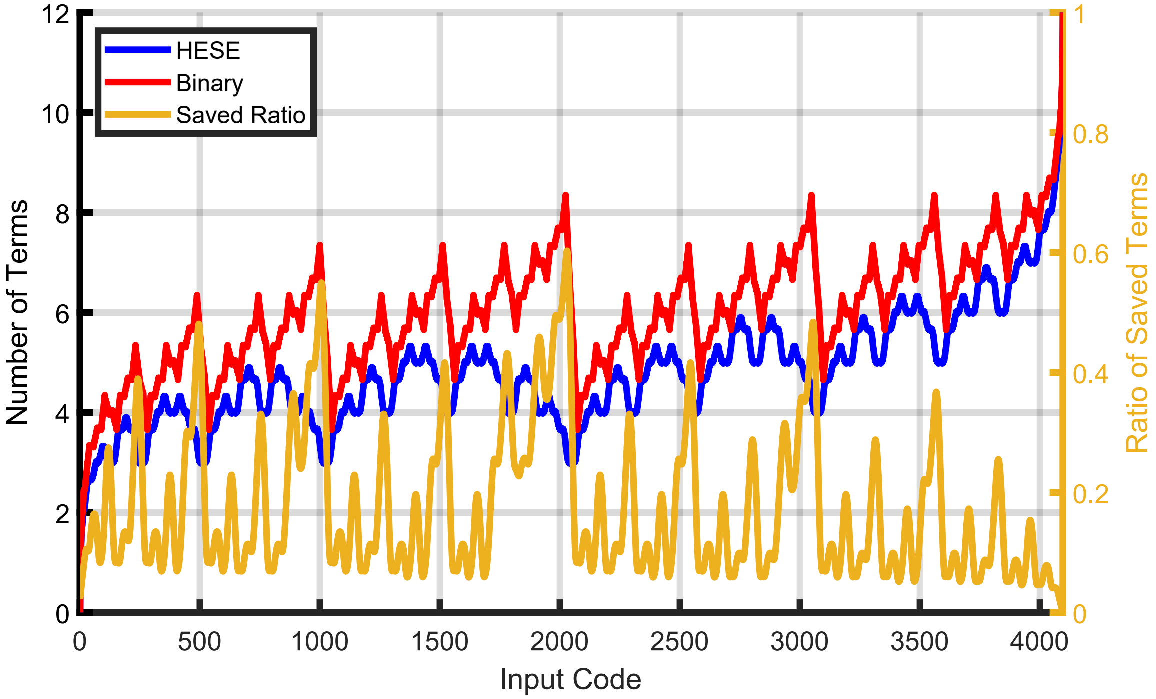

Figure 3-14: Simulated number of terms of the HESE and the binary encoding for all the 12-bit digital codes. On average, the HESE saves 23% of the terms compared to the binary encoding. Note that the terms refer to the non-zero digits. For HESE, the terms include both positive and negative ones. For binary, the terms include only positive ones since there are no negative ones in the binary encoding.

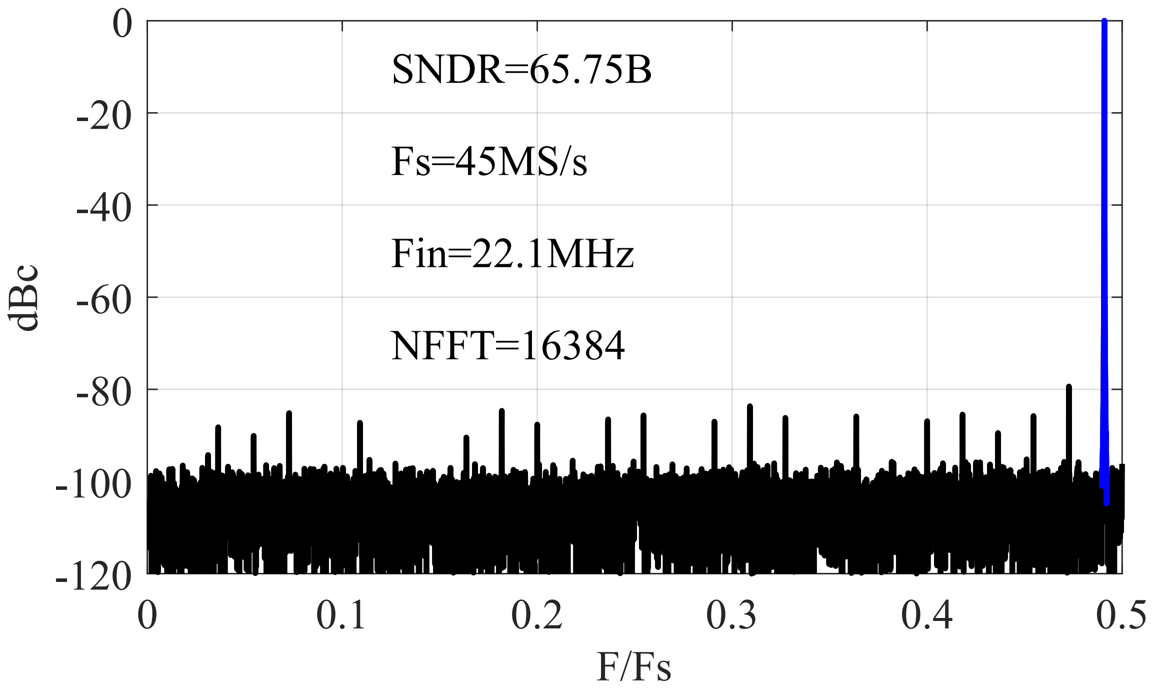

Figure 3-15: Measured spectrum of the HESE SAR

general direction for the ADC design in ANNs leveraging bit-level sparsity. The core area is 0.072mm^2.

Chapter 4

RaM-SAR: The first secure ADC for high-speed applications

4.1 Introduction

The design and implementation of secure analog-to-digital converters (ADCs) have become a crucial challenge in the field of integrated circuit design, especially in the context of low-power, high-speed, and small form-factor systems such as those used in IoT and wearable devices. The security of these systems is often threatened by side-channel attacks, such as power analysis and electromagnetic (EM) analysis, which can extract sensitive information from the power consumption or electromagnetic emissions of the system.

To address this challenge, we present RaM-SAR, a secure random-mapping SAR ADC that provides resistance against both power and EM side-channel attacks. RaM-SAR is a 12-bit, 25 MS/s ADC with an energy consumption of 11.3 fJ per conversion step. This is achieved through the use of a novel random-mapping technique, where each conversion is randomly mapped to one of the thousands of conversion sequences. This randomization of the power supply traces helps to protect against neural network-based power and EM side-channel attacks.

The RaM-SAR ADC offers several advantages over prior works in the field of secure ADCs. Firstly, it provides protection with much lower energy and area overheads compared to prior works. This is a significant advantage as it enables the integration of secure ADCs into systems with limited power and space resources, such as IoT and wearable devices. Secondly, the prototype RaM-SAR ADC, implemented in 65nm CMOS, demonstrates significant improvements in terms of performance. It has a 12.5 times higher bandwidth and 4.8 times better energy efficiency compared to prior secure ADCs. This makes the RaM-SAR ADC a competitive option for secure high-speed data conversion applications.

A SAR ADC architecture is a popular choice for high-speed, low-power, and high-resolution ADCs. The SAR ADC architecture is well-suited for the random-mapping technique used in RaM-SAR as it can handle large numbers of conversion sequences with low overhead. The random-mapping technique involves randomly selecting one of the thousands of conversion sequences for each digital output, thus randomizing the power consumption patterns and making it more difficult for attackers to extract sensitive information.

To evaluate the performance of the RaM-SAR ADC, a prototype was implemented in 65nm CMOS technology. The prototype was tested in a variety of scenarios to assess its accuracy, speed, and energy consumption, as well as its resistance against power and EM side-channel attacks. The results showed that the RaM-SAR ADC achieved high accuracy with a 10.9 ENOB. It also demonstrated high-speed performance, with a conversion rate of 25 MS/s. Additionally, the energy efficiency of 11.3 fJ per conversion step was found to be significantly lower than that of prior works, making it suitable for low-power applications. The RaM-SAR ADC was found to be resilient against power and EM side-channel attacks, with randomization of the power supply traces making it difficult for attackers to extract sensitive information.

The RaM-SAR ADC presents a significant contribution to the field of secure ADCs. It offers protection against power and EM side-channel attacks with much lower energy and area overheads compared to prior works. The high accuracy, speed, and energy efficiency of the RaM-SAR ADC make it a competitive option for secure high-speed data conversion applications. The random-mapping technique used in RaM-SAR provides a promising direction for future work in the field of secure ADCs

The RaM-SAR reduces the area and energy overhead and increases the sampling rate, which broadens the secure ADCs into video applications.

4.2 Contributions

This ADC has been specifically designed to provide resistance against both power and electromagnetic (EM) side-channel attacks, which are typical threats to low-power, high-speed, and small form-factor systems such as those used in IoT and wearable devices.

The RaM-SAR ADC achieves security through a novel random-mapping technique. This process involves randomly mapping each conversion to one of thousands of conversion sequences, which randomizes the power supply traces, making it more difficult for attackers to extract sensitive information. This random-mapping technique is applied to a Successive Approximation Register (SAR) ADC architecture, which is suitable for low-power applications.

Compared to prior works, the RaM-SAR ADC offers significant advantages including lower energy and area overheads, which enables its integration into systems with limited power and space resources. It also demonstrates significant performance improvements, with a higher bandwidth and better energy efficiency compared to prior secure ADCs.

The research also discusses the major leakage sources of unprotected SAR ADCs, which include the CDAC, comparator, and SAR logic, and explains how the design of RaM-SAR mitigates these leakage sources. The research presents a threat model for ADC power and electromagnetic side-channel attacks and discusses the methods for carrying out these attacks. The research also discusses Convolutional Neural Network-based Power Side-Channel Attack (CNN-PSA) and Convolutional Neural Network-based Electro-Magnetic Side-Channel Attack (CNN-EMSA), two proposed side-channel attacks that leverage the power of convolutional neural networks.

4.3 Proposed RaM-SAR

In this work, we propose an energy-efficient and high-speed secure SAR based on random mapping (RaM-SAR) to randomize the conversion scheme (Figure. 2-3, right). It is based on the LSB-first SAR [60]. LSB-first SAR saves energy when the signal acitivty is low. In a typical case where the inputs are the same for consecutive samples, the LSB-first SAR finishes the conversion in 2 bit-cycles. Conversion energy for redundant bit-cycles are saved. The LSB-first SAR has 3 phases, DIR, ToMSB, and ToLSB phase. After sampling the input to the CDAC, the LSB-first SAR decides the direction of the bit flipping in the DIR phase. The sampled input is compared to the initial guess of the LSB-first SAR. The initial guess is set to the previous A/D conversion to save energy in low-activity signals. For EEG, input change between samples is usually much smaller than the full scale, and sometimes the input doesn't change between samples. During the DIR phase, the lower bound of the initial guess is tested first. The upper bound of the initial guess is tested by flipping the dummy LSB capacitor. If the input is larger than the initial guess, the test voltage is moving up. If the input is smaller than the initial guess, the test voltage is moving down. During the ToMSB phase, the test bits are flipped from LSB to MSB until the test voltage overshoots the input. During the ToLSB phase, a conventional binary search is conducted from the overshoot bit to the LSB.

The LSB-first SAR has different conversion switching sequences for each digital output but it's deterministic as the initial guess is always the previous digital output. Therefore, the LSB-first SAR is not inherently safe from attacks. However, by randomizing the initial guess instead, the conversion sequence for a given input is also randomized. The LSB-first SAR needs up to 25 bit-cycles for a 12bit conversion, which limits its usage in high sampling rate applications. As will be explained, the RaM-SAR uses only 15 bit-cycles in the worst case. A RaM-SAR prototype fabricated in a 65nm CMOS achieved a FoM of 11.3fJ/conv.-step.

Figure. 4-1 shows the single-ended version of the proposed RaM-SAR architecture along with its block design. A differential version is implemented. The SAR ADC is

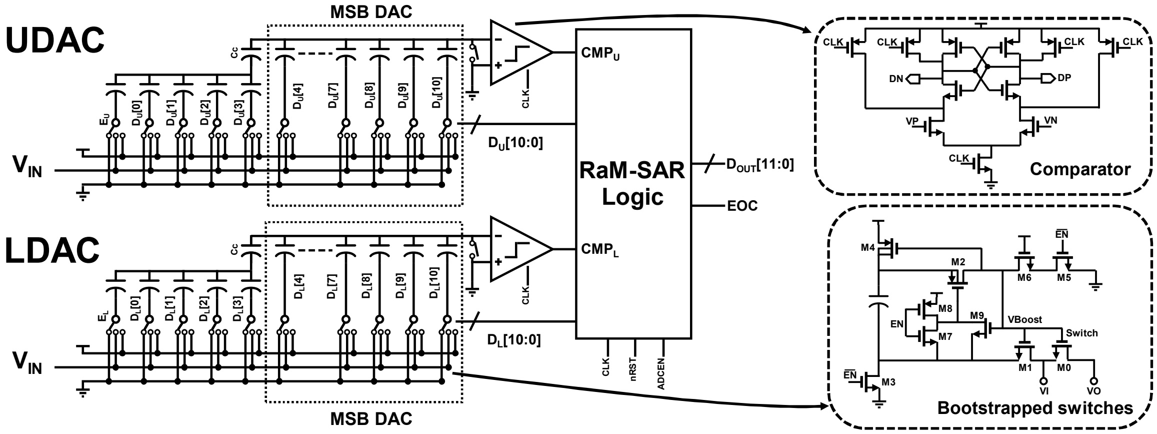

Figure 4-1: Proposed RaM-SAR architecture and block diagram

split into two half DACs and two half-sized comparators to provide two thresholds. Two comparators' outputs are fed into the RaM-SAR logic. The DAC is controlled by the logic output. Strong-Arm comparator is used. The sampling switches are bootstrapped to increase the linearity.

Two thresholds reduce the worst-case number of cycles needed per conversion from 25 to 15. As in the direct HESE SAR ADC, noise averaging between the two halves is performed in the last two bit-cycles to eliminate the noise penalty due to half-size DACs and comparators. Therefore, there is no extra area/power overhead in this split DAC design. Foreground calibration removes the effect of capacitor mismatches.

Figure. 4-2 shows example conversions of RaM-SAR. In the first phase (P1), the differential input voltage is sampled onto the bottom plates of the CDAC (Figure. 4-3). DRND, the 11b pseudo-random number generated by on-chip linear feedback shift registers, is set to the random start. The thresholds of the comparators are set according to the random start. If both outputs of the comparators are high, then the random start was too low, so the direction of bit-cycling DIR is set to 1 to increase the thresholds. The inverse holds if both outputs of the comparators are low. The DIR controls both extra LSBs. Otherwise, the random start is correct, and the RaM-SAR combines the DACs for LSB decision as in Figure. 4-4.

The conversion sequence is randomized by the random start, which significantly weakens the correlation between power/EM side-channel leakage and digital outputs.

Samp.1 Conversion Table

	Cycle	$D_L[5:0] + E_L$ $D_U[5:0] + E_U$	CMP_L CMP_U
P1[1]	1	000000 + 00 000000 + 10	Hi Hi
P2[1]	2	000010 + 10 000110 + 10	Hi Hi
	3	001110 + 10 011110 + 10	Hi Hi
	4	111110 + 10 111110 + 10	Lo Lo
P3[1]	5	101010 + 10 110010 + 10	Lo Lo
	6	100010 + 10 100110 + 10	Lo Lo
P4[2]	7	100010 + 01	Lo
	8	100000 + 01	Lo
	End	100000	

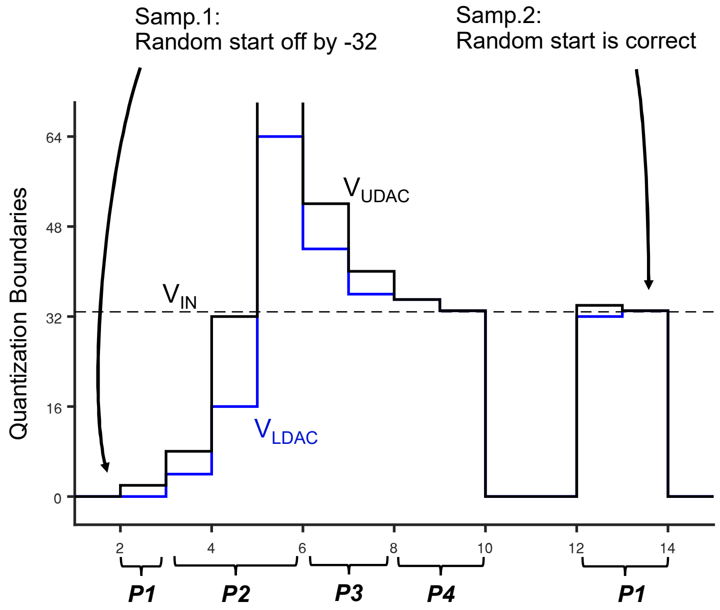

[1] LSB in P1~P3 is in gray code since the weight of E_L and E_U is 2LSB in the 12b resolution when the DAC is split into UDAC and LDAC.

[2] E_L is always 0 after combination of UDAC and LDAC in P4. The weight of E_U now is one LSB due to the combination in P4.

Figure 4-2: Example conversions of RaM-SAR

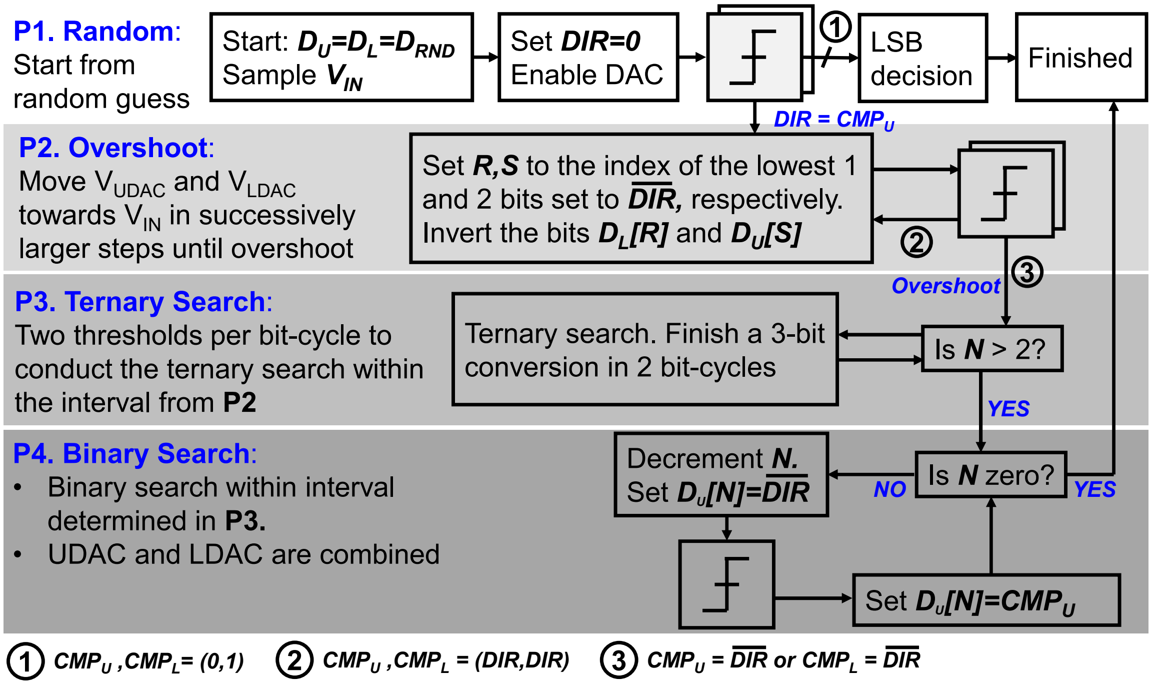
P1. Random:
Start from random guess
Start: D_U=D_L=D_{RND}
Sample V_{IN}
Set DIR=0
Enable DAC
1
LSB decision
Finished
DIR = CMP_U
P2. Overshoot:
Move V_{UDAC} and V_{LDAC} towards V_{IN} in successively larger steps until overshoot
Set R,S to the index of the lowest 1 and 2 bits set to $\overline{DIR}$, respectively.
Invert the bits $D_L[R]$ and $D_U[S]$
2
3
Overshoot
P3. Ternary Search:
Two thresholds per bit-cycle to conduct the ternary search within the interval from P2
Ternary search. Finish a 3-bit conversion in 2 bit-cycles
Is N > 2?
YES
P4. Binary Search:
• Binary search within interval determined in P3.
• UDAC and LDAC are combined
Decrement N.
Set $D_u[N]=\overline{DIR}$
NO
Is N zero?
YES
Set $D_u[N]=CMP_U$
1 CMP_U ,CMP_L= (0,1)
2 CMP_U ,CMP_L = (DIR,DIR)
3 CMP_U = $\overline{DIR}$ or CMP_L = $\overline{DIR}$

Figure 4-3: Flowchart of RaM-SAR

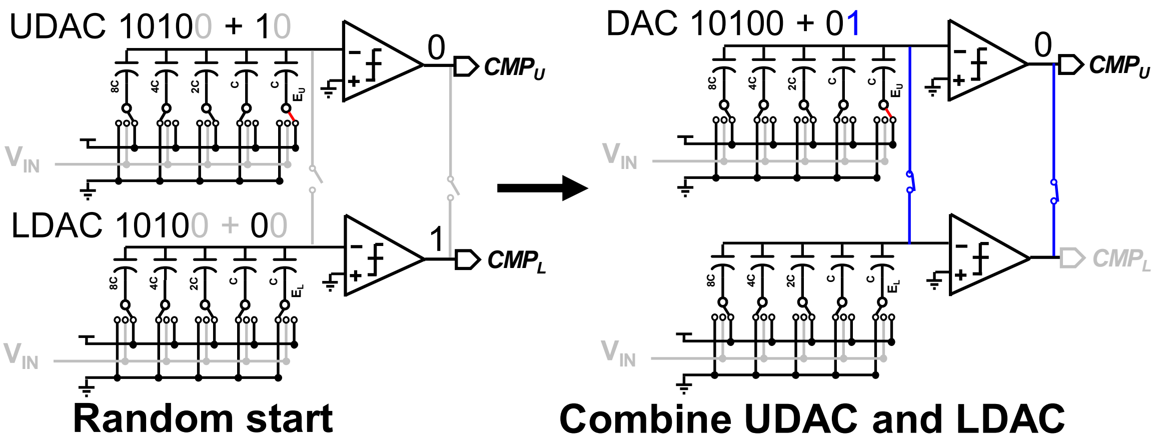

Figure 4-4: Conversion for a correct random start

The two-threshold, two half-DAC arrangement further randomizes the power supply traces. R and S are two registors to support the conversion scheme. In the second phase (P2), R and S are set to the index of the lowest one and two-bits of the random start that are not currently set to DIR, respectively. Two thresholds make P2 faster compared to [60]. Bit R for the lower DAC and bits S for upper DAC are inverted to move in the desired direction. This is repeated until one of the outputs of the comparators flips, indicating that the target value has overshot. The rest of the MSBs are now finalized for this conversion. Let N be the current bit under test. The conversion proceeds to the LSB in the Ternary Search Phase (P3) first and then the Binary Search Phase (P4) if N is less than 3. In the P3, 3b conversion is finished in 2 bit-cycles. The bit-cycling continues as the ToLSB phase in [60] in the P4 and the DACs are combined. When bit-cycling has gone back down to the LSB, conversion is completed, and the DACs are purged.

