

Late Breaking Results: Practical Power Side-Channel Attack on Analog Compute-in-Memory Macro

Simon Wilhelmstätter*, Johannes Stark*, Devanshi Upadhyaya†, Maël Gay†, Ilia Polian†, Maurits Ortmanns*

*Institute of Microelectronics, University of Ulm, Germany, firstname.lastname@uni-ulm.de

†Institute for Computer Architecture and Computer Engineering, University of Stuttgart, Germany, firstname.lastname@informatik.uni-stuttgart.de

Abstract—Analog compute-in-memory (ACIM) architectures emerged as energy efficient matrix–vector multiplication accelerators utilized for neural network inference in power-constrained environments. However, the security implications of ACIM hardware remain almost entirely unexplored. In particular, no previous work has evaluated information leakage of fabricated ACIM hardware through power side-channels. This work presents the first measured power side-channel analysis of a fabricated 28 nm ACIM macro. Using a convolutional neural network (CNN) workload, we show that power traces exhibit strong data-dependent information leakage that allows accurate reconstruction of private input images, with a mean structural similarity index measure (MSSIM) of up to 0.71.

Index Terms—Analog Compute-in-Memory, Hardware-Security, Power Side-Channel Analysis

I. INTRODUCTION

Compute-in-memory architectures have emerged as a solution to overcome data-movement bottlenecks in modern neural network accelerators. By performing matrix–vector multiplications directly within memory arrays in the analog domain, ACIM achieves excellent energy efficiency and throughput [1]. Among different implementation types, charge-domain ACIM (CD-ACIM) emerged as the most promising architecture that can be manufactured with today’s technologies, due to its high linearity, noise robustness, and scaling properties [2].

Despite these advances, the security implications of ACIM hardware remain largely unexplored. Previous research analyzed the susceptibility of resistive random-access memory based ACIM on power side-channel analysis, but relied purely on simulated power traces, where all non-idealities were precisely controlled [3], [4]. *This work presents the first power side-channel analysis performed on measured power traces from a CD-ACIM chip. In a proof-of-concept experiment in which a CNN inference with private input image is attacked, the private image can be reconstructed with high accuracy. This motivates the practical relevance of security-aware design methodology for next-generation ACIM hardware.*

A. Analog Compute-in-Memory Macro

The evaluated accelerator is a CD-ACIM macro manufactured in 28 nm CMOS, see Fig. 1. It implements charge-domain matrix–vector multiplication using a capacitive voltage-divider in switched-capacitor operation mode, inspired by the PIMCA macro [5]. The logical matrix size is 256 rows \times 32 columns.

Acknowledgment: This work was supported by the German National Science Foundation (DFG) Project SeMSiNN (OR 245/21-1, PO 1220/20-1) within the Priority Program SPP 2253 Nano Security.

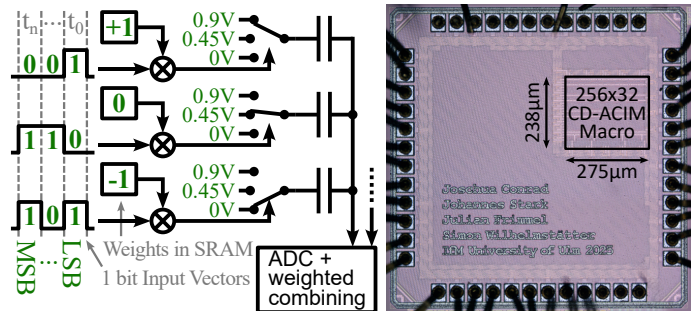


Fig. 1. Working principle of CD-ACIM (left) and photo of the 28nm CMOS macro (right).

An input vector contains 256 1 bit elements, and a weight matrix contains 32 vectors, each of which contains signed 6 bit operands (1 sign bit + 5 magnitude bits, so $5 \cdot 32 = 160$ physical summation columns). In every computation cycle, $256 \cdot 32 \cdot 5$ ternary results are computed and applied to the voltage dividers for analog summation along each column. The results of each of the 5 physical columns belonging to the same multi-bit weight vector are added with the respective binary weighting towards the multi-bit results. Multi-bit input vectors are processed bit-serial, from the LSB-vector to the MSB-vector, and the individual results are again combined with the respective binary weighting. The same supply voltages as in [5] are used to represent the ternary results of individual bitwise multiplications, namely $0\text{ V} \cong -1$, $0.45\text{ V} \cong 0$ and $0.9\text{ V} \cong +1$. Each voltage divider is pulled to 0.45 V during the reset phase. The macro processes 1 bit input vectors at a rate of 200 MHz.

