



Delft University of Technology

## MANTIS

### A Mixed-Signal Near-Sensor Convolutional Imager SoC Using Charge-Domain 4b-Weighted 5-to-84-TOPS/W MAC Operations for Feature Extraction and Region-of-Interest Detection

Lefebvre, Martin; Bol, David

#### DOI

[10.1109/JSSC.2024.3484766](https://doi.org/10.1109/JSSC.2024.3484766)

#### Publication date

2024

#### Document Version

Final published version

#### Published in

IEEE Journal of Solid-State Circuits

#### Citation (APA)

Lefebvre, M., & Bol, D. (2024). MANTIS: A Mixed-Signal Near-Sensor Convolutional Imager SoC Using Charge-Domain 4b-Weighted 5-to-84-TOPS/W MAC Operations for Feature Extraction and Region-of-Interest Detection. *IEEE Journal of Solid-State Circuits*, 60(3), 934-948.  
<https://doi.org/10.1109/JSSC.2024.3484766>

#### Important note

To cite this publication, please use the final published version (if applicable).  
Please check the document version above.

#### Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

#### Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.  
We will remove access to the work immediately and investigate your claim.

***Green Open Access added to TU Delft Institutional Repository***

***'You share, we take care!' - Taverne project***

**<https://www.openaccess.nl/en/you-share-we-take-care>**

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

# MANTIS: A Mixed-Signal Near-Sensor Convolutional Imager SoC Using Charge-Domain 4b-Weighted 5-to-84-TOPS/W MAC Operations for Feature Extraction and Region-of-Interest Detection

Martin Lefebvre<sup>ID</sup>, *Member, IEEE*, and David Bol<sup>ID</sup>, *Senior Member, IEEE*

**Abstract**—Recent advances in artificial intelligence (AI) have prompted the search for enhanced algorithms and hardware to support the deployment of machine learning (ML) at the edge. More specifically, in the context of the Internet of Things (IoT), vision chips must be able to fulfill the tasks of low to medium complexity, such as feature extraction (FE) or region-of-interest (RoI) detection, with a sub-mW power budget imposed by the use of small batteries or energy harvesting. Mixed-signal vision chips relying on in- or near-sensor processing have emerged as an interesting candidate because of their favorable tradeoff between energy efficiency (EE) and computational accuracy compared with digital systems for these specific tasks. In this article, we introduce a mixed-signal convolutional imager system-on-chip (SoC) codenamed MANTIS, featuring a unique combination of large  $16 \times 16$  4b-weighted filters, operation at multiple scales, and double sampling, well suited to the requirements of medium-complexity tasks. The main contributions are (i) circuits called DS3 units combining delta-reset sampling (DRS), image down-sampling (DS), and voltage downshifting and (ii) charge-domain multiply-and-accumulate (MAC) operations based on switched-capacitor (SC) amplifiers and charge sharing in the capacitive DAC of the successive-approximation (SAR) ADCs. MANTIS achieves peak EEs normalized to 1b operations of 4.6 and 84.1 TOPS/W at the accelerator and SoC levels, while computing feature maps (fmaps) with a root-mean-square error (RMSE) ranging from 3 to 11.3%. It also demonstrates a face RoI detection with a false negative rate (FNR) of 11.5%, while discarding 81.3% of image patches and reducing the data transmitted off chip by  $13\times$  compared with the raw image.

**Index Terms**—Charge domain, CMOS image sensor (CIS), convolutional neural network (CNN), feature extraction (FE), mixed signal, multiply-and-accumulate (MAC) operations, near sensor, region-of-interest (RoI) detection, system-on-chip (SoC).

## I. INTRODUCTION

RECENT years have seen artificial intelligence (AI) rise as a key component of numerous engineered systems,

Received 26 June 2024; revised 5 September 2024 and 14 October 2024; accepted 18 October 2024. Date of publication 11 November 2024; date of current version 26 February 2025. This article was approved by Associate Editor Yoonmyung Lee. (*Corresponding author: Martin Lefebvre.*)

Martin Lefebvre is with the Institute of Information and Communication Technologies, Electronics and Applied Mathematics, Université catholique de Louvain, 1348 Louvain-la-Neuve, Belgium, and also with the Department of Microelectronics, Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: m.lefebvre@tudelft.nl).

David Bol is with the Institute of Information and Communication Technologies, Electronics and Applied Mathematics, Université catholique de Louvain, 1348 Louvain-la-Neuve, Belgium.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/JSSC.2024.3484766>.

Digital Object Identifier 10.1109/JSSC.2024.3484766

0018-9200 © 2024 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See <https://www.ieee.org/publications/rights/index.html> for more information.