Implementation of the circuit components including CDAC, comparator, and sampling networks are identical to those of the direct HESE SAR ADC.

4.4 Information Leakage source of the unprotected SAR ADC

To understand the major leakage source of the unprotected SAR ADC, the relationships between the A/D conversion output and the power/EM leakage of CDAC, comparator, and SAR logic are discussed.

CDAC is one of the three major leakage sources of the unprotected SAR ADC. The switching energy correlates to the bit-decision result of the comparator [61]. The second major leakage source is the comparator. Generally, the comparator consumes more power as the differential input voltage becomes smaller. When the CLK in Figure 3-12 is high, the NMOS input pair takes more time to turn on the regeneration pairs and the regeneration latch takes more time to resolve its outputs as the differential input becomes smaller. The third major leakage source is the unprotected SAR logic. The logic consumes a different amount of power depending on the bit-decision result. The logic leaks the entire A/D conversion result from MSB to LSB.

4.5 Neural-network-based ADC PSA and EMSA

4.5.1 Threat Model

The ADC PSA and EMSA are carried out by attackers who have physical access to the same type of target ADCs for the profiled attacks. These attackers can gather power or electromagnetic leakage from the ADCs, potentially compromising the security and integrity of the data being processed. In this analysis, we will explore the intricacies of these attacks and the challenges faced by both attackers and ADC designers in protecting sensitive information.

For the ADC PSA, a resistor is inserted between the off-chip power supply and the power supply pin of the ADC to measure the voltage drop across the resistor. This method allows attackers to monitor the ADC's power consumption and potentially derive information about the A/D conversion process. Similarly, for the ADC

EMSA, a sensitive EM current probe is positioned on top of the ADC for non-contact measurement, capturing electromagnetic emissions from the ADC's operation.

Profiled attacks are required for both ADC PSA and EMSA because the ADC leakage behavior is highly dependent on the specific circuit implementation. Identifying a general model linking the side-channel leakage to the corresponding A/D conversion for arbitrary ADCs can be challenging, necessitating the use of ADC-specific attack approaches.

To execute a profiled ADC PSA, an attacker may procure a training ADC that is nominally identical to the target ADC (e.g., an off-the-shelf product bearing the same part number as the target ADC). Utilizing this training ADC, the attacker can develop the ADC power analysis model by examining the relationship between the power traces and the corresponding A/D conversion results. Since the training ADC contains the same circuits as the target ADC, the developed power analysis model can be employed to attack the target ADC. For testing RaM-SAR, we randomly select a chip as the training ADC and attack three other chips for testing. Unlike digital hardware (e.g., microprocessors or FPGAs) that can be easily updated to incorporate different algorithm implementations, most analog/mixed-signal circuits are application-specific integrated circuits (ASICs) that cannot be modified after fabrication or product release.

Although an attacker might be able to acquire a training ADC, it is reasonable to assume that they will not have access to the transistor-level implementation details of the ADC. Manufacturers typically do not disclose such information in their datasheets, and reverse engineering of integrated circuits (ICs) can be prohibitively expensive. An ADC-specific power model is not practical. Neural-network-based models are stronger because they can learn the correlation between the collected traces and the corresponding A/D conversions.

Furthermore, the analog supply voltage (VDDA) and the digital supply voltage (VDDD) are typically separated. This standard practice serves to shield sensitive analog circuitry from its noisy digital counterpart. Attackers can exclude the disturbance from the digital circuitry by only collecting traces from the VDDA. The

IO supply voltage (VDDIO) is not available in a system-on-chip (SoC) environment, as the ADC delivers its bits to the processor directly. Probing the VDDIO is not practical.

4.5.2 CNN-PSA and CNN-EMSA

Analog-to-Digital Converter (ADC) Convolutional Neural Network-based Power Side-Channel Attack (CNN-PSA) [16] and ADC Convolutional Neural Network-based Electro-Magnetic Side-Channel Attack (CNN-EMSA) [8] are proposed side-channel attacks that leverage the power of convolutional neural networks (CNNs). These attacks are profiled attacks, consisting of two phases: the profiling phase and the attacking phase. Both approaches employ similar conversion analysis models, but they differ in their input formats. ADC PSA utilizes power traces, while ADC EMSA relies on EM leakage. The objective of both attacks is to uncover the relationship between the side-channel leakage and the corresponding A/D conversion outputs.

For the analysis model, one possible method is to draw upon domain knowledge of ADC design for heuristic feature engineering. However, this can be an extremely challenging task. The behavior of an ADC is highly dependent on its circuit implementation, and the detailed structure of the ADC is typically not publicly available. Consequently, a data-driven solution emerges as a more promising approach. By analyzing the collected leakage traces from the profiling phase, the CNNs autonomously identify the most informative parts of the raw traces. This process is achieved by minimizing the loss function using gradient descent and back-propagation algorithms. The loss function quantifies the discrepancy between the parameterized analysis model and the real model.

Our CNN architecture consists of frontend convolutional layers, max-pooling layers, and a flattened layer. Each convolutional layer contains five filters, with a filter size of 5 and a stride of 1. The activation function employed is the Rectified Linear Unit (ReLU). Following each convolutional layer is a max-pooling layer, with a pooling size and stride of 5 each. Our CNN architecture has 99.8 thousand trainable parameters for each bit-wise CNN. For each bit-wise CNN inference, 643.6 thousand

Testing setup

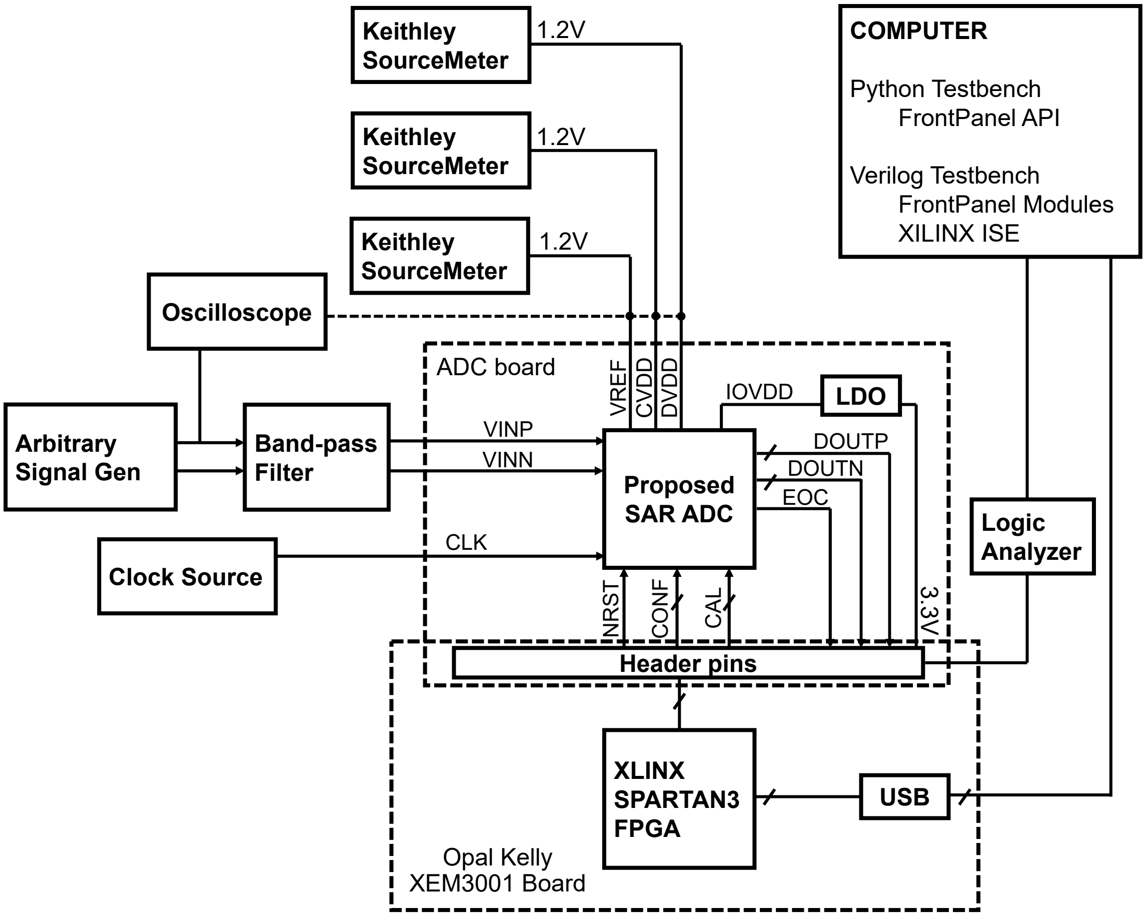

Figure 4-5: Testbench Diagram

of floating-point operations are needed.

4.6 Measurement Results

The testbench diagram is shown in Figure 4-5 and the measurement setup is shown in Figure 4-6. An FPGA board (Opal Kelly XEM6001) was used to communicate with the ADC test chip. This involves sending the NRST, CONF and CAL signals to reset the chip, configure the ADC mode, send calibration codes, and receive the output codes. 3 SourceMeters (Keithley 2400) were used to monitor the chip's core power consumption. A decoupling capacitor of 0.1uF was used on each supply node. An arbitrary waveform generator (Tektronix AFG3102) is used to generate the ADC

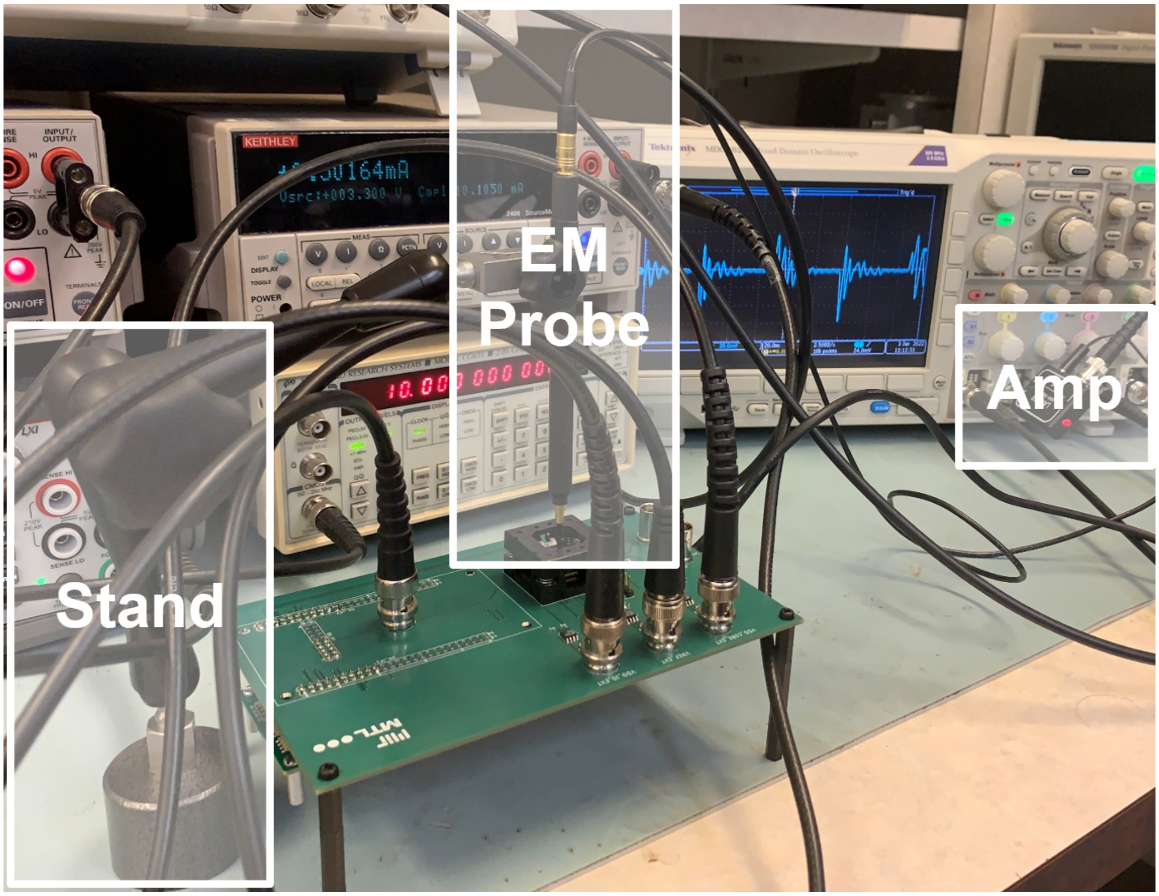

Figure 4-6: Photo of the setup

input signal and band-pass-filtered (from KR Electronics) for harmonics reduction. Then, the filtered signal is fed into a balun (Mini-Circuits ADTT1-6+) for single-ended to differential conversion. A low-jitter clock source (SRS CG635) supplies the SAR ADC clock signal. An oscilloscope with an integrated logic analyzer (Tektronix MDO3024) captures the ADC and the power traces for PSA. A QFN60 socket (3M 260-4204-01) is employed to facilitate the chip measurement process. On-board LDOs (Linear Technology LT3021) supply power to the SAR ADC. A stand is used to hold the EM Probe (LANGER). The signal from the probe is amplified and recorded by the aforementioned oscilloscope.

Fabricated in the 65nm LP process, the RaM-SAR takes 0.072mm2 (Figure. 4-7). The fabricated chips are wire-bonded in 60pin QFN package. The prototype demonstrates significant improvements with 12.5× higher bandwidth and 4.8× better