B. Threat Model

As ACIM macros target applications in energy-constrained environments, they are most commonly deployed in embedded systems. An attacker is expected to have access to a sample instance that can be used to gain knowledge about the macro, e.g., by reverse-engineering the raw silicon die. It is therefore assumed that the macro’s matrix size and the bit-serial processing scheme are known. In our attack scenario, the attacker only needs access to the external supply pins of a packaged chip to collect power traces.

C. Power Trace Measurement Setup

Power traces are measured for the two supply voltage domains 0.45 V and 0.9 V using a $10\ \Omega$ shunt resistor on each supply line. The AC voltage drop over the shunt resistor is fed into a 25 dB low noise amplifier with 2.5 GHz bandwidth,




Fig. 2. Examples of the measured power traces with the 1 bit input vector segments $s_{i,j}$ highlighted (red).

before it is sampled with a Tektronix MSO 70804 oscilloscope. The oscilloscope outputs 25 GS/s at 8 bit resolution with an 8 GHz bandwidth, greatly exceeding the 200 MHz chip clock to ensure that no relevant signal components are lost.

II. EXPERIMENTS AND OBSERVATION

We evaluate the side-channel leakage of the CD-ACIM macro using a magnetic resonance imaging (MRI) image classification task implemented with a VGG-16 CNN [6]. All convolution operations are performed on the macro. The *im2col* transformation is used to reformat the 2D convolution into a sequence of matrix-vector multiplications. The generated input vectors are zero-padded to a length of 256 to match the input dimension of the macro. The weights of the CNN are extracted from a pretrained VGG-16 and quantized to signed 6 bit resolution before being programmed into the macro. The input images are grayscale MRI scans with 8 bit pixel resolution. Each multi-bit input vector is split into 8 1 bit input activation vectors, which the macro processes in bit-serial fashion from LSB to MSB.

Our goal is to reconstruct the input image solely from the measured power traces. We therefore target the power trace stemming from the inference of the first layer. The network’s first layer uses a 3×3 convolution filter with stride 1 and padding 1, so the height H and width W of the result map after the first layer stay unchanged. We measure the power traces for both supply domains: 0.9 V domain (driving +1 ternary results) and 0.45 V domain (reset bias and 0 results). Each power trace is evaluated separately.

For the evaluation, the full trace from the first layer’s inference is split into $H \cdot W \cdot 8$ many segments $s_{i,j}$, $i \in [0, H \cdot W - 1]$, $j \in [0, 7]$. The samples from each segment correspond to the current that was drawn from supply while the macro processed a 1 bit input activation vector for output pixel position i and input activation magnitude j (see Fig. 2). On average, a segment contains 125 oscilloscope samples. For each segment $s_{i,j}$ (i.e., 1 bit vector), we compute a scalar score $\tilde{a}_{i,j}$ that estimates the ternary result activity inside the macro (i.e., how many results were +1/-1 for this 1b vector-matrix multiplication). The score serves as a proxy for the share of non-zero bits in the 1 bit input activation vector that is processed during the segment $s_{i,j}$. It is computed as the maximum of absolute values over all samples in a segment, normalized to a range of $[0, 1]$ (cf. eq. (1) and eq. (4)).

For each pixel index i , the 8 scalar scores $\tilde{a}_{i,0}.. \tilde{a}_{i,7} \in [0, 1]$ associated with each 8 bit input vector are recombined to an

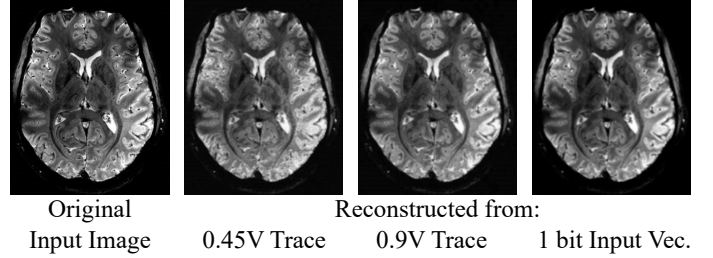


Fig. 3. Overview of the images: original input image [8] and reconstructed from the measured power traces and from the ideal 1 bit input activation vectors.