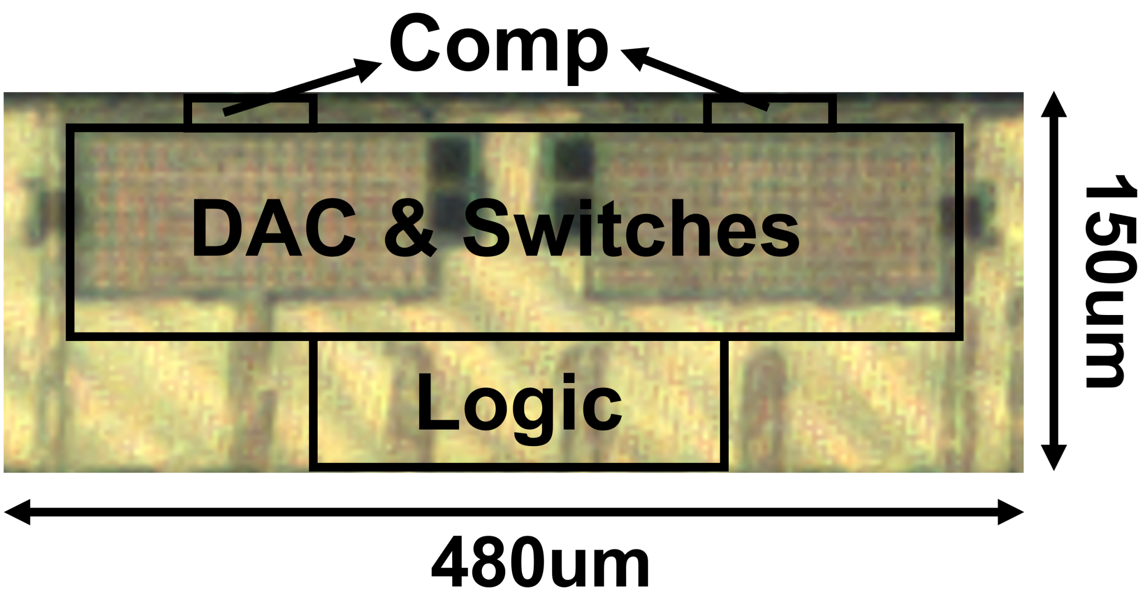

Figure 4-7: Chip micrograph

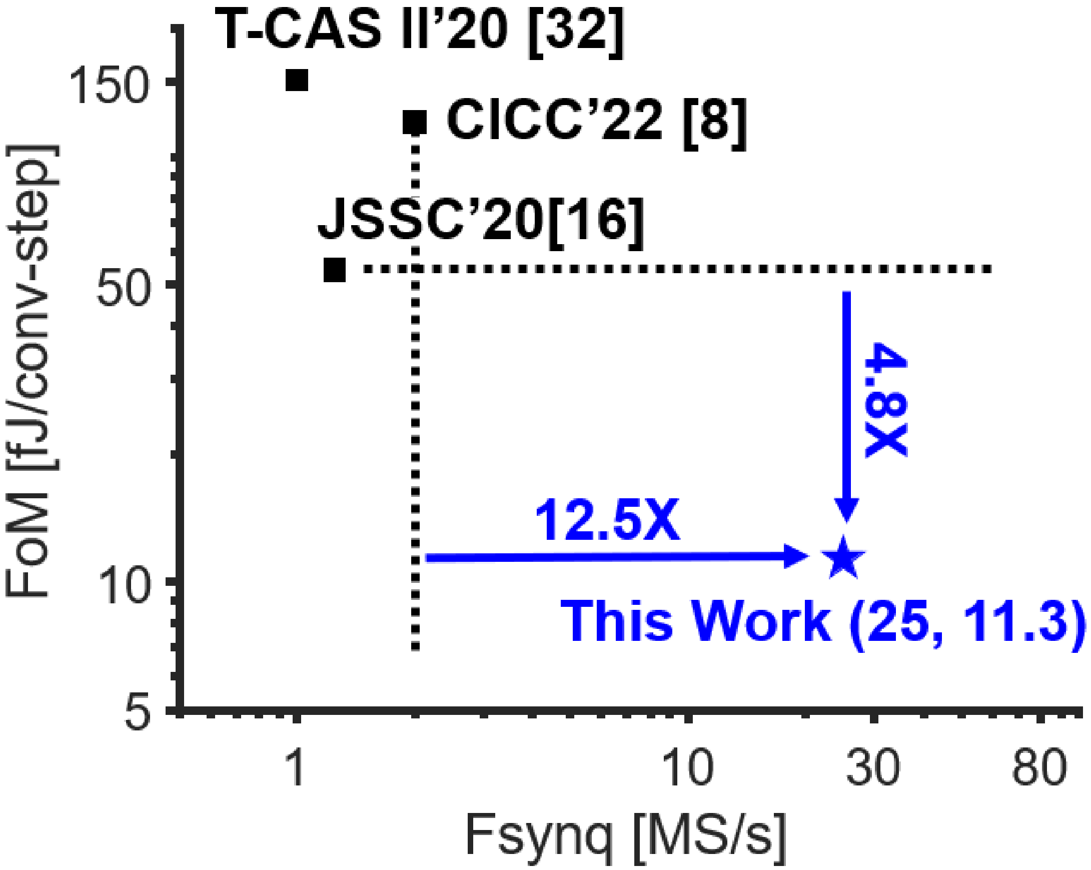

Figure 4-8: FoM vs Speed

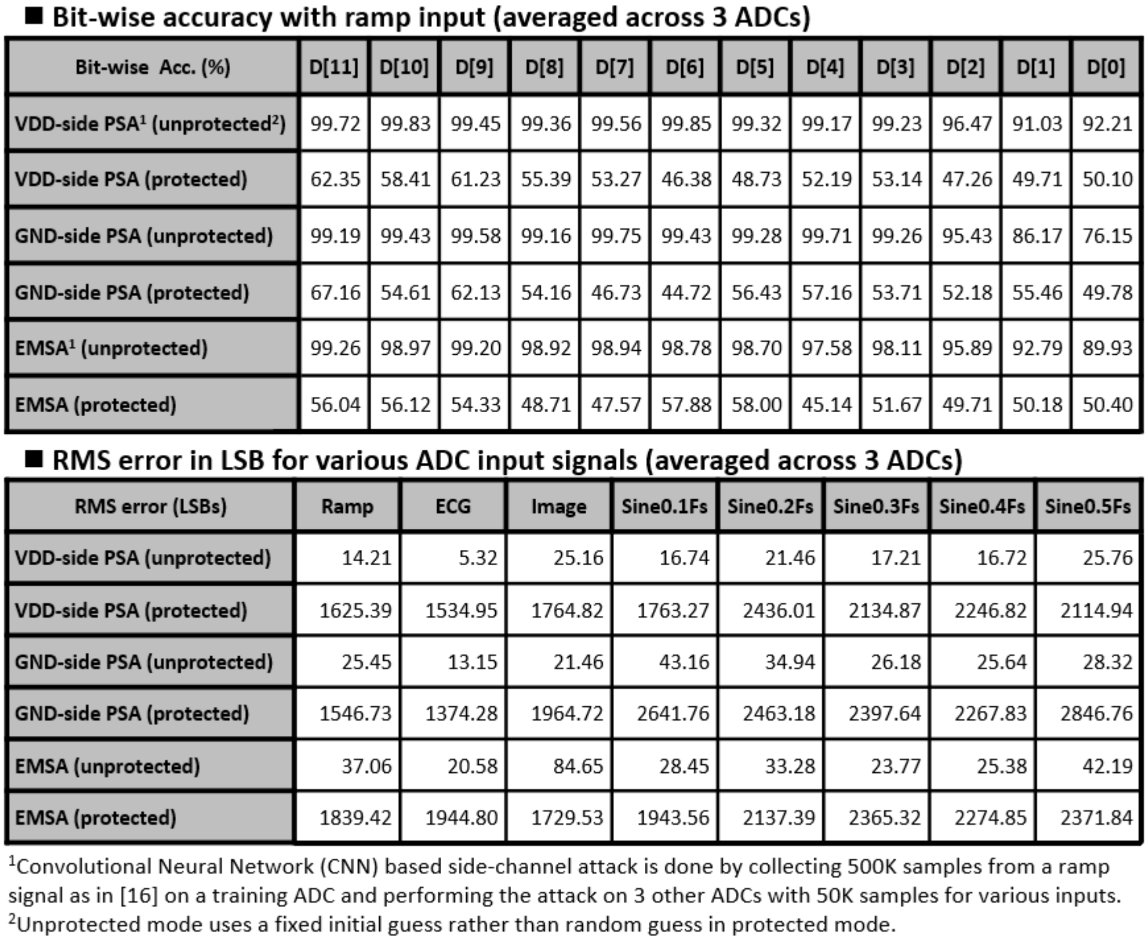

■ Bit-wise accuracy with ramp input (averaged across 3 ADCs)

Bit-wise Acc. (%)	D[11]	D[10]	D[9]	D[8]	D[7]	D[6]	D[5]	D[4]	D[3]	D[2]	D[1]	D[0]
VDD-side PSA[1] (unprotected[2])	99.72	99.83	99.45	99.36	99.56	99.85	99.32	99.17	99.23	96.47	91.03	92.21
VDD-side PSA (protected)	62.35	58.41	61.23	55.39	53.27	46.38	48.73	52.19	53.14	47.26	49.71	50.10
GND-side PSA (unprotected)	99.19	99.43	99.58	99.16	99.75	99.43	99.28	99.71	99.26	95.43	86.17	76.15
GND-side PSA (protected)	67.16	54.61	62.13	54.16	46.73	44.72	56.43	57.16	53.71	52.18	55.46	49.78
EMSA[1] (unprotected)	99.26	98.97	99.20	98.92	98.94	98.78	98.70	97.58	98.11	95.89	92.79	89.93
EMSA (protected)	56.04	56.12	54.33	48.71	47.57	57.88	58.00	45.14	51.67	49.71	50.18	50.40

■ RMS error in LSB for various ADC input signals (averaged across 3 ADCs)

RMS error (LSBs)	Ramp	ECG	Image	Sine0.1Fs	Sine0.2Fs	Sine0.3Fs	Sine0.4Fs	Sine0.5Fs
VDD-side PSA (unprotected)	14.21	5.32	25.16	16.74	21.46	17.21	16.72	25.76
VDD-side PSA (protected)	1625.39	1534.95	1764.82	1763.27	2436.01	2134.87	2246.82	2114.94
GND-side PSA (unprotected)	25.45	13.15	21.46	43.16	34.94	26.18	25.64	28.32
GND-side PSA (protected)	1546.73	1374.28	1964.72	2641.76	2463.18	2397.64	2267.83	2846.76
EMSA (unprotected)	37.06	20.58	84.65	28.45	33.28	23.77	25.38	42.19
EMSA (protected)	1839.42	1944.80	1729.53	1943.56	2137.39	2365.32	2274.85	2371.84

[1] Convolutional Neural Network (CNN) based side-channel attack is done by collecting 500K samples from a ramp signal as in [16] on a training ADC and performing the attack on 3 other ADCs with 50K samples for various inputs.

[2] Unprotected mode uses a fixed initial guess rather than random guess in protected mode.

Figure 4-9: Power/EM attacks of unprotected vs. protected mode of the RaM-SAR ADC

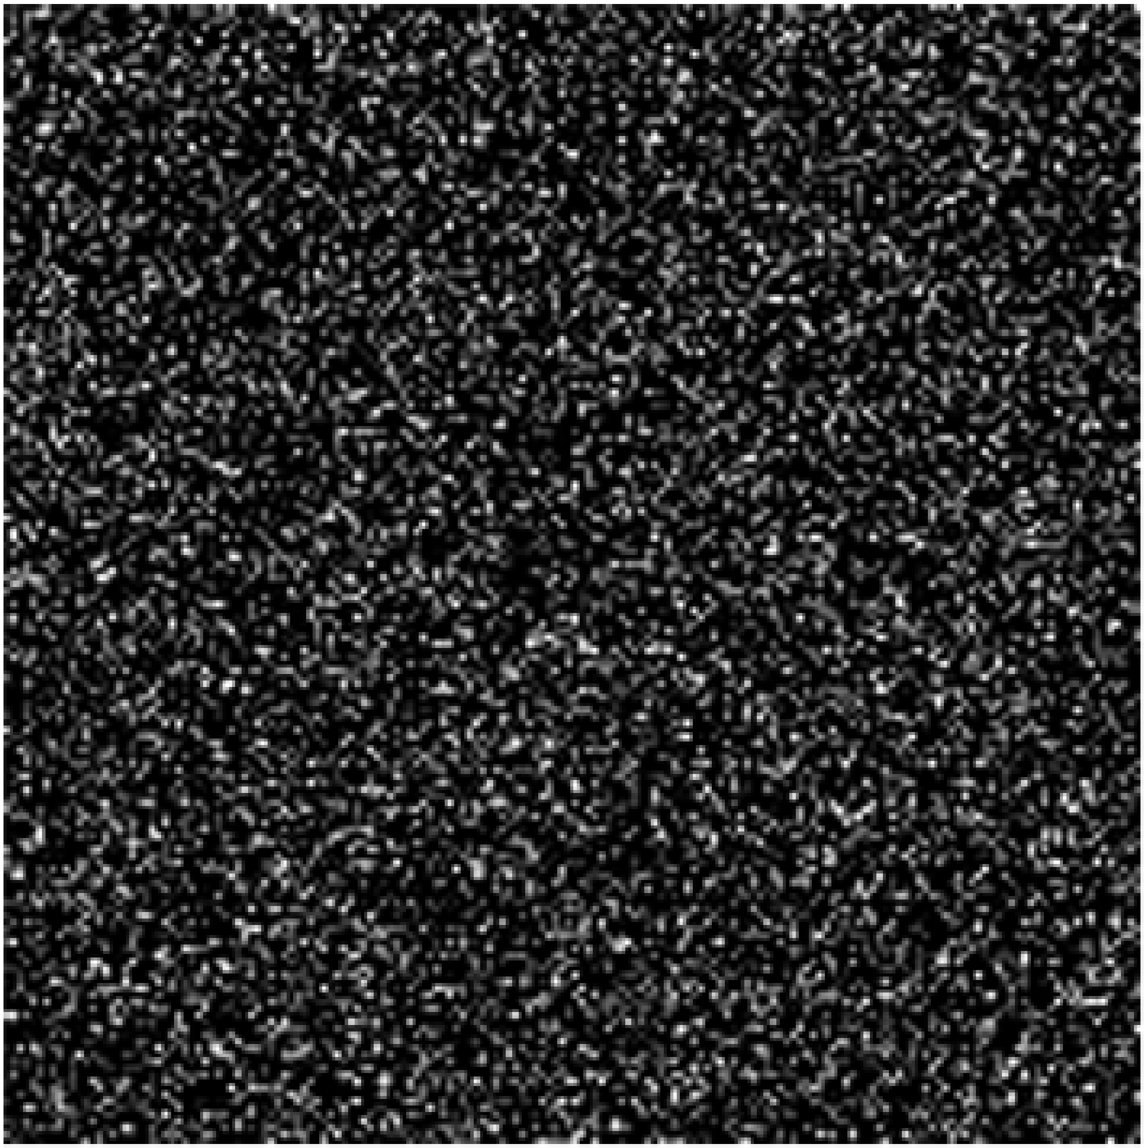

Figure 4-10: Example image of PSA on the protected SAR ADC

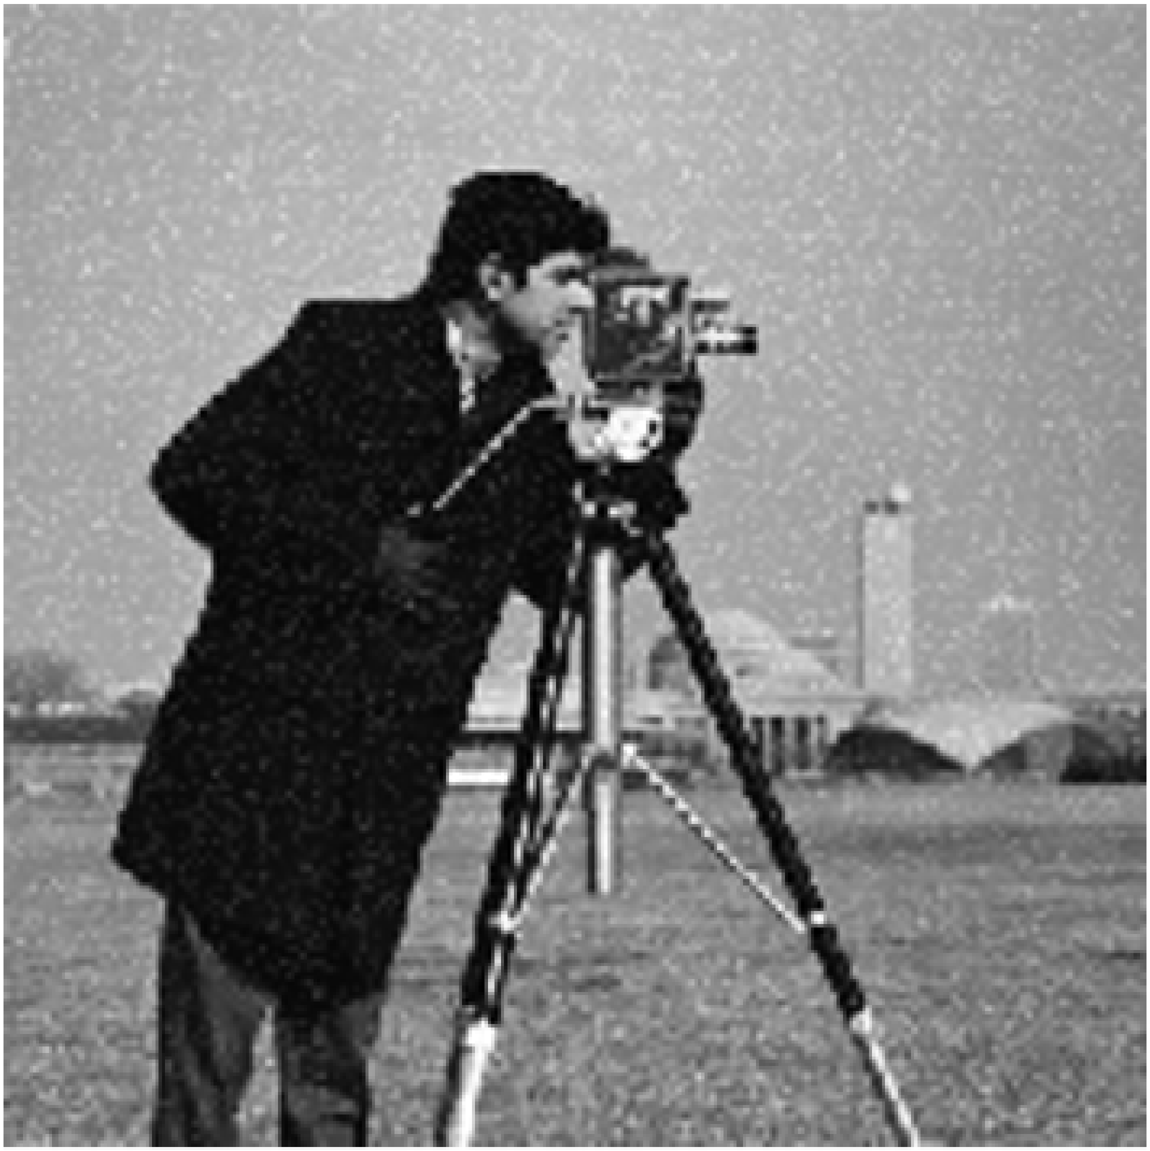

Figure 4-11: Example image of PSA on the unprotected SAR ADC

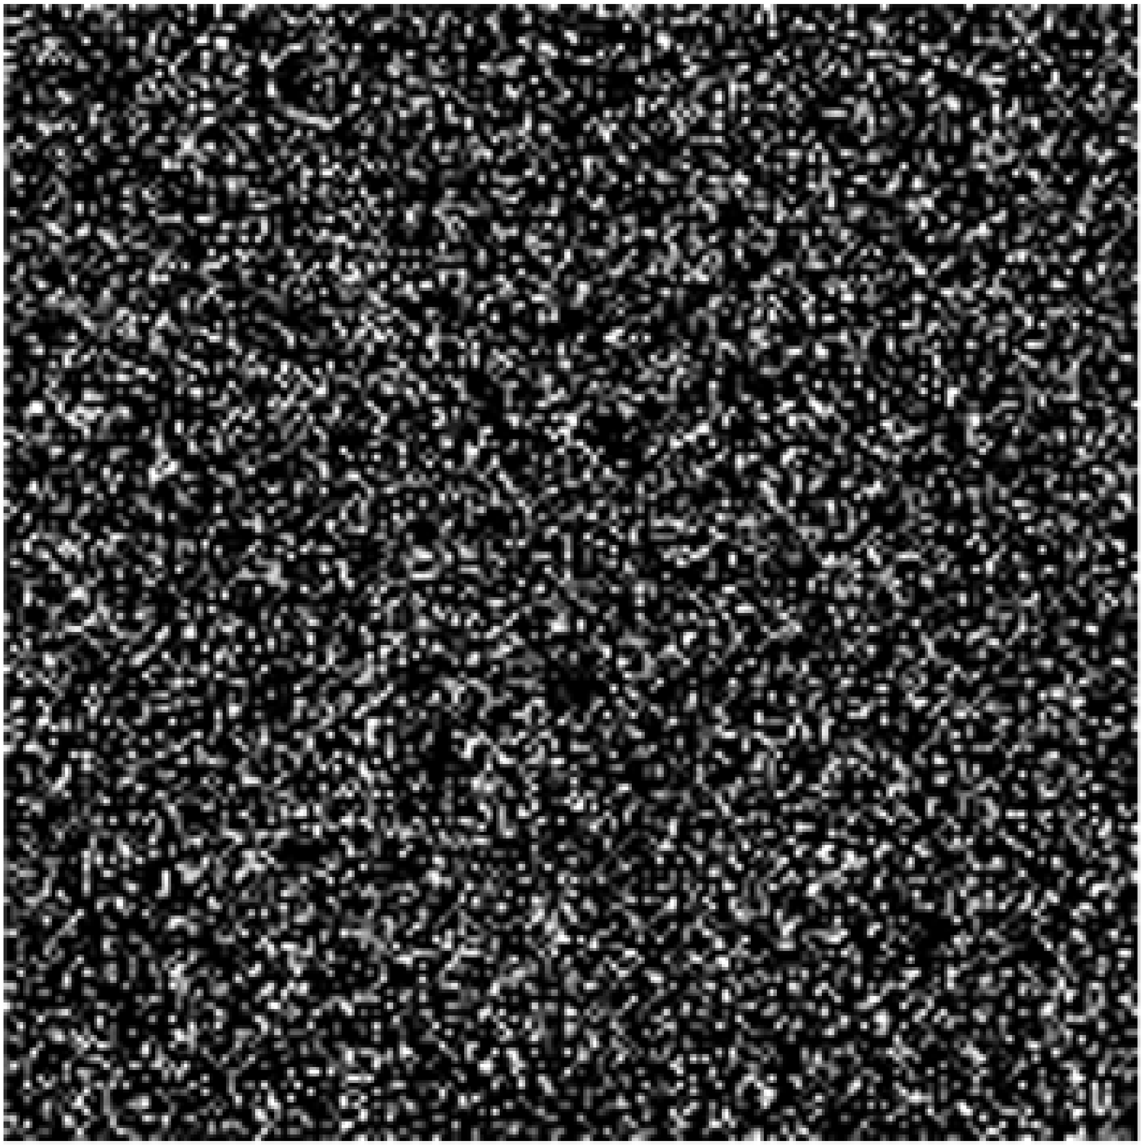

Figure 4-12: Example image of EMSA on the protected SAR ADC

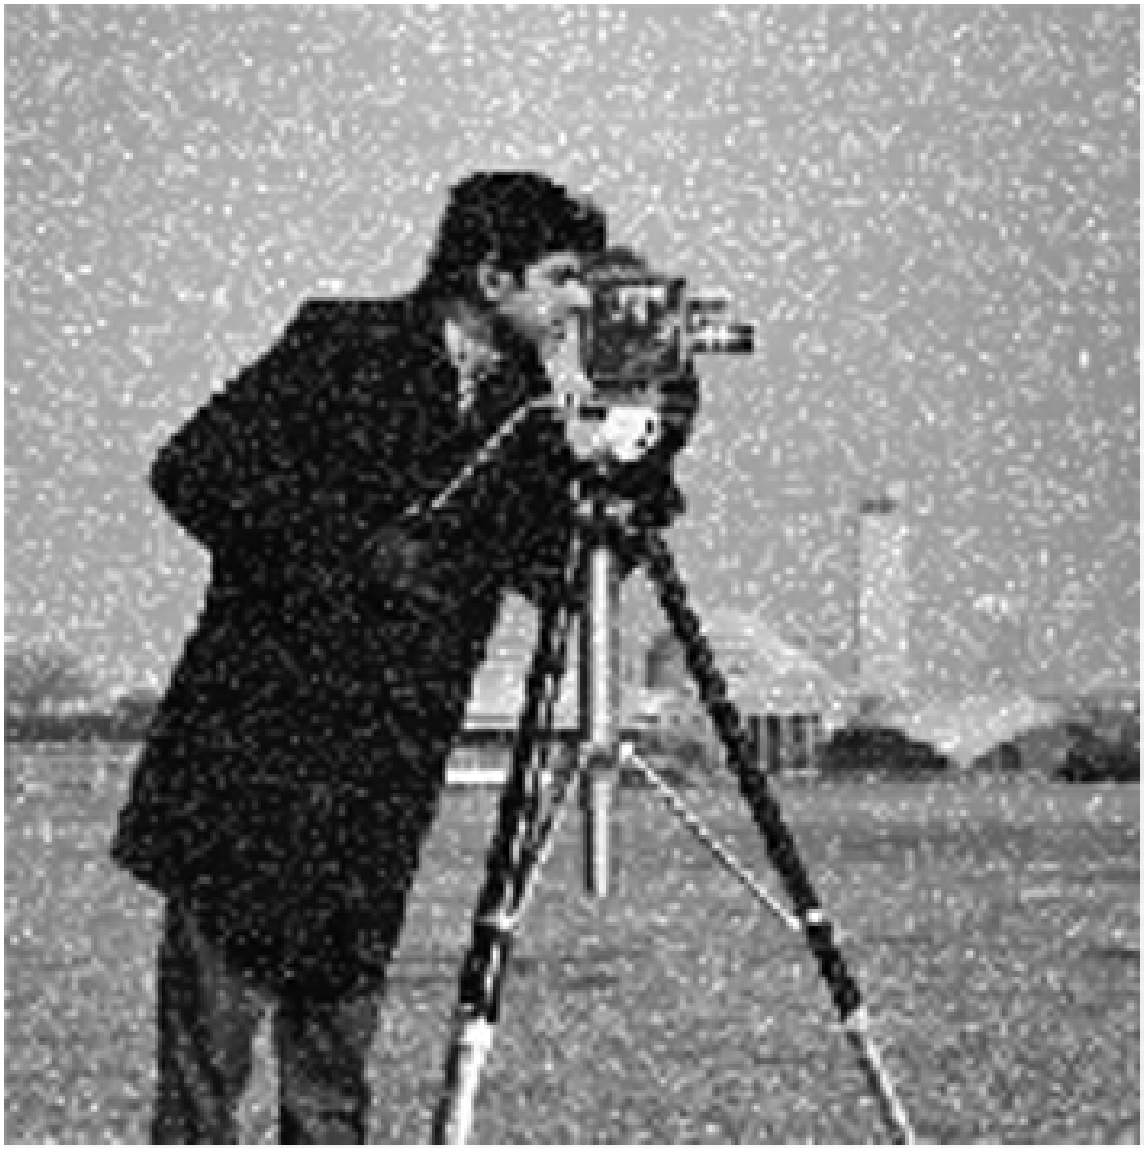

Figure 4-13: Example image of EMSA on the unprotected SAR ADC

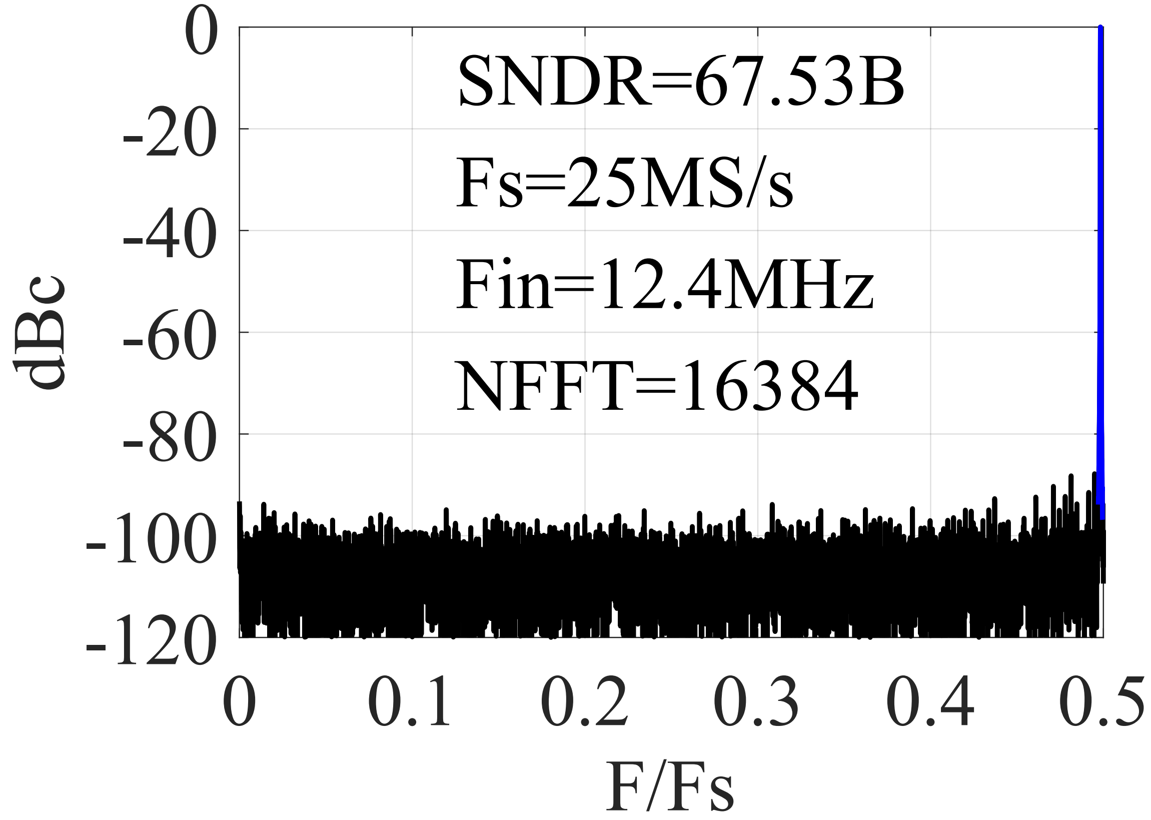

Figure 4-14: Measured FFT plot of RaM-SAR

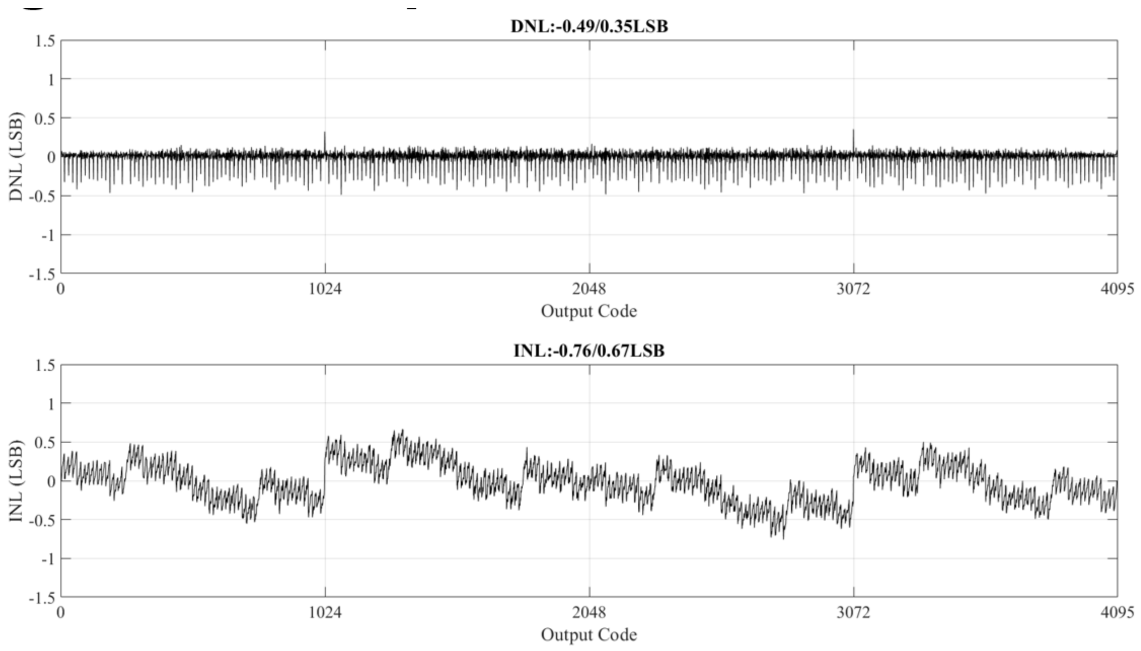

Figure 4-15: Measured DNL/INL of the unprotected mode SAR ADC

energy efficiency over prior secure ADCs (Figure. 4-8). The sampling rate of RaM-SAR is 25MS/s and the FoM is 11.3 fJ/c.-s. The sampling rate of [32] is 1.25MS/s and the FoM is 150 fJ/c.-s The sampling rate of [8] is 2MS/s and the FoM is 120.7 fJ/c.-s. The sampling rate of [16] is 1.25MS/s and the FoM is 54.3 fJ/c.-s.

In our experiment, we performed Convolutional Neural Network (CNN) based Power Side-Channel Attacks (PSA) and Electro-Magnetic Side-Channel Attacks (EMSA) against both unprotected and protected Analog-to-Digital Converters (ADCs), as illustrated in Figure 4-9. The CNN-based side-channel attacks involve the collection of 500,000 samples from a ramp signal, as demonstrated in [16], on a training ADC. These samples are then used to perform the attack on three other ADCs, each with 50,000 samples collected for various input signals.

The unprotected mode relies on a fixed initial guess, in contrast to the random initial guess employed in the protected mode. To conduct the VDD-side PSA, we measured the current profile of the VDD pin by connecting it with a 30ohm series resistor. Similarly, for the GND-side PSA, the current profile of the GND pin was collected using a 30ohm series resistor. For the EMSA, the electromagnetic profile of the chip was gathered using an EM probe placed near the ADC.

Different CNN models were utilized for separate bits in the ADC. The bit-wise accuracy with a ramp input was calculated as an average across the three tested ADCs. To evaluate the Root Mean Square (RMS) error, the error was normalized with respect to the full scale of a 12-bit ADC.

We conducted tests on various input signals, including ramp, Electrocardiogram (ECG), image, and sine waves. These diverse inputs allowed us to thoroughly evaluate the performance and security of both unprotected and protected ADCs under different conditions. Our experiment aimed to demonstrate the effectiveness of the protection mechanisms implemented in the secure ADC, as well as identify any potential weaknesses or vulnerabilities that could be exploited by adversaries using CNN-based side-channel attacks.

By comparing the outcomes of these tests, we aimed to gain insights into the relative security and performance of unprotected and protected ADCs. This information can help guide future design improvements and contribute to the development of more secure and efficient ADCs for various applications. The results of our experiment demonstrate the effectiveness of the secure ADC in preventing CNN-based PSA and EMSA, while still maintaining high performance and energy efficiency.

Figure. 4-9 demonstrates the effectiveness of our protection scheme in safeguarding the ADC against VDD-side PSA, GND-side PSA, and EMSA. In evaluating the bit-wise accuracy of the ADCs, a 100% accuracy indicates that the attack algorithm accurately guesses the digital bit 100% of the time, meaning that the ADC is not secure. Conversely, a 50% accuracy implies that the attack algorithm performs at the same level as a random selector, signifying that the ADC is secure.

D11 represents the Most Significant Bit (MSB), while D0 denotes the Least Significant Bit (LSB). When assessing the unprotected ADC under VDD-side PSA, we observed that the bit accuracy exceeded 99% for bits D3 through D11 and surpassed 90% for bits D0 to D2. This outcome reveals that the unprotected ADC is not secure under VDD-side PSA.

However, for the protected ADC under VDD-side PSA, the bit accuracy hovered around 50% for all bits except the MSB. This result indicates that the protected ADC is secure under VDD-side PSA. In the case of the unprotected ADC under GND-side PSA, the bit accuracy approached 100% for the majority of bits, suggesting that the unprotected ADC is not secure under GND-side PSA.

On the other hand, for the protected ADC under GND-side PSA, the bit accuracy remained around 50% for all bits except the MSB, confirming that the protected ADC is secure under GND-side PSA. For the unprotected ADC subjected to EMSA, the bit accuracy was nearly 100% for most bits, indicating that the unprotected ADC is not secure under EMSA.

In contrast, for the protected ADC under EMSA, the bit accuracy was approximately 50% for all bits, including the MSB. This finding demonstrates that the protected ADC is secure under EMSA. The results of our study highlight the importance of incorporating a robust protection scheme into ADC designs to ensure their security against various side-channel attacks.

As part of the evaluation process, an example image is fed into the ADCs to assess the efficacy of the protection mechanism in preserving the confidentiality of the input data. This approach allows us to evaluate the performance of the ADCs under realistic conditions and determine their susceptibility to side-channel attacks.

After conducting the EMSA on the unprotected ADC, the attacker can clearly discern the image's content, featuring a man with a camera, as shown in Figure 4-13. However, when examining the protected ADC, the attacker can only observe random noise, as depicted in Figure 4-10. In contrast, with the protection enabled, the EMSA result appears to be random and does not disclose any useful information about the original image, as demonstrated in Figure 4-12.

Further evaluation of the ADC's performance reveals a Spurious-Free Dynamic Range (SFDR) of 86.6dB, as shown in Figure 4-14. The Differential Non-Linearity (DNL) measures -0.49/+0.35LSB, while the Integral Non-Linearity (INL) is -0.76/+0.67LSB 4-15.

The ADC achieves a FoM of 11.3fJ/c.-s, which is comparable with the state-of-the-art energy-efficient non-secure ADCs [62][63], with both PSA and EMSA resilience (Figure. 4-16). The energy efficiency of the secure ADC can still be improved.

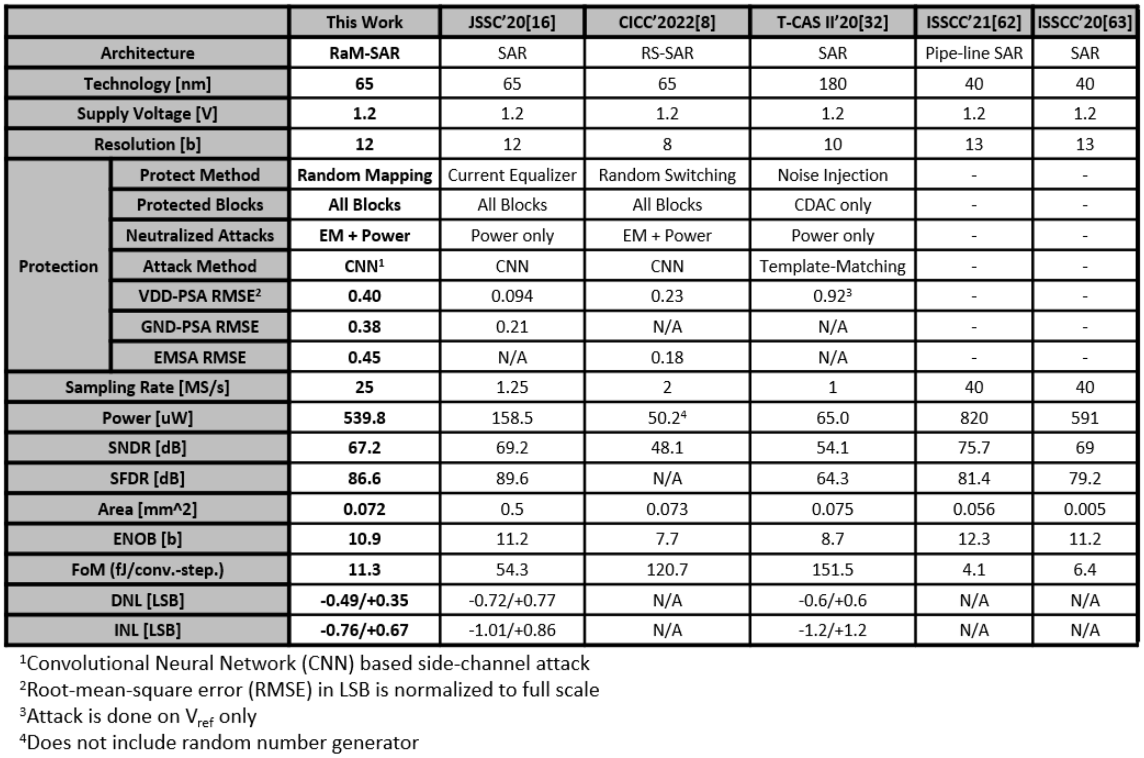

		This Work	JSSC'20[16]	CICC'2022[8]	T-CAS II'20[32]	ISSCC'21[62]	ISSCC'20[63]
Architecture		**RaM-SAR**	SAR	RS-SAR	SAR	Pipe-line SAR	SAR
Technology [nm]		**65**	65	65	180	40	40
Supply Voltage [V]		**1.2**	1.2	1.2	1.2	1.2	1.2
Resolution [b]		**12**	12	8	10	13	13
Protection	Protect Method	**Random Mapping**	Current Equalizer	Random Switching	Noise Injection	-	-
	Protected Blocks	**All Blocks**	All Blocks	All Blocks	CDAC only	-	-
	Neutralized Attacks	**EM + Power**	Power only	EM + Power	Power only	-	-
	Attack Method	**CNN[1]**	CNN	CNN	Template-Matching	-	-
	VDD-PSA RMSE[2]	**0.40**	0.094	0.23	0.92[3]	-	-
	GND-PSA RMSE	**0.38**	0.21	N/A	N/A	-	-
	EMSA RMSE	**0.45**	N/A	0.18	N/A	-	-
Sampling Rate [MS/s]		**25**	1.25	2	1	40	40
Power [uW]		**539.8**	158.5	50.2[4]	65.0	820	591
SNDR [dB]		**67.2**	69.2	48.1	54.1	75.7	69
SFDR [dB]		**86.6**	89.6	N/A	64.3	81.4	79.2
Area [mm^2]		**0.072**	0.5	0.073	0.075	0.056	0.005
ENOB [b]		**10.9**	11.2	7.7	8.7	12.3	11.2
FoM (fJ/conv.-step.)		**11.3**	54.3	120.7	151.5	4.1	6.4
DNL [LSB]		**-0.49/+0.35**	-0.72/+0.77	N/A	-0.6/+0.6	N/A	N/A
INL [LSB]		**-0.76/+0.67**	-1.01/+0.86	N/A	-1.2/+1.2	N/A	N/A

[1] Convolutional Neural Network (CNN) based side-channel attack
[2] Root-mean-square error (RMSE) in LSB is normalized to full scale
[3] Attack is done on V_{ref} only
[4] Does not include random number generator

Figure 4-16: Comparison Table

Chapter 5

Sniff-SAR: Detection-driven and un-trainable secure ADC

5.1 Introduction

One of the main challenges in designing secure ADCs is balancing the trade-off between security and energy efficiency. Secure ADCs often have non-negligible energy and area overheads, which is not ideal for resource-constrained IoT applications. The protection schemes in secure ADCs are typically always-on, even when side-channel attacks are not being performed.

Another challenge is the increasing power of neural network-based side-channel attacks, which render existing protection mechanisms less robust. To address these challenges, this work proposes the first detection-driven secure ADC that protects against both power and EM side-channel attacks. The ADC operates in an energy-efficient switching mode under normal conditions. When an EMSA or PSA is detected, a secure switching mechanism is enabled, rendering the ADC practically untrainable by neural network-based attacks.