TABLE I
MSSIM SCORES FOR THE RECONSTRUCTED IMAGES FROM FIG. 3.

	1 bit Input Vec.	0.45 V Trace	0.9 V Trace
Input Image	0.92	0.68	0.71
1 bit Input Vec.	1.00	0.74	0.79

8 bit scalar value b_i using binary weighted summation (cf. eq. (5)). This value b_i is an estimate for each pixel’s brightness, and the total set of brightness scores $\{b_0..b_{H \cdot W - 1}\}$ is the flattened, reconstructed grayscale image.

$$a_{i,j} = \max_{125 \text{ Samples}} (|s_{i,j}|), \quad (1)$$

$$a_{\max} = \max_{i,j} (a_{i,j}), \quad (2)$$

$$a_{\min} = \min_{i,j} (a_{i,j}), \quad (3)$$

$$\tilde{a}_{i,j} = \frac{a_{i,j} - a_{\min}}{a_{\max} - a_{\min}} \in [0, 1], \quad (4)$$

$$b_i = \sum_{j=0}^7 2^j \cdot \tilde{a}_{i,j} \in [0, 255]. \quad (5)$$

The resulting brightness estimates b_i , rearranged as 2D grayscale image, for both power traces are shown in Fig. 3. As an ideal reference, the reconstructed image from the true shares of non-zero elements in each 1 bit input activation vector is also shown (Fig. 3, right). It can be seen that the reconstructed images are different but still cover all characteristic structures of the input image. To quantify the similarity of the reconstructed images, the MSSIM [7] scores are given in Table I. The scores are substantially above 0 for all cases, and the result for 0.9 V outperforms the result for 0.45 V.

These results practically confirm for the first time that CD-ACIM leaks sufficient information through its dynamic supply current to enable accurate input recovery, a serious security concern for real-world applications.

III. CONCLUSION

We presented the first measured power side-channel analysis on a CD-ACIM accelerator and showed that supply power traces alone enable accurate reconstruction of private MRI input images under real-world CNN workloads. These results demonstrate that CD-ACIM architectures leak exploitable information through power side-channels and highlight the need for dedicated countermeasures in future designs.

REFERENCES

- [1] J. Lee, B. Zhang, and N. Verma, "A Switched-Capacitor SRAM In-memory Computing Macro with High-precision, High-efficiency Differential Architecture," in *2024 IEEE European Solid-State Electronics Research Conference (ESSERC)*, pp. 357–360, 2024.
- [2] Y. He, X. Hu, H. Jia, and J.-s. Seo, "SRAM- and eDRAM-Based Compute-in-Memory Designs, Accelerators, and Evaluation Frameworks: Macro-Level and System-Level Optimization and Evaluation," *IEEE Solid-State Circuits Magazine*, vol. 17, no. 2, pp. 39–48, 2025.
- [3] Z. Wang, Y. Wu, Y. Park, S. Yoo, X. Wang, J. K. Eshraghian, and W. D. Lu, "PowerGAN: A Machine Learning Approach for Power Side-Channel Attack on Compute-in-Memory Accelerators," *Advanced Intelligent Systems*, vol. 5, no. 12, p. 2300313, 2023.
- [4] Z. Wang, F.-h. Meng, Y. Park, J. K. Eshraghian, and W. D. Lu, "Side-channel attack analysis on in-memory computing architectures," *IEEE Trans. Emerg. Topics Comput.*, pp. 1–13, 2023.
- [5] B. Zhang, S. Yin, M. Kim, J. Saikia, S. Kwon, S. Myung, H. Kim, S. J. Kim, J.-S. Seo, and M. Seok, "PIMCA: A Programmable In-Memory Computing Accelerator for Energy-Efficient DNN Inference," *IEEE Journal of Solid-State Circuits*, vol. 58, no. 5, pp. 1436–1449, 2023.
- [6] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," *3rd International Conference on Learning Representations (ICLR 2015)*, 2015.
- [7] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [8] CEA, "A world premiere: the living brain imaged with unrivaled clarity thanks to the world's most powerful MRI machine," Apr. 2024.