The proposed detection-driven secure ADC [64] offers several advantages. First, it reduces the energy and area overheads associated with conventional secure ADCs by only activating the secure switching mechanism when an attack is detected. This makes it more suitable for resource-constrained IoT applications. Second, the secure

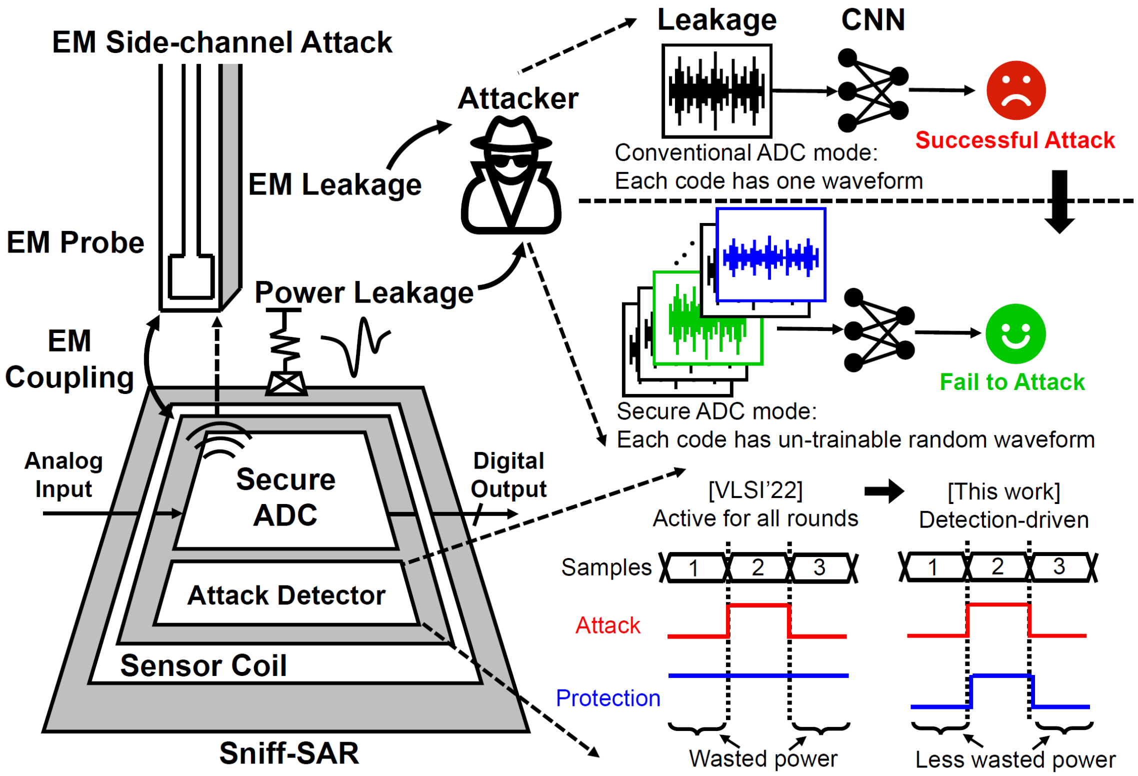

Figure 5-1: Side-channel security challenges of ADCs and detection-driven protection based on randomization

switching mechanism is designed to be resistant to neural network-based side-channel attacks, ensuring a higher level of security.

To implement the detection-driven secure ADC, several design considerations must be taken into account. The detection mechanism should be able to accurately identify both EMSA and PSA with minimal false positives and negatives. Additionally, the secure switching mechanism must be carefully designed to prevent information leakage while maintaining the ADC's performance.

The detection-driven secure ADC offers a promising solution for addressing the challenges associated with securing ADCs in resource-constrained IoT applications. By operating in an energy-efficient switching mode and activating the secure switching mechanism only when an attack is detected, the proposed ADC reduces energy and area overheads while providing robust protection against both power and EM side-channel attacks. Moreover, by leveraging machine learning techniques for attack detection and designing secure switching mechanisms that are resistant to neural network-based attacks, the detection-driven secure ADC offers a higher level of security and robustness for IoT applications.

5.2 Previous Work

Recently, detection-driven techniques have been proposed to reduce the overhead of side-channel attack resiliency. [65] introduces EQZ-LDO, an innovative digital low drop-out regulator (LDO) with equalizer designed to provide resilience against side-channel attacks (SCAs). A key feature of the proposed solution is its attack detection capability, combined with a detection-driven protection mechanism. This approach ensures that the protection system is activated only when an SCA is detected. As a result, the energy-delay-product (EDP) overhead is minimized, amounting to a mere 0.5% increase.

In the context of Internet of Things (IoT) devices, the EQZ-LDO offers an efficient and effective means of enhancing security. By spreading the EDP overhead across the device's lifetime, it enables a more sustainable approach to resource allocation.

In addition, the detection-driven protection scheme ensures that performance is not compromised during periods when no attacks are occurring.

The EQZ-LDO digital low drop-out regulator offers a powerful solution for SCA resilience in IoT devices. Its unique combination of attack detection and detection-driven protection allows it to provide strong security without significant EDP overhead. This approach ensures that both performance and security are effectively balanced, making it an attractive option for IoT device developers and manufacturers seeking to protect their products from side-channel attacks. Secure ADC can take a similar concept to reduce the overhead of the security feature. However, the EQZ-LDO cannot protect against EMSA.

EMSA is a more practical approach to capturing the side-channel information of the ADCs. [66] presents the development of a cryptographic engine (CE) designed to resist local electromagnetic analysis attacks (L-EMAs). The core of this innovative solution is an LC-oscillator-based tamper-access sensor that detects the approach of a micro EM probe, ensuring the protection of secret key information from potential attacks.

The fully-digital sensor circuit features a reference-free dual-coil sensing scheme and a ring-oscillator-based one-step digital sensor calibration. These design elements contribute to a reduced sensor area overhead of just 1.6%. Furthermore, the sensor is designed to operate intermittently, interleaving between CE operations. This approach allows for power savings and minimizes performance penalties, with power consumption reduced by 7.6% and performance penalty limited to 0.2%.

The development of a cryptographic engine resistant to local EM-analysis attacks represents a significant step forward in the field of cybersecurity. The use of an LC-oscillator-based tamper-access sensor, combined with a fully-digital sensor circuit and an intermittent operation mode, ensures that secret key information remains secure while minimizing overhead and performance penalties. The successful demonstration of L-EMA attack detection and key protection in a prototype highlights the potential of this approach in securing cryptographic systems against increasingly sophisticated attacks.

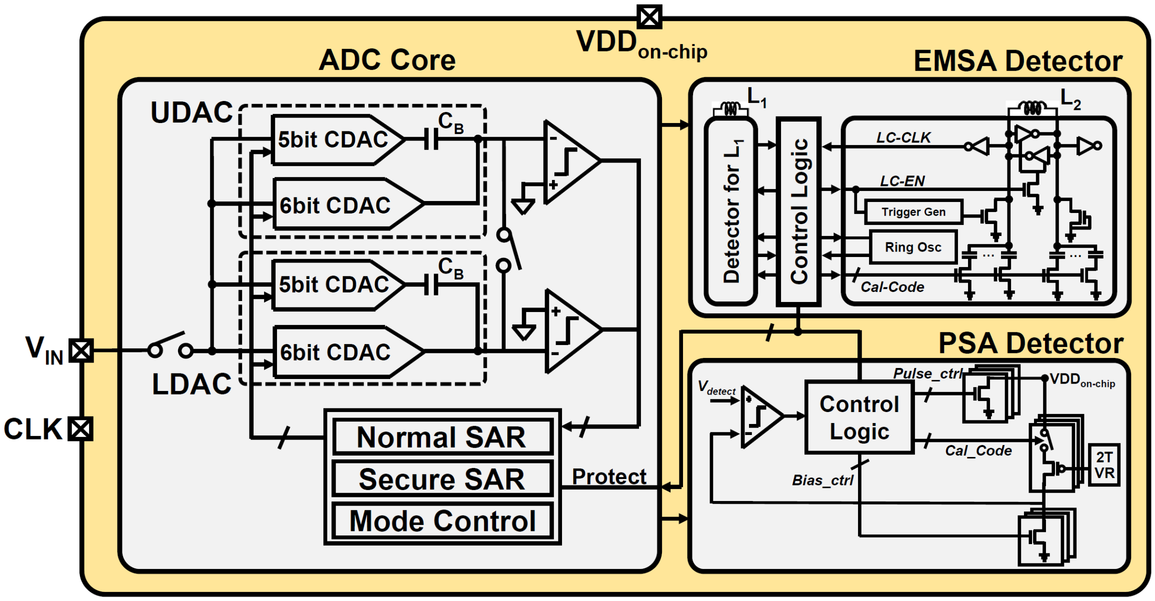

Figure 5-2: System architecture of the side-channel-secure ADC with detection-driven protection

The proposed work integrates the PSA [65] and EMSA [66] detectors into a SAR ADC (Sniff-SAR) with minor modifications. The Sniff-SAR operates in an energy-efficient mode and provides secures conversions only when PSA or EMSA is detected, which reduces the overhead of secure ADCs.

5.3 Contributions

This research proposes a novel ADC design, known as Sniff-SAR, that enhances security in resource-constrained IoT applications by balancing energy efficiency and protection against side-channel attacks. The Sniff-SAR uses a detection-driven approach to combat electromagnetic side-channel attacks (EMSA) and power side-channel attacks (PSA). Under normal conditions, the ADC operates in an energy-efficient mode, but when a threat is detected, it switches to a secure mode, reducing energy and area overheads associated with conventional secure ADCs.

The EMSA and PSA detectors within the ADC core monitor for side-channel attacks and activate the secure mode when necessary, providing robust protection. The ADC's secure mode employs a unique conversion plan that leverages a random number generator, making it resistant to side-channel attacks and impractical for training neural network-based attacks.

The Sniff-SAR builds upon previous work, integrating PSA and EMSA detectors into a SAR ADC with minor modifications. The design also maintains high performance and energy efficiency by operating in an energy-efficient mode under normal conditions and only activating the secure mode when an attack is detected.

The resulting prototype can detect a 30-ohm series resistor for PSA and an EM probe at a distance of 0.16mm for EMSA, demonstrating its robustness against security threats. This detection-driven secure ADC design provides a promising solution for secure, energy-efficient ADCs in IoT applications.

5.4 Proposed Sniff-SAR

5.4.1 Architecture

Figure. 5-2 shows the system architecture of the Sniff-SAR with detection-driven protection. The EMSA and PSA detectors capture the attempt of side-channel attacks. The ADC core performs the analog-to-digital conversion. The ADC core normally operates in the unprotected SAR mode which is faster and more energy efficient. EMSA and PSA detectors check for side-channel attacks periodically and once they notice that the ADC is under attack, the ADC activates the secure SAR mode against both EMSA and PSA.

Figure 5-2 presents the system architecture of the Sniff-SAR, featuring detection-driven protection. This architecture consists of two primary components: the attack detectors and the ADC core.

The EMSA and PSA detectors are responsible for monitoring and detecting attempts of side-channel attacks. These detectors periodically check for any signs of Electro-Magnetic Side-Channel Attacks (EMSA) or Power Side-Channel Attacks (PSA), which could threaten the security of the ADC.

The ADC core is responsible for performing the A/D conversion process. Under normal operating conditions, the ADC core functions in the unprotected SAR mode, which offers higher sampling speed and greater energy efficiency compared to the secure SAR mode. This allows the Sniff-SAR to conserve resources and maintain optimal performance when not under attack.

When the EMSA and PSA detectors identify a side-channel attack attempt, the ADC activates the secure SAR mode to counter both EMSA and PSA threats. This mode introduces additional security measures and countermeasures to protect sensitive information and maintain the integrity of the ADC's operation.

By implementing this detection-driven protection, the Sniff-SAR achieves a balance between performance and security. The ADC can operate efficiently in the unprotected SAR mode when no attacks are detected, reducing energy consumption and ensuring optimal performance. However, when an attack is detected, the ADC can swiftly transition to the secure SAR mode, providing robust protection against both EMSA and PSA without compromising the ADC's overall performance.

The Sniff-SAR's detection-driven protection offers a flexible and efficient solution for safeguarding ADCs in resource-constrained IoT applications. By operating in an energy-efficient mode and activating secure SAR mode only when necessary, the Sniff-SAR can effectively protect against both EMSA and PSA while maintaining high performance and energy efficiency. The core circuitry including CDAC, comparator, and sampling networks is identical to the RaM-SAR.

5.4.2 Attack Detectors

To detect electromagnetic side-channel attacks (EMSA), the ADC core employs dual sensor coils, L1 and L2, placed over the core to form two LC oscillators [66]. When an EM probe approaches the ADC, the oscillation frequencies of the two LC oscillators diverge. Digital counters in the control logic detect this divergence and activate the secure successive approximation register (SAR) mode in the ADC, offering enhanced protection against EMSA.

The sensor coils are implemented in two distinct metal layers, providing orthogonal edges to cover multiple attack vectors. Local ring oscillators serve as a process, voltage, and temperature (PVT) monitor to calibrate the LC oscillators, eliminating the need for an external clock reference, as shown in [66].

In addition to the EMSA detector, the design includes a power side-channel attack (PSA) detector. The PSA detector generates a current pulse to sense the IR drop across the non-negligible resistance employed for PSA detection. This IR drop is amplified by a common-gate, subthreshold-biased PMOS array [65]. The PMOS array is biased by an embedded 2T voltage reference [65].

The amplifier is designed to deliver a DC gain of 30dB and a unity-gain bandwidth of 350kHz, effectively suppressing mid-to-high frequency supply noise. When the amplified voltage falls below a predetermined Vdetect threshold, the PSA detector interprets it as a PSA attempt and activates the secure SAR mode in the ADC core, enhancing the overall security against power side-channel attacks.

The key features of the Sniff-SAR include an ADC core, dual sensor coils for EMSA detection, and a dedicated PSA detector. These components work together to enhance the security of the ADC core without incurring additional area or power overhead. The innovative approach demonstrated in this work has the potential to significantly improve the resilience of ADCs and other electronic systems against increasingly sophisticated side-channel attacks, ensuring the confidentiality and integrity of sensitive information in various applications.

5.4.3 Conversion Plan

The secure mode offers 3.6 x 10^{16} unique switching traces using a true random number generator, which makes it highly resistant to side-channel attacks, including PSA and EMSA. To put this into perspective, it would take approximately 2,900 years to collect all possible switching traces at 100 times each, with a sampling rate of 40 Mega Samples per second (MS/s). Consequently, training an ADC using neural networks to perform PSA or EMSA would be impractical in this case. Pseudorandom number generator is used to prove the concept.

The secure SAR ADC consists of three phases: random start, search, and LSB. The DU and DL registers correspond to the UDAC and LDAC, respectively. The

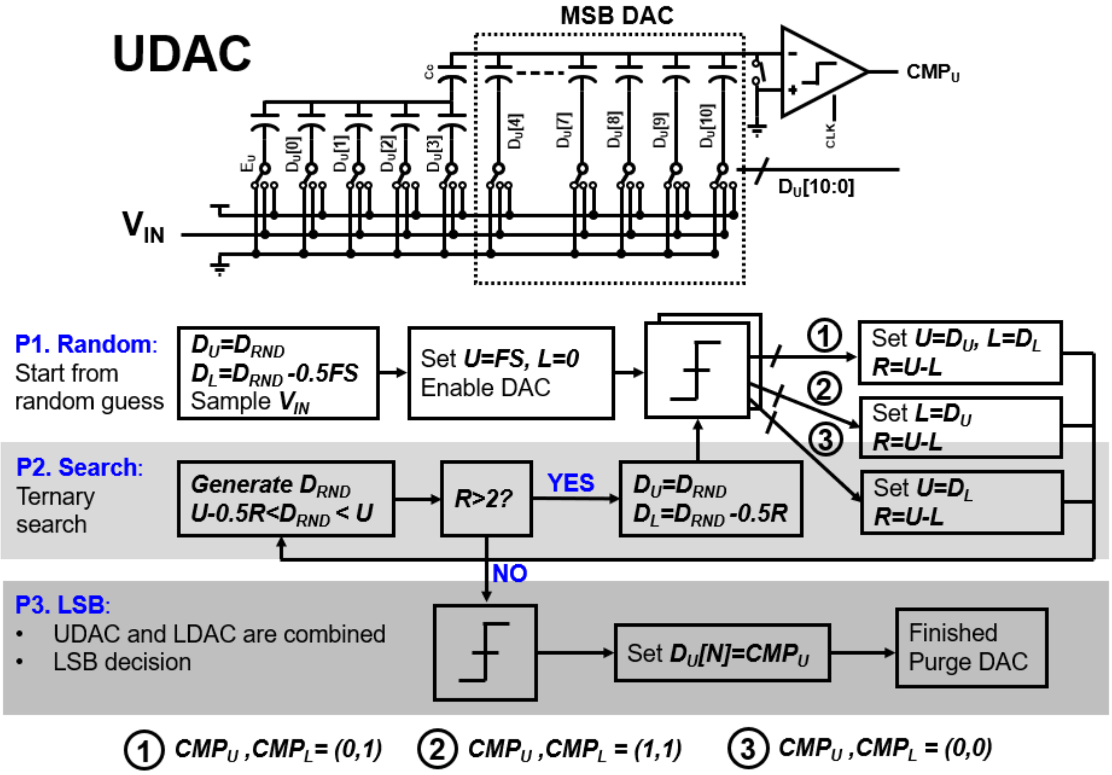

Figure 5-3: DAC schematic and flowchart of the secure SAR

random start phase begins with the DU register set to a pseudo-random number (DRND) between half of the Full Scale (FS) and the FS. The DL register is then set to half FS below DU. The upper bound (U) and the lower bound (L) of the Range of Uncertainty (R) are initially set to FS and 0, respectively.

CMPU and CMPL denote the outputs of the upper and lower comparators, respectively. The values of U, L, and R are updated according to CMPU and CMPL. If CMPU equals 0 and CMPL equals 1, the input voltage lies between DU and DL, so U and L are set to DU and DL, respectively. If CMPU and CMPL both equal 1, the input voltage is larger than DU, and L is updated to DU. If CMPU equals 0 and CMPL equals 0, the input voltage is smaller than DL, and U is updated to DL. The Range of Uncertainty (R) is updated accordingly.

In the search phase, DU is set to a pseudo-random number (DRND) larger than both the center of the Range of Uncertainty and U. The DL register is set to 0.5R below DU. If R falls below 2, the secure SAR enters the final LSB phase.

During the LSB phase, the UDAC and LDAC are combined to make a binary decision. It is worth noting that the unprotected SAR ADC utilizes ternary search, which allows for three-bit decisions in only two cycles. This process enhances the performance and efficiency of the ADC.

The prototype is capable of detecting a 30-ohm series resistor for PSA and an EM probe at a distance of 0.16mm. The maximum detectable probe distance is accurately measured using a 3D positioning stage. This design ensures that the detector can identify even subtle variations in the probe's position, significantly enhancing the security of the ADC against electromagnetic attacks.

The quiescent power consumption of the EMSA and PSA detectors is 21µW and 1.1µW, respectively, which demonstrates the energy-efficient operation of these detectors. Such low power consumption is crucial for the practical implementation of these security features in various applications, particularly in power-constrained environments such as Internet of Things (IoT) devices.

5.5 Measurement Setup

In this experiment, as depicted in Figure. 5-4, we used a comprehensive measurement setup to ensure the accuracy of the results. The foundation of the setup is an optical breadboard (B1824FX), which provides a stable platform for positioning the test PC board and the positioning stage. To enhance the precision of the positioning stage, an MT405 adapter is utilized.

A crucial component of this setup is the KDC101 motor controller, responsible for the accurate 3D positioning of the EM probe. This allows for precise control and movement of the probe to obtain reliable measurements. To facilitate communication between the PC and the motor controller, a KCH301 USB controller hub and power supply are employed.

Furthermore, the MT3-Z8 three-axis motorized translation stage plays a vital role in holding the EM probe securely in place. This ensures that the probe remains stable during the experiment, minimizing potential errors or inaccuracies.

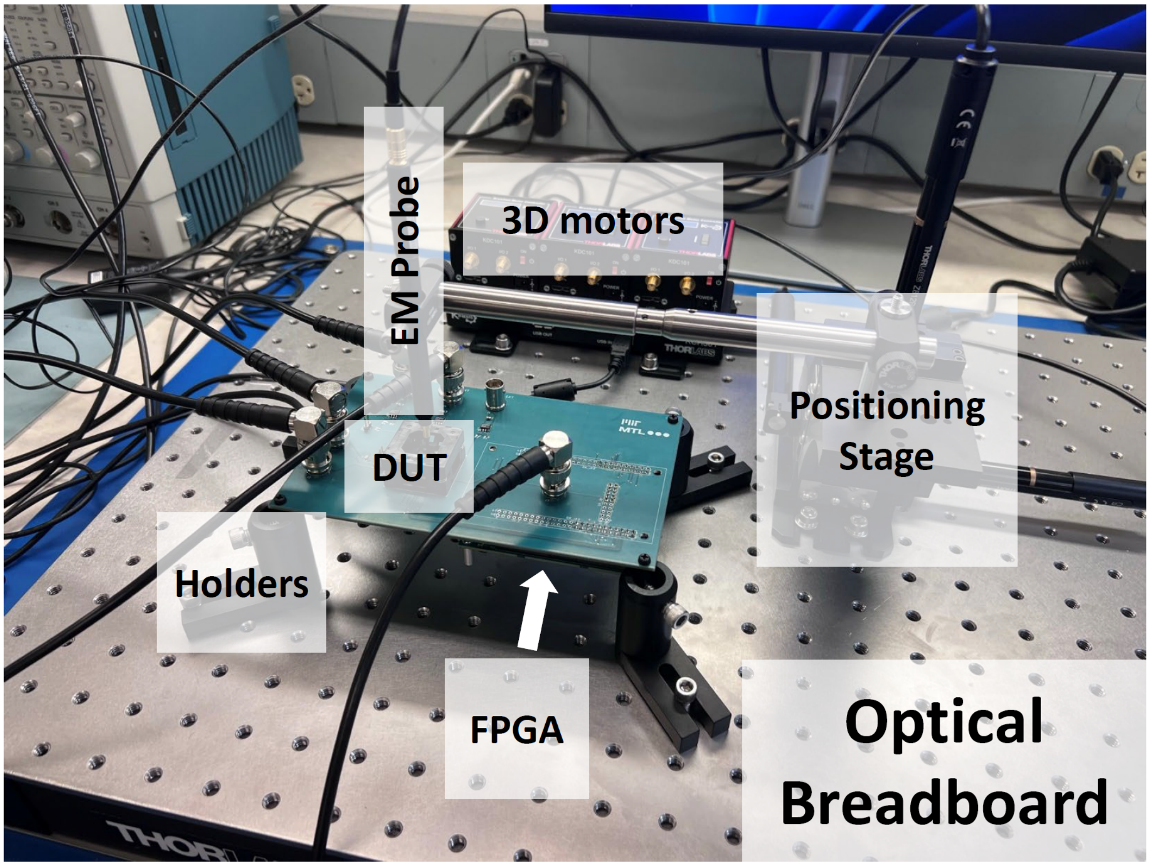

Figure 5-4: Measurement Setup

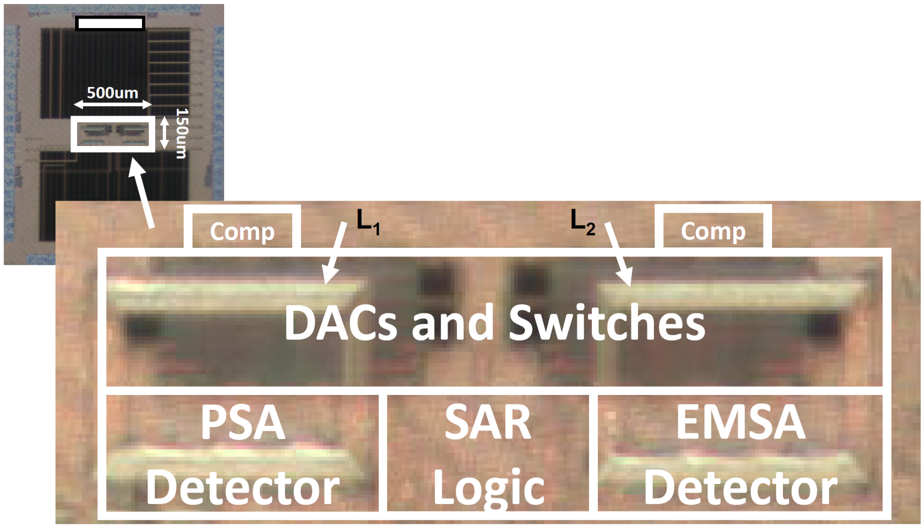

Figure 5-5: Die micrograph

5.6 Measurement Results

Figure. 5-6 demonstrates the effectiveness of our protection scheme in safeguarding the ADC against VDD-side PSA, GND-side PSA, and EMSA. The bit-wise accuracy for unprotected mode is close to 100% and the bit-wise accuracy for protected mode is close to 50%. This means that in the unprotected mode, the attacker can acquire the digital output accurately. For the protected mode, the ADC offers strong protection.

The Sniff-SAR is designed using the 65nm LP process. The ADC features a compact footprint, occupying only 0.075mm^2 as shown in Figure. 5-5. In comparison to previous work [7], which employed always-on protection, the Sniff-SAR introduces a detection-driven protection mechanism that is both energy-efficient and highly secure.

A unique switching scheme has been implemented in the Sniff-SAR, making it virtually untrainable using neural networks. This significantly enhances the ADC's resistance against advanced side-channel attacks. By incorporating duty cycling for both EMSA and PSA detectors, the Sniff-SAR achieves the highest sampling rate and best figure of merit (FoM) among secure ADCs.

The secure SAR mode of the Sniff-SAR has an FoM of 9.8fJ/conversion-step, which is comparable to state-of-the-art energy-efficient unprotected ADCs using similar technology. This remarkable performance highlights the potential of the Sniff-SAR

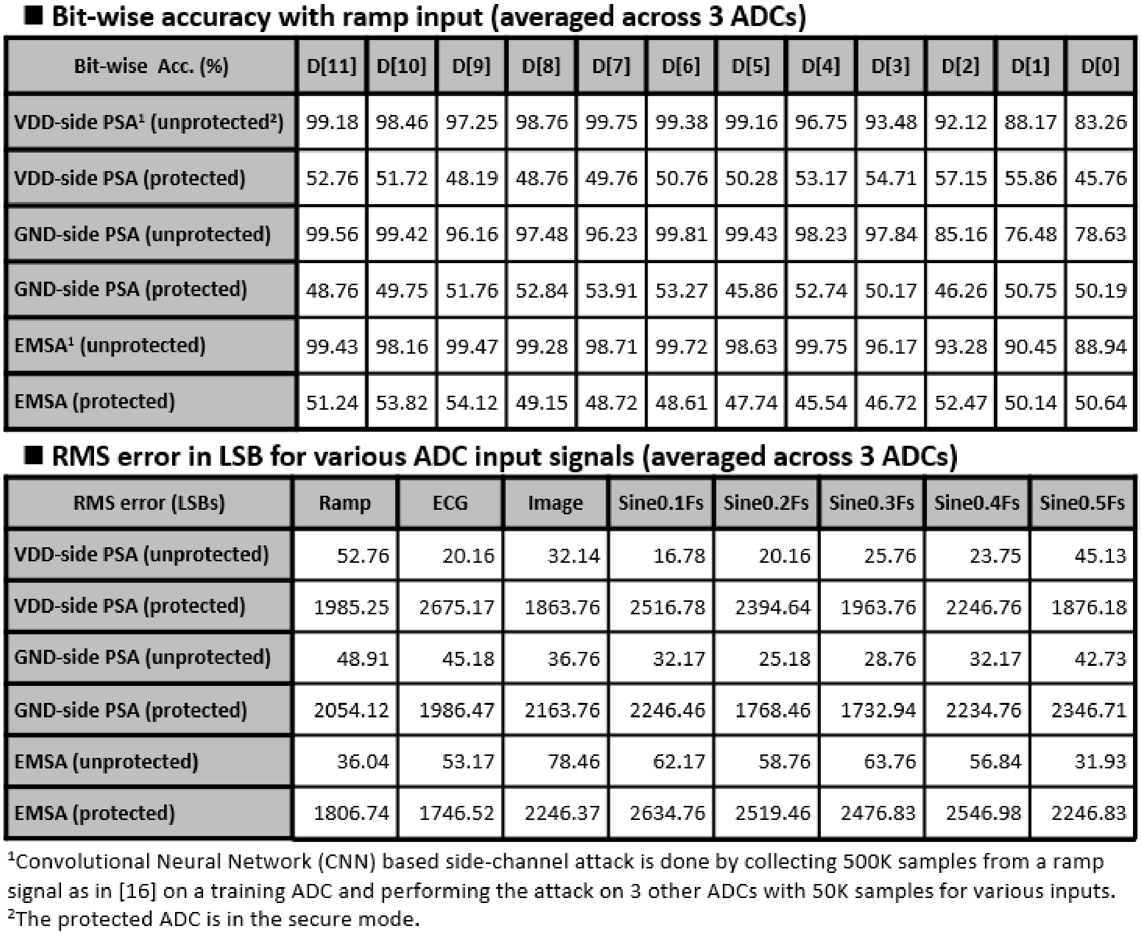

■ Bit-wise accuracy with ramp input (averaged across 3 ADCs)

Bit-wise Acc. (%)	D[11]	D[10]	D[9]	D[8]	D[7]	D[6]	D[5]	D[4]	D[3]	D[2]	D[1]	D[0]
VDD-side PSA[1] (unprotected[2])	99.18	98.46	97.25	98.76	99.75	99.38	99.16	96.75	93.48	92.12	88.17	83.26
VDD-side PSA (protected)	52.76	51.72	48.19	48.76	49.76	50.76	50.28	53.17	54.71	57.15	55.86	45.76
GND-side PSA (unprotected)	99.56	99.42	96.16	97.48	96.23	99.81	99.43	98.23	97.84	85.16	76.48	78.63
GND-side PSA (protected)	48.76	49.75	51.76	52.84	53.91	53.27	45.86	52.74	50.17	46.26	50.75	50.19
EMSA[1] (unprotected)	99.43	98.16	99.47	99.28	98.71	99.72	98.63	99.75	96.17	93.28	90.45	88.94
EMSA (protected)	51.24	53.82	54.12	49.15	48.72	48.61	47.74	45.54	46.72	52.47	50.14	50.64

■ RMS error in LSB for various ADC input signals (averaged across 3 ADCs)

RMS error (LSBs)	Ramp	ECG	Image	Sine0.1Fs	Sine0.2Fs	Sine0.3Fs	Sine0.4Fs	Sine0.5Fs
VDD-side PSA (unprotected)	52.76	20.16	32.14	16.78	20.16	25.76	23.75	45.13
VDD-side PSA (protected)	1985.25	2675.17	1863.76	2516.78	2394.64	1963.76	2246.76	1876.18
GND-side PSA (unprotected)	48.91	45.18	36.76	32.17	25.18	28.76	32.17	42.73
GND-side PSA (protected)	2054.12	1986.47	2163.76	2246.46	1768.46	1732.94	2234.76	2346.71
EMSA (unprotected)	36.04	53.17	78.46	62.17	58.76	63.76	56.84	31.93
EMSA (protected)	1806.74	1746.52	2246.37	2634.76	2519.46	2476.83	2546.98	2246.83

[1] Convolutional Neural Network (CNN) based side-channel attack is done by collecting 500K samples from a ramp signal as in [16] on a training ADC and performing the attack on 3 other ADCs with 50K samples for various inputs.
[2] The protected ADC is in the secure mode.

Figure 5-6: Power/EM attacks of unprotected vs. protected mode of the Sniff-SAR

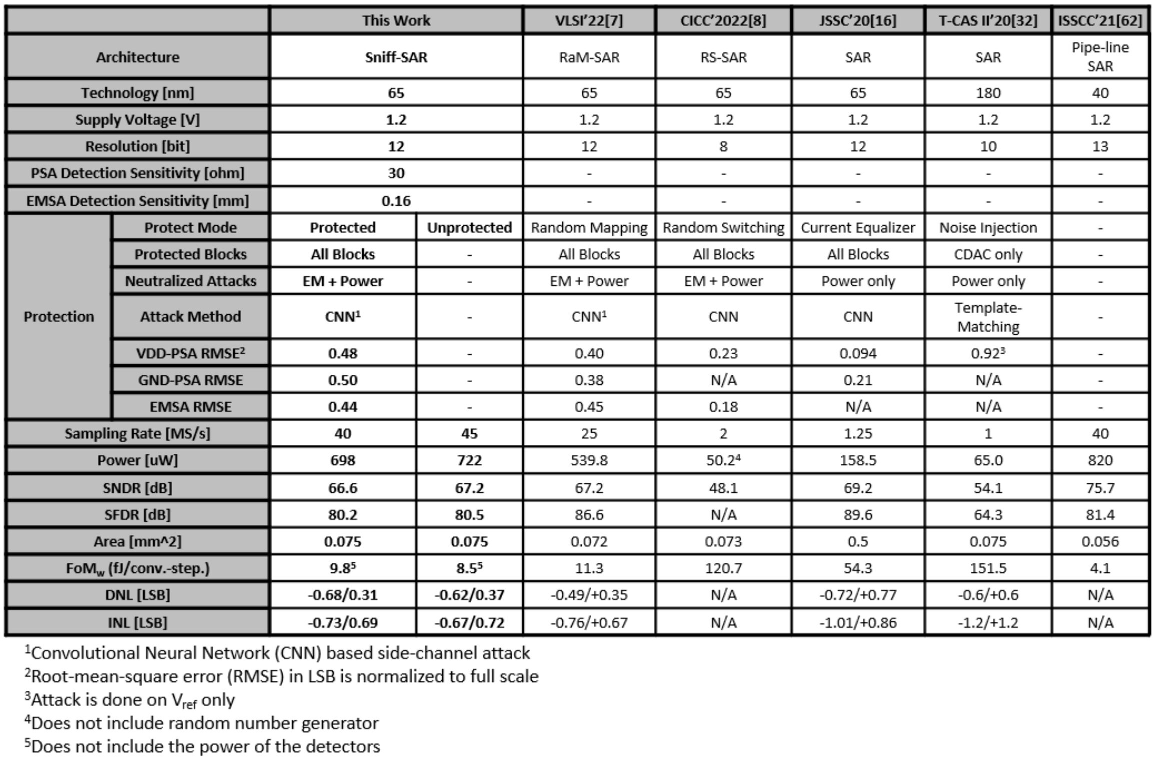

		This Work		VLSI'22[7]	CICC'2022[8]	JSSC'20[16]	T-CAS II'20[32]	ISSCC'21[62]
Architecture		Sniff-SAR		RaM-SAR	RS-SAR	SAR	SAR	Pipe-line SAR
Technology [nm]		65		65	65	65	180	40
Supply Voltage [V]		1.2		1.2	1.2	1.2	1.2	1.2
Resolution [bit]		12		12	8	12	10	13
PSA Detection Sensitivity [ohm]		30		-	-	-	-	-
EMSA Detection Sensitivity [mm]		0.16		-	-	-	-	-
Protection	Protect Mode	Protected	Unprotected	Random Mapping	Random Switching	Current Equalizer	Noise Injection	-
	Protected Blocks	All Blocks	-	All Blocks	All Blocks	All Blocks	CDAC only	-
	Neutralized Attacks	EM + Power	-	EM + Power	EM + Power	Power only	Power only	-
	Attack Method	CNN[1]	-	CNN[1]	CNN	CNN	Template-Matching	-
	VDD-PSA RMSE[2]	0.48	-	0.40	0.23	0.094	0.92[3]	-
	GND-PSA RMSE	0.50	-	0.38	N/A	0.21	N/A	-
	EMSA RMSE	0.44	-	0.45	0.18	N/A	N/A	-
Sampling Rate [MS/s]		40	45	25	2	1.25	1	40
Power [uW]		698	722	539.8	50.2[4]	158.5	65.0	820
SNDR [dB]		66.6	67.2	67.2	48.1	69.2	54.1	75.7
SFDR [dB]		80.2	80.5	86.6	N/A	89.6	64.3	81.4
Area [mm^2]		0.075	0.075	0.072	0.073	0.5	0.075	0.056
FoM$_w$ (fJ/conv.-step.)		9.8[5]	8.5[5]	11.3	120.7	54.3	151.5	4.1
DNL [LSB]		-0.68/0.31	-0.62/0.37	-0.49/+0.35	N/A	-0.72/+0.77	-0.6/+0.6	N/A
INL [LSB]		-0.73/0.69	-0.67/0.72	-0.76/+0.67	N/A	-1.01/+0.86	-1.2/+1.2	N/A

[1] Convolutional Neural Network (CNN) based side-channel attack

[2] Root-mean-square error (RMSE) in LSB is normalized to full scale

[3] Attack is done on V$_{ref}$ only

[4] Does not include random number generator

[5] Does not include the power of the detectors

Figure 5-7: Performance summary and comparison with the state-of-the-art secure ADCs and energy-efficient ADCs

to provide both security and energy efficiency in a wide range of applications.

To evaluate the performance of the secure SAR mode, an example image 5-10 is fed into the Sniff-SAR ADC. The obtained results (Figure. 5-12) show that with the protection enabled, the EMSA outcome appears almost random and does not reveal any useful information about the original image. This indicates that the secure SAR mode effectively protects the ADC from side-channel attacks when compared to unprotected SAR.

The Sniff-SAR successfully detects and protects against both power and electromagnetic side-channel attacks. The low power consumption of the EMSA and PSA detectors, combined with the effective performance of the secure SAR mode, demonstrates the potential of this design to enhance the security and resilience of electronic systems in a wide range of applications.

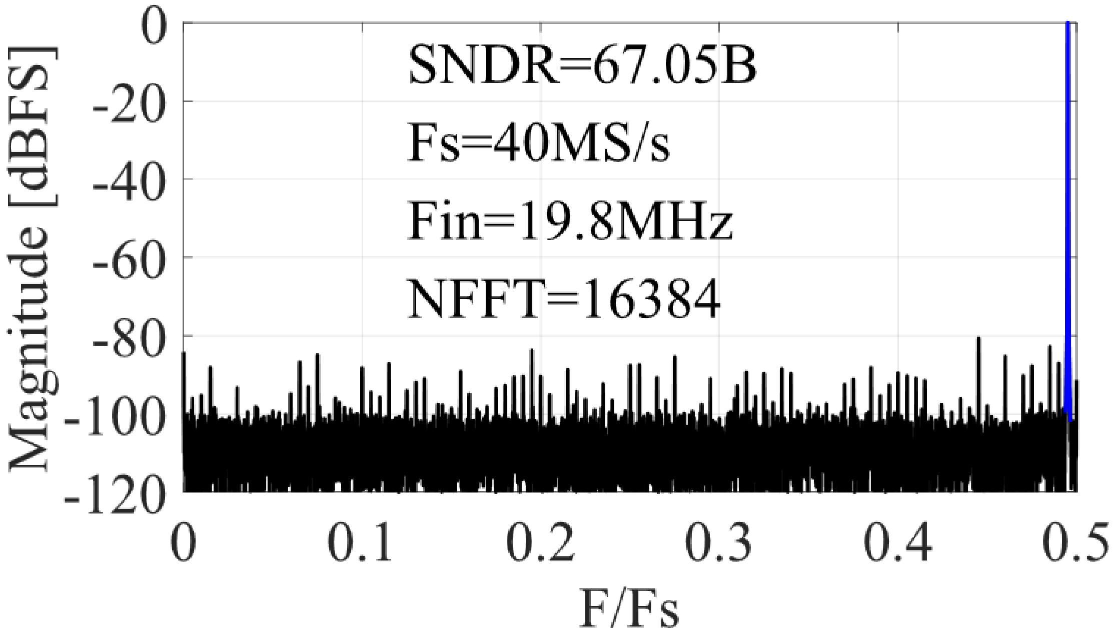

Figure 5-8: Measured output spectra of the secure SAR

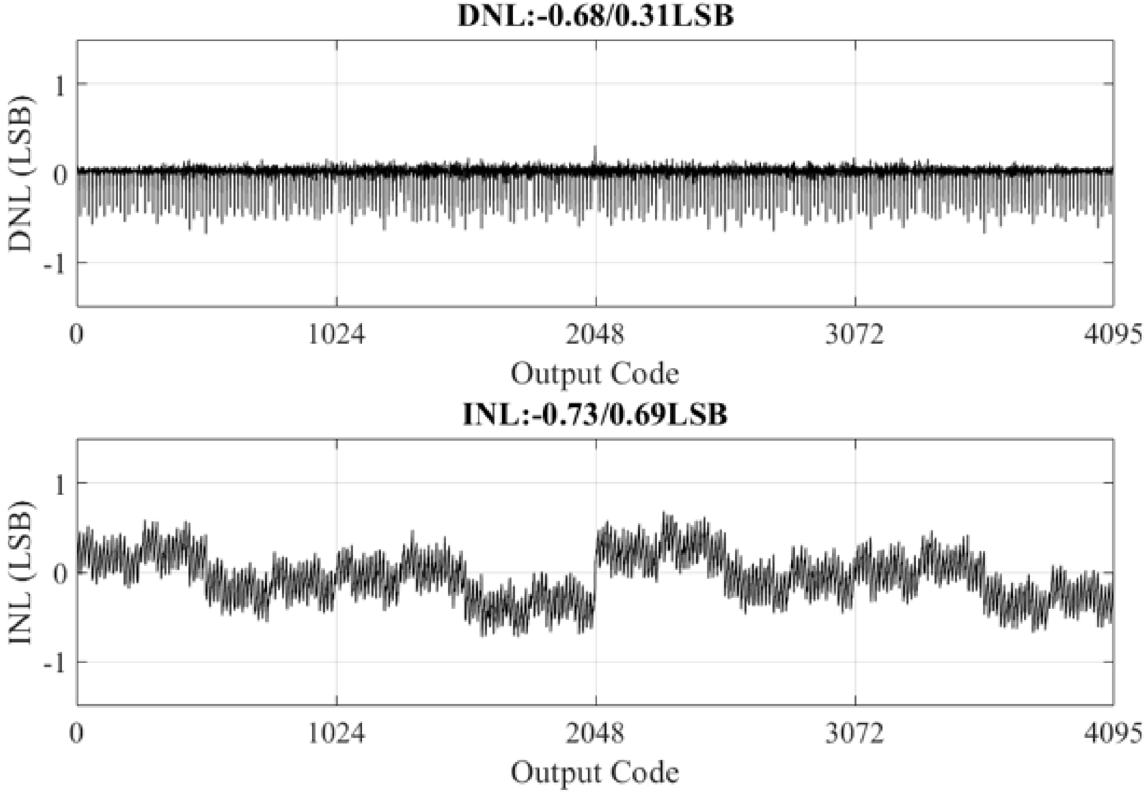

Figure 5-9: Measured DNL/INL of the unprotected mode SAR ADC after calibration

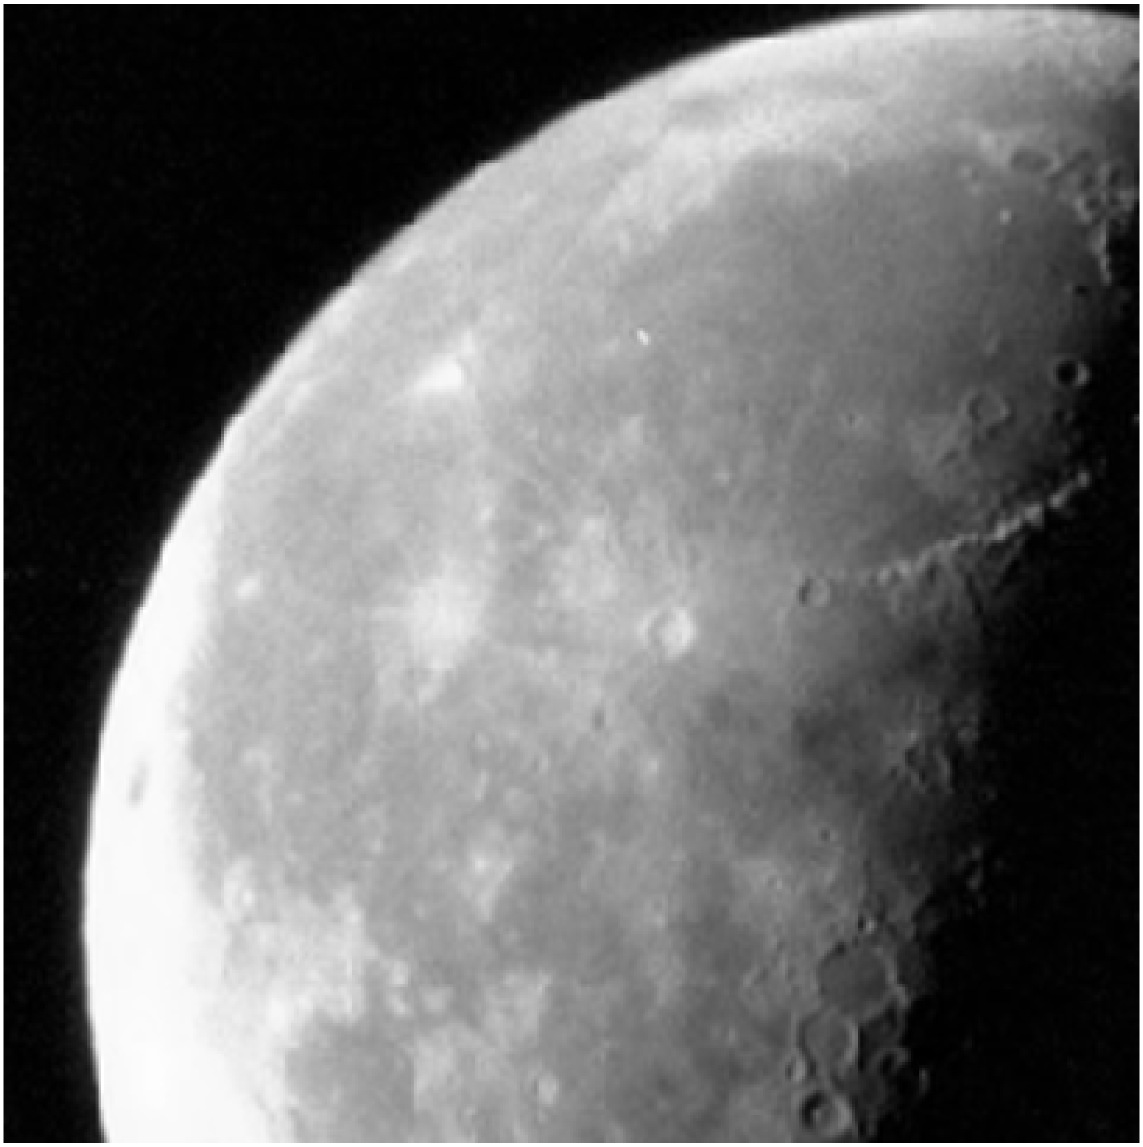

Figure 5-10: Example image

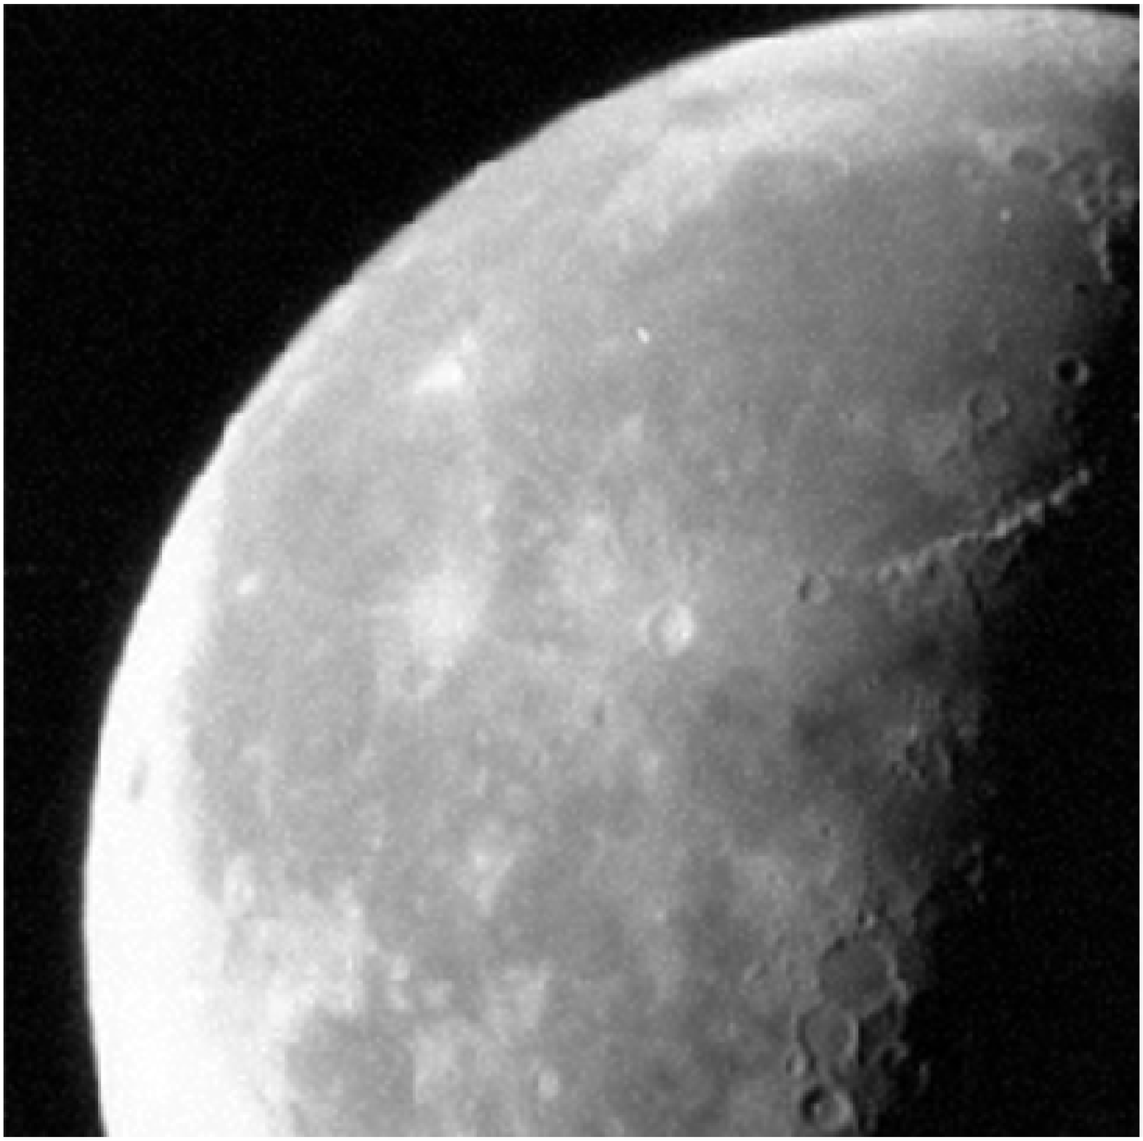

Figure 5-11: EMSA result on the unprotected SAR ADC

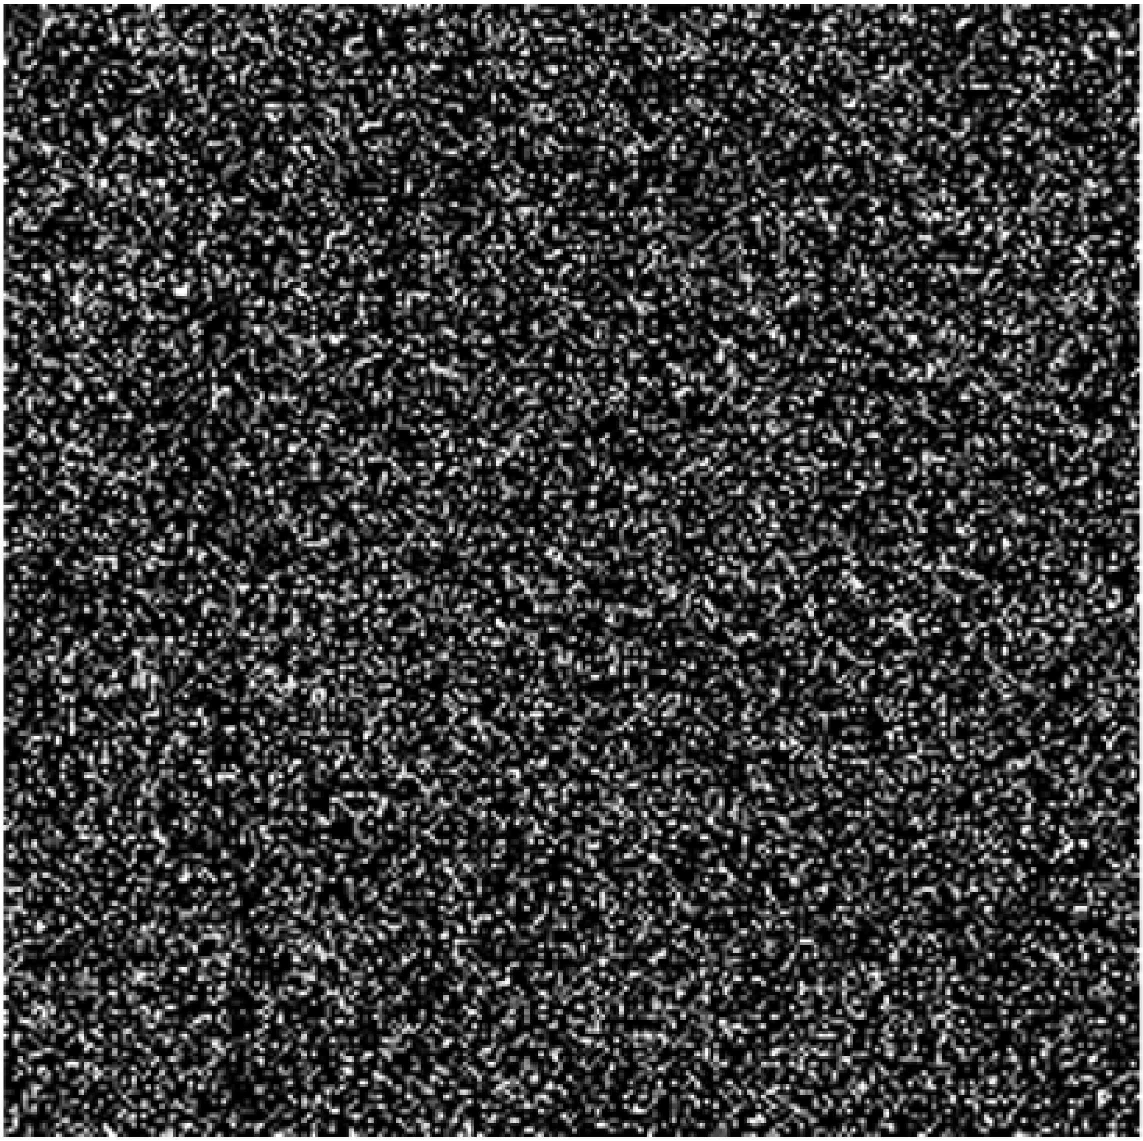

Figure 5-12: EMSA result on the protected SAR ADC

Chapter 6

Conclusions and Future Work

As the concluding chapter, this chapter discusses the contributions of the thesis and suggests future research directions.

6.1 Thesis Contributions

6.1.1 Direct HESE SAR

This research presents a sparsity-aware analog-to-digital converter (ADC) for analog neural networks, specifically a bit-level sparsity-aware SAR ADC that directly produces Hybrid Encoding for Signed Expressions (HESE). The primary goal of the ADC is to support large artificial neural networks (ANNs) while maintaining high accuracy. With a 12-bit resolution, the HESE ADC is designed to efficiently process data, capitalizing on the sparsity inherent in the neural networks.

The HESE ADC employs a unique architecture that incorporates two thresholds for 2-bit look-ahead (LA) functionality, which facilitates more efficient processing by anticipating the next two bits of the data. Additionally, noise averaging (NA) is implemented during the final two bit-cycles of the ADC's operation, reducing the impact of noise on the overall performance of the system.

The proposed sparsity-aware ADC has significant advantages over traditional ADCs, particularly in the context of large-scale ANNs. By exploiting sparsity in the data, the HESE ADC can achieve improved efficiency, leading to reduced power consumption and faster processing times. This is essential for modern applications, such as machine learning and artificial intelligence, where rapid and efficient data processing is crucial. The HESE SAR ADC exhibits a figure of merit (FoM) of 15.2 fJ/conv.-step when operating at 45MS/s with a core area of just $0.072\mathbf{mm^2}$.

6.1.2 RaM-SAR

ADCs are susceptible to side-channel attacks, where adversaries exploit the correlation between power consumption or electromagnetic emissions and the A/D conversion to extract sensitive information. To tackle this challenge and enhance the security of ADCs, we introduce the RaM-SAR

RaM-SAR is a secure random-mapping SAR ADC that offers robust resistance against both power and electromagnetic side-channel attacks. This 12-bit, 25 MS/s ADC achieves an energy consumption of only 11.3 fJ per conversion step, making it suitable for energy-constrained applications. The key to RaM-SAR's heightened security is a unique random-mapping technique that randomly assigns each conversion to one of the thousands of different conversion sequences. This randomization effectively disrupts the predictable patterns that adversaries rely on to perform successful side-channel attacks, thereby safeguarding sensitive data.

The random-mapping technique used in RaM-SAR is designed to provide comprehensive protection against various side-channel attacks, including both power and electromagnetic attacks. By disrupting the correlations between the leaked signals and the internal operations of the ADC, RaM-SAR can thwart neural network-based power and electromagnetic side-channel attacks, which are becoming increasingly sophisticated and challenging to defend against.

The low energy consumption per conversion step enables RaM-SAR to be employed in a wide range of applications where power constraints are a critical consideration. Moreover, the 25 MS/s sampling rate ensures that RaM-SAR can deliver the performance required for high-speed data acquisition and processing tasks.

The innovative design of RaM-SAR is an attractive solution for enhancing the security of ADCs in IoT and wearable devices. By incorporating RaM-SAR into these systems, designers can provide a crucial layer of protection against increasingly sophisticated side-channel attacks, ensuring the privacy and integrity of sensitive data.

6.1.3 Sniff-SAR

As the demand for secure ADCs has grown, there has been an increasing need for ADCs that can detect and defend against power and electromagnetic (EM) side-channel attacks. In response to this challenge, we have developed Sniff-SAR, an innovative 9.8 fJ/conversion-step, 12-bit secure ADC designed to offer robust protection against side-channel attacks through the use of EMSA (electromagnetic side-channel analysis) and PSA (power side-channel analysis) detectors.

Under normal circumstances, the ADC core of Sniff-SAR performs the analog-to-digital conversion process while operating in an unprotected SAR (successive approximation register) mode. This mode is both faster and more energy-efficient compared to the secure mode, making it the default choice for routine operation. However, Sniff-SAR's EMSA and PSA detectors periodically scan for signs of side-channel attacks. Upon detecting a potential threat, the ADC automatically activates its secure SAR mode, which is specifically designed to defend against both EMSA and PSA attacks.

The secure SAR mode of Sniff-SAR offers an extraordinary degree of protection, generating 3.6 x 10^{16} different switching traces. This vast number of traces makes it practically infeasible for attackers to train neural networks for PSA or EMSA attacks.

The development of Sniff-SAR is an important milestone in the pursuit of secure ADCs, offering a versatile solution that can significantly enhance the security of a broad range of electronic systems. The innovative features of Sniff-SAR, such as its detection-driven protection and secure SAR mode, ensure that it remains a formidable defense against side-channel attacks.

Furthermore, the flexible and adaptable nature of Sniff-SAR allows it to be easily integrated into existing systems and technologies. This seamless compatibility makes Sniff-SAR an appealing choice for designers looking to improve the security of ADCs in electronic systems, ranging from IoT devices to wearable technology and beyond.

Sniff-SAR represents a significant advancement in the field of secure ADC design. By incorporating detection-driven protection against side-channel attacks and utilizing a secure SAR mode that produces an enormous number of switching traces, Sniff-SAR offers an effective and adaptable solution for enhancing the security of ADCs across a wide array of electronic systems. With its combination of energy efficiency, speed, and compatibility with existing technologies, Sniff-SAR has the potential to play a pivotal role in the ongoing quest to safeguard sensitive data and protect against increasingly sophisticated side-channel attacks.

6.2 Future directions

6.2.1 SCA with different Neural Networks for time series data

Previous Neural-Network-Based ADC SCAs are mostly using CNNs and MLP. CNNs are more powerful than MLP in various image tasks. CNNs are a specialized class of neural networks that have proven to be highly effective in solving problems associated with spatial data, such as images. CNNs have gained widespread popularity in the field of computer vision due to their ability to automatically learn features from raw image data and their capability to achieve state-of-the-art performance in various image recognition and classification tasks. The key component of a CNN is the convolutional layer, which performs local operations on the input data to learn spatial hierarchies and capture spatial patterns. This property makes CNNs particularly well-suited for processing images and other grid-like data structures.

On the other hand, Recurrent Neural Networks (RNNs) [67] are specifically designed to address problems involving temporal or sequential data, such as text, speech, and time series. The ADC leakage traces are time series. RNN-based ADC SCA can be more powerful than CNN-based ADC CSA. RNNs possess a unique architecture that allows them to maintain internal states and capture dependencies across variable-length sequences. This is achieved through the incorporation of recurrent connections, which enable information to persist across time steps, allowing the network to model complex temporal relationships. RNNs have found numerous applications in natural language processing, speech recognition, and various sequence prediction tasks.

Transformers [68] can take the ADC SCA a step further. Transformers have emerged as a highly efficient alternative to RNN-based models, such as Long Short-Term Memory (LSTM) networks when it comes to handling sequential data. One key advantage of transformers over RNNs is their ability to process the entire ADC leakage sequence in parallel, rather than sequentially. This is made possible by the self-attention mechanism employed by transformers, which allows them to directly capture dependencies between all pairs of tokens in the leakage sequence, irrespective of their positions. As a result, transformers can achieve significantly faster processing speeds compared to RNN-based models, which must process input data sequentially, making it challenging to fully exploit the parallel computing capabilities of modern hardware like GPUs.

Training LSTMs can be more difficult in comparison to transformer networks, mainly due to the higher number of parameters typically involved in LSTM networks. LSTMs rely on gating mechanisms to control the flow of information through their hidden states, which necessitates additional learnable parameters for each gate. These parameters not only increase the complexity of the model but can also lead to issues such as vanishing or exploding gradients, making it challenging to optimize the network during training. Moreover, the inherently sequential nature of LSTMs can result in longer training times, as they cannot fully leverage the benefits of parallelization.

In contrast, transformer networks usually have a more streamlined architecture, characterized by a lower number of parameters and a greater capacity for parallelization. These factors contribute to more efficient and effective training, enabling transformers to converge faster and often achieve superior performance on various tasks. The success of transformer-based models, such as BERT [69] and GPT-3 [70], has further demonstrated their ability to outperform RNN-based models in many natural language processing tasks, solidifying the transformer as the state-of-the-art architecture for sequential data processing. It is promising to use transformer networks for ADC SCA. Far-field EMSA can be enabled with transformer networks.

6.2.2 Non-profiled attacks with transfer learning

The conventional machine learning approach assumes that both training and testing data originate from the same domain, resulting in similar data distribution characteristics and input feature spaces. Nonetheless, real-world scenarios often do not abide by this assumption, and collecting appropriate training data can be challenging or costly [71].

Therefore, there has been a growing interest in developing high-performance learners that can learn from easily accessible data originating from different domains. This new approach, known as transfer learning, aims to transfer knowledge and insights from one domain to another, leveraging the similarities between the domains and exploiting the knowledge that has already been acquired in one domain to improve learning in the other domain. The goal of transfer learning is to overcome the scarcity or difficulty of collecting training data in the target domain by utilizing the knowledge obtained from a different but related domain.

An interesting avenue for ADC SCA is through the use of transfer learning. Profiled ADC SCAs present a unique challenge, particularly as the part number of the ADC under attack may be unknown. Additionally, collecting appropriate training data for the specific ADC under attack can be a challenging and costly endeavor. With transfer learning, however, attackers can leverage knowledge and insights gained from similar ADCs to circumvent the need for training data for each new ADC. By using transfer learning, attackers can effectively transfer knowledge and insights from related but distinct domains, providing a promising approach to improve the performance of machine learning models in scenarios where the specific target ADC is unknown or training data is scarce.

6.2.3 Remote ADC SCA

ADC SCA can be more powerful if the SCA can be done remotely without the notice of the ADC user. Far-field EMSA and remote PSA are two possible methods.

In [72], the authors present a novel deep learning-based side-channel attack on AES-128 encryption using far-field electromagnetic emissions. Their neural networks, trained on traces from five Bluetooth devices at different distances from the target, effectively recover encryption keys even at 15 meters away in an office environment. This research highlights the potential of far-field EMSA. The sensitive ADC input might be recovered by far-field EMSA without the notice of the ADC user.

In [73], the researchers discuss the growing adoption of heterogeneous computing, which has led to the integration of Field Programmable Gate Arrays (FPGAs) into cloud data centers and flexible System-on-Chips (SoCs). They reveal that the integrated FPGA introduces a new security vulnerability by enabling software-based power side-channel attacks without needing physical proximity to the target system. The paper demonstrates the construction of an on-chip power monitor using ring oscillators (ROs) on a modern FPGA and its ability to observe power consumption of other modules on the FPGA or SoC. The researchers then show the successful power analysis attack on an RSA cryptomodule using the RO-based FPGA power monitor. They also demonstrate that the FPGA-based power monitor can observe the power consumption of a CPU on the same SoC and break timing-channel protection for an RSA program running on the CPU. This work highlights remote power side-channel attacks using an FPGA, challenging the common assumption that such attacks require specialized equipment and physical access to the victim hardware when an integrated FPGA is involved. Similarly, ADCs that are integrated with FPGA might be under similar remote PSA.

6.2.4 Low-overhead ADC SCA Countermeasure

This thesis has proposed two schemes to reduce the overhead of the ADC SCA countermeasure. The energy efficiency, area, and bandwidth can be further improved for secure ADCs. The secure ADCs can be widely adopted in real-world applications if this overhead is removed.

There are numerous possible methods that can be further investigated and developed to address this challenge. First and foremost, a more efficient random conversion scheme could be proposed, which would enable ADCs to process input data in a more secure manner without compromising their overall performance. This could involve the use of advanced cryptographic techniques, such as true random number generation, to ensure that the ADC's conversion process is both unpredictable and resistant to side-channel attacks.

Second, signature attenuation and information masking can be integrated into conventional on-chip power management blocks, with the goal of concealing ADC activities without incurring any additional performance overhead. By obfuscating the power traces and electromagnetic emissions generated during the ADC's operation, these techniques can help prevent adversaries from exploiting side-channel information to reveal the ADC's internal states.

Bibliography

[1] HT Kung, Bradley McDanel, and Sai Qian Zhang. Term quantization: furthering quantization at run time. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–14, 2020.

[2] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.

[3] Gordon E Moore et al. Cramming more components onto integrated circuits, 1965.

[4] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.

[5] Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal processing magazine*, 29(6):82–97, 2012.

[6] Ruhi Sarikaya, Geoffrey E Hinton, and Anoop Deoras. Application of deep belief networks for natural language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 22(4):778–784, 2014.

[7] Ruicong Chen, Hanrui Wang, Anantha Chandrakasan, and Hae-Seung Lee. Ramsar: A low energy and area overhead, 11.3 fj/conv.-step 12b 25ms/s secure random-mapping sar adc with power and em side-channel attack resilience. In *2022 IEEE Symposium on VLSI Technology and Circuits (VLSI Technology and Circuits)*, pages 94–95. IEEE, 2022.

[8] Maitreyi Ashok, Edlyn V Levine, and Anantha P Chandrakasan. Randomized switching sar (rs-sar) adc protections for power and electromagnetic side channel security. In *2022 IEEE Custom Integrated Circuits Conference (CICC)*, pages 1–2. IEEE, 2022.

[9] Patrick McDaniel and Stephen McLaughlin. Security and privacy challenges in the smart grid. *IEEE security & privacy*, 7(3):75–77, 2009.

[10] Data Encryption Standard et al. Federal information processing standards publication 46. *National Bureau of Standards, US Department of Commerce*, 23:1–18, 1977.

[11] Eric Rescorla, Hannes Tschofenig, and Nagendra Modadugu. The datagram transport layer security (dtls) protocol version 1.3. *Internet Engineering Task Force*, 2019.

[12] Paul Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In *Advances in Cryptology—CRYPTO'99: 19th Annual International Cryptology Conference Santa Barbara, California, USA, August 15–19, 1999 Proceedings 19*, pages 388–397. Springer, 1999.

[13] Eric Brier, Christophe Clavier, and Francis Olivier. Correlation power analysis with a leakage model. In *Cryptographic Hardware and Embedded Systems-CHES 2004: 6th International Workshop Cambridge, MA, USA, August 11-13, 2004. Proceedings 6*, pages 16–29. Springer, 2004.

[14] Benjamin Timon. Non-profiled deep learning-based side-channel attacks with sensitivity analysis. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pages 107–131, 2019.

[15] Vinod V Gadde, Hiromitsu Awano, and Makoto Ikeda. An encryption-authentication unified a/d conversion scheme for iot sensor nodes. In *2018 IEEE Asian Solid-State Circuits Conference (A-SSCC)*, pages 123–126. IEEE, 2018.

[16] Taehoon Jeong, Anantha P Chandrakasan, and Hae-Seung Lee. S2adc: A 12-bit, 1.25-ms/s secure sar adc with power side-channel attack resistance. *IEEE Journal of Solid-State Circuits*, 56(3):844–854, 2020.

[17] B. Murmann. ADC Perfoarmance Survey 1997-2020. [Online] Available: `http://web.stanford.edu/~murmann/adcsurvey.html`, 2021.

[18] Ruicong Chen, HT Kung, Anantha Chandrakasan, and Hae-Seung Lee. A bit-level sparsity-aware sar adc with direct hybrid encoding for signed expressions for aiot applications. In *Proceedings of the ACM/IEEE International Symposium on Low Power Electronics and Design*, pages 1–6, 2022.

[19] Wm A Wulf and Sally A McKee. Hitting the memory wall: Implications of the obvious. *ACM SIGARCH computer architecture news*, 23(1):20–24, 1995.

[20] Mark Horowitz. 1.1 computing's energy problem (and what we can do about it). In *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, pages 10–14. IEEE, 2014.

[21] Ali Shafiee, Anirban Nag, Naveen Muralimanohar, Rajeev Balasubramonian, John Paul Strachan, Miao Hu, R Stanley Williams, and Vivek Srikumar. Isaac: A convolutional neural network accelerator with in-situ analog arithmetic in crossbars. *ACM SIGARCH Computer Architecture News*, 44(3):14–26, 2016.

[22] Brady D Lund and Ting Wang. Chatting about chatgpt: how may ai and gpt impact academia and libraries? *Library Hi Tech News*, 2023.

[23] Vivienne Sze, Yu-Hsin Chen, Tien-Ju Yang, and Joel S Emer. Efficient processing of deep neural networks: A tutorial and survey. *Proceedings of the IEEE*, 105(12):2295–2329, 2017.

[24] Vadim Gutnik and Anantha P Chandrakasan. Embedded power supply for low-power dsp. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 5(4):425–435, 1997.

[25] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.

[26] Werner Schindler, Kerstin Lemke, and Christof Paar. A stochastic model for differential side channel cryptanalysis. In *Cryptographic Hardware and Embedded Systems–CHES 2005: 7th International Workshop, Edinburgh, UK, August 29–September 1, 2005. Proceedings 7*, pages 30–46. Springer, 2005.

[27] Richard Gilmore, Neil Hanley, and Maire O'Neill. Neural network based attack on a masked implementation of aes. In *2015 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*, pages 106–111. IEEE, 2015.

[28] Mehdi-Laurent Akkar and Christophe Giraud. An implementation of des and aes, secure against some attacks. In *Cryptographic Hardware and Embedded Systems—CHES 2001: Third International Workshop Paris, France, May 14–16, 2001 Proceedings 3*, pages 309–318. Springer, 2001.

[29] Kris Tiri and Ingrid Verbauwhede. A digital design flow for secure integrated circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 25(7):1197–1208, 2006.

[30] Carlos Tokunaga and David Blaauw. Secure aes engine with a local switched-capacitor current equalizer. In *2009 IEEE International Solid-State Circuits Conference-Digest of Technical Papers*, pages 64–65. IEEE, 2009.

[31] Monodeep Kar, Arvind Singh, Sanu Mathew, Anand Rajan, Vivek De, and Saibal Mukhopadhyay. 8.1 improved power-side-channel-attack resistance of an aes-128 core via a security-aware integrated buck voltage regulator. In *2017 IEEE International Solid-State Circuits Conference (ISSCC)*, pages 142–143. IEEE, 2017.

[32] Takuji Miki, Noriyuki Miura, Hiroki Sonoda, Kento Mizuta, and Makoto Nagata. A random interrupt dithering sar technique for secure adc against reference-charge side-channel attack. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 67(1):14–18, 2019.

[33] Susana Borromeo, Cristina Rodriguez-Sanchez, Felipe Machado, Juan Antonio Hernandez-Tamames, and Roberto de la Prieta. A reconfigurable, wearable, wireless ecg system. In *2007 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 1659–1662. IEEE, 2007.

[34] Hae-Sung Lee, D. Hodges, and P. Gray. A self calibrating 12b 12 µ s cmos adc. In *1984 IEEE International Solid-State Circuits Conference - (ISSCC). Digest of Technical Papers*, volume XXVII, pages 64–65, 1984.

[35] Arvind Singh, Monodeep Kar, Sanu Mathew, Anand Rajan, Vivek De, and Saibal Mukhopadhyay. 25.3 a 128b aes engine with higher resistance to power and electromagnetic side-channel attacks enabled by a security-aware integrated all-digital low-dropout regulator. In *2019 IEEE International Solid-State Circuits Conference-(ISSCC)*, pages 404–406. IEEE, 2019.

[36] Archisman Ghosh, Debayan Das, Josef Danial, Vivek De, Santosh Ghosh, and Shreyas Sen. 36.2 an em/power sca-resilient aes-256 with synthesizable signature attenuation using digital-friendly current source and ro-bleed-based integrated local feedback and global switched-mode control. In *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, volume 64, pages 499–501. IEEE, 2021.

[37] Debayan Das, Josef Danial, Anupam Golder, Nirmoy Modak, Shovan Maity, Baibhab Chatterjee, Donghyun Seo, Muya Chang, Avinash Varna, Harish Krishnamurthy, et al. 27.3 em and power sca-resilient aes-256 in 65nm cmos through> 350× current-domain signature attenuation. In *2020 IEEE International Solid-State Circuits Conference-(ISSCC)*, pages 424–426. IEEE, 2020.

[38] K. L. Loh. 1.2 fertilizing aiot from roots to leaves. In *2020 IEEE International Solid- State Circuits Conference - (ISSCC)*, pages 15–21, 2020.

[39] Ruicong Chen. *Activity-Scaling SAR with Direct Hybrid Encoding for Signed Expressions for AIoT Applications*. PhD thesis, Massachusetts Institute of Technology, 2021.

[40] W. Shan, M. Yang, J. Xu, Y. Lu, S. Zhang, T. Wang, J. Yang, L. Shi, and M. Seok. 14.1 a 510nw 0.41v low-memory low-computation keyword-spotting chip using serial fft-based mfcc and binarized depthwise separable convolutional neural network in 28nm cmos. In *2020 IEEE International Solid- State Circuits Conference - (ISSCC)*, pages 230–232, 2020.

[41] J. K. Brown, D. Abdallah, J. Boley, N. Collins, K. Craig, G. Glennon, K. Huang, C. J. Lukas, W. Moore, R. K. Sawyer, Y. Shakhsheer, F. B. Yahya, A. Wang, N. E. Roberts, D. D. Wentzloff, and B. H. Calhoun. 27.1 a 65nm energy-harvesting ulp soc with 256kb cortex-m0 enabling an 89.1µw continuous machine health monitoring wireless self-powered system. In *2020 IEEE International Solid- State Circuits Conference - (ISSCC)*, pages 420–422, 2020.

[42] Mohamed R. Abdelhamid, Ruicong Chen, Joonhyuk Cho, Anantha P. Chandrakasan, and Fadel Adib. Self-reconfigurable micro-implants for cross-tissue wireless and batteryless connectivity. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking - (MobiCOM)*. Association for Computing Machinery, 2020.

[43] Yu-Hsin Chen, Tushar Krishna, Joel S Emer, and Vivienne Sze. Eyeriss: An energy-efficient reconfigurable accelerator for deep convolutional neural networks. *IEEE journal of solid-state circuits*, 52(1):127–138, 2016.

[44] Bradley McDanel, HT Kung, and Sai Qian Zhang. Saturation rram leveraging bit-level sparsity resulting from term quantization. In *2021 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5. IEEE, 2021.

[45] Teyuh Chou, Wei Tang, Jacob Botimer, and Zhengya Zhang. Cascade: Connecting rrams to extend analog dataflow in an end-to-end in-memory processing paradigm. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 114–125, 2019.

[46] Ruicong Chen, HT Kung, Anantha Chandrakasan, and Hae-Seung Lee. A bit-level sparsity-aware sar adc with direct hybrid encoding for signed expressions leveraging algorithm-circuit co-design. *MIT MTL Annual Research Report*, 2022.

[47] Ruicong Chen, Anantha Chandrakasan, and Hae-Seung Lee. Direct hybrid encoding for signed expressions sar adc for analog neural networks. *MIT MTL Annual Research Report*, 2021.

[48] HT Kung. High-order-bit first conversion for signed-digit representations. In *Annual GOMACTech Conference*. IEEE, 2021.

[49] Yanfei Chen, Xiaolei Zhu, Hirotaka Tamura, Masaya Kibune, Yasumoto Tomita, Takayuki Hamada, Masato Yoshioka, Kiyoshi Ishikawa, Takeshi Takayama, Junji Ogawa, et al. Split capacitor dac mismatch calibration in successive approximation adc. In *2009 IEEE Custom Integrated Circuits Conference*, pages 279–282. IEEE, 2009.

[50] Ming Ding, Pieter Harpe, Yao-Hong Liu, Benjamin Busze, Kathleen Philips, and Harmke de Groot. 26.2 a 5.5 fj/conv-step 6.4 ms/s 13b sar adc utilizing a redundancy-facilitated background error-detection-and-correction scheme. In *2015 IEEE International Solid-State Circuits Conference-(ISSCC) Digest of Technical Papers*, pages 1–3. IEEE, 2015.

[51] V Hariprasath, Jon Guerber, SH Lee, and UK Moon. Merged capacitor switching based sar adc with highest switching energy-efficiency. *Electronics letters*, 46(9):620, 2010.

[52] George Wegmann, Eric A Vittoz, and Fouad Rahali. Charge injection in analog mos switches. *IEEE Journal of Solid-state circuits*, 22(6):1091–1097, 1987.

[53] Un-Ku Moon and Bang-Sup Song. Background digital calibration techniques for pipelined adcs. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, 44(2):102–109, 1997.

[54] Sunghyuk Lee. *Techniques for low-power high-performance adcs.* PhD thesis, Massachusetts Institute of Technology, 2014.

[55] Andrew M Abo and Paul R Gray. A 1.5-v, 10-bit, 14.3-ms/s cmos pipeline analog-to-digital converter. *IEEE Journal of Solid-State Circuits*, 34(5):599–606, 1999.

[56] Behzad Razavi. The strongarm latch [a circuit for all seasons]. *IEEE Solid-State Circuits Magazine*, 7(2):12–17, 2015.

[57] Yannis P Tsividis, Paul R Gray, David A Hodges, and Jacob Chacko. A segmented μ-255 law pcm voice encoder utilizing nmos technology. *IEEE Journal of Solid-State Circuits*, 11(6):740–747, 1976.

[58] Qiang Ma, Linfu Xiao, Yiu-Cheong Tam, and Evangeline FY Young. Simultaneous handling of symmetry, common centroid, and general placement constraints. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 30(1):85–95, 2010.

[59] Naveen Verma and Anantha P Chandrakasan. An ultra low energy 12-bit rate-resolution scalable sar adc for wireless sensor nodes. *IEEE Journal of Solid-State Circuits*, 42(6):1196–1205, 2007.

[60] F. M. Yaul and A. P. Chandrakasan. A 10 bit sar adc with data-dependent energy reduction using lsb-first successive approximation. *IEEE Journal of Solid-State Circuits - (JSSC)*, 49(12):2825–2834, 2014.

[61] Brian P Ginsburg and Anantha P Chandrakasan. 500-ms/s 5-bit adc in 65-nm cmos with split capacitor array dac. *IEEE Journal of Solid-State Circuits*, 42(4):739–747, 2007.

[62] Xiyuan Tang, Xiangxing Yang, Jiaxin Liu, Wei Shi, David Z Pan, and Nan Sun. 27.4 a 0.4-to-40ms/s 75.7 db-sndr fully dynamic event-driven pipelined adc with 3-stage cascoded floating inverter amplifier. In *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, volume 64, pages 376–378. IEEE, 2021.

[63] Jiaxin Liu, Xiyuan Tang, Wenda Zhao, Linxiao Shen, and Nan Sun. A 13-bit 0.005-mm 2 40-ms/s sar adc with kt/c noise cancellation. *IEEE Journal of Solid-State Circuits*, 55(12):3260–3270, 2020.

[64] Ruicong Chen, Anantha Chandrakasan, and Hae-Seung Lee. Sniff-sar: A 9.8fj/c.-s 12b secure adc with detection driven protection against power and em side-channel attack. In *2023 IEEE Custom Integrated Circuit Conference (CICC), accepted for publication.* IEEE, 2023.

[65] Sung Justin Kim, Dongkwun Kim, Ayushparth Sharma, and Mingoo Seok. Eqz-ldo: A near-zero edp overhead,> 10m-attack-resilient, secure digital ldo featuring attack-detection and detection-driven protection for a correlation-power-analysis-resilient iot device. In *2021 Symposium on VLSI Circuits*, pages 1–2. IEEE, 2021.

[66] Noriyuki Miura, Daisuke Fujimoto, Daichi Tanaka, Yu-ichi Hayashi, Naofumi Homma, Takafumi Aoki, and Makoto Nagata. A local em-analysis attack resistant cryptographic engine with fully-digital oscillator-based tamper-access sensor. In *2014 symposium on VLSI circuits digest of technical papers*, pages 1–2. IEEE, 2014.

[67] Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329*, 2014.

[68] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[69] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[70] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

[71] Karl Weiss, Taghi M Khoshgoftaar, and DingDing Wang. A survey of transfer learning. *Journal of Big data*, 3(1):1–40, 2016.

[72] Ruize Wang, Huanyu Wang, and Elena Dubrova. Far field em side-channel attack on aes using deep learning. In *Proceedings of the 4th ACM Workshop on Attacks and Solutions in Hardware Security*, pages 35–44, 2020.

[73] Mark Zhao and G Edward Suh. Fpga-based remote power side-channel attacks. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 229–244. IEEE, 2018.