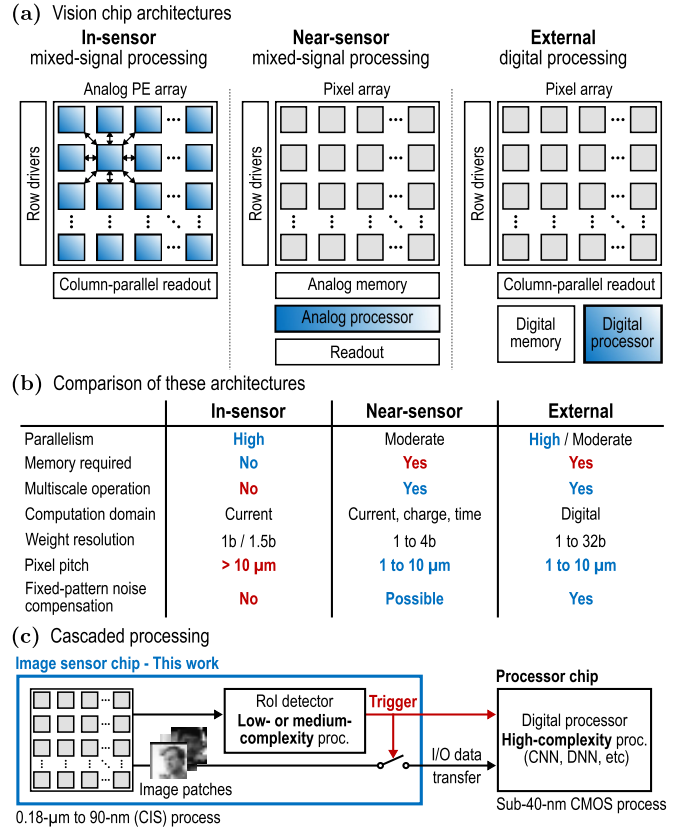


Fig. 1. (a) Vision chip architectures ranging from mixed-signal processing in or near the pixel array to conventional digital processing outside of it. (b) Strengths and limitations of these architectures. (c) Envisioned system based on a cascaded processing scheme similar to [4], in which only relevant image patches are transmitted from the image sensor to the digital processor.

reaching an unprecedented level of pervasiveness at the applications level. Among them, the Internet of Things (IoT) has elicited a particular interest as the large amount of data generated by sensor nodes calls for the development of specialized machine learning (ML) algorithms and hardware to efficiently process data at the edge, a concept coined as edge AI or tiny ML. More specifically, in the context of vision sensors, edge devices must be able to solve vision tasks of low to medium complexity, e.g., feature extraction (FE) and region-of-interest (RoI) detection, within a sub-mW power budget, as IoT nodes are often supplied by limited-capacity

batteries. Mixed-signal vision chips have, thus, emerged as a suitable candidate, since they outperform digital chips in terms of energy efficiency (EE) while maintaining a sufficient computational accuracy. This improved EE also stems from a reduced number of analog-to-digital (A2D) conversions compared with digital implementations, leading to energy and area savings.

Mixed-signal vision chip architectures can be divided into two main categories, namely, in-sensor [1], [2], [3] and near-sensor [4], [5], [6], [7], [8] vision chips, respectively, implemented with analog processing elements (PEs) inside or in the periphery of the pixel array. A third category, referred to as hybrid vision chips [9], [10], is not represented in Fig. 1(a) but simply combines elements from both categories. On the one hand, in-sensor vision chips are massively parallel and do not require any memory, be it analog or digital. However, connections between pixel-level PEs are usually local and limited to neighboring pixels, hampering the calculation of image-level features and, thereby, the operation at multiple spatial scales. In addition, pixel-level PEs also lead to a relatively large pixel pitch above  $10\ \mu\text{m}$ . At last, in-sensor processing is often limited to low-complexity tasks, as it relies on binary (1b) or ternary (1.5b) weights and on raw [1], [2], [3] or amplified [9] photocurrents subject to significant fixed-pattern noise (FPN), i.e., local mismatch between pixel responses. On the other hand, near-sensor and hybrid vision chips usually present a decreased throughput compared with in-sensor ones, but are better suited to the execution of medium-complexity tasks because of the use of large-size 1.5b Haar-like filters [4], [5], [9], [10] or to an increased 4b filter weight resolution [7], [8]. They make use of conventional pixel structures, such as a three- or four-transistor (3T or 4T) active pixel sensor (APS) or pulse-width-modulated (PWM) digital pixel sensor, which are compatible with double sampling techniques to compensate FPN. The 3T/4T pixels, respectively, use rolling and global shutters and can either rely on a voltage-based readout with a source follower (SF) or on a time-based one with a PWM structure, allowing to reduce the supply voltage without degrading the output dynamic range. Besides, near-sensor and hybrid vision chips can operate at multiple spatial scales because of image downsampling (DS) [4], [5] or filter dilation [9]. Finally, an analog memory generally based on capacitors (caps) is required to store a few rows of the image, ultimately leading to power and/or area overheads. Nevertheless, existing works fall short of preserving EE while simultaneously supporting medium-complexity tasks, which require sufficient computational accuracy brought by FPN-compensated inputs and increased weight resolution, as well as multiscale operation and large filters for tasks as RoI detection.

In this work, we present a mixed-signal near-sensor convolutional imager system-on-chip (SoC) codenamed MANTIS, fabricated in United Microelectronics Corporation (UMC)  $0.11\text{-}\mu\text{m}$  CMOS technology and supporting both FE and RoI detection. It includes two main contributions providing an effective answer to the aforementioned limitations of the existing vision chips. First, circuits called DS3 units combine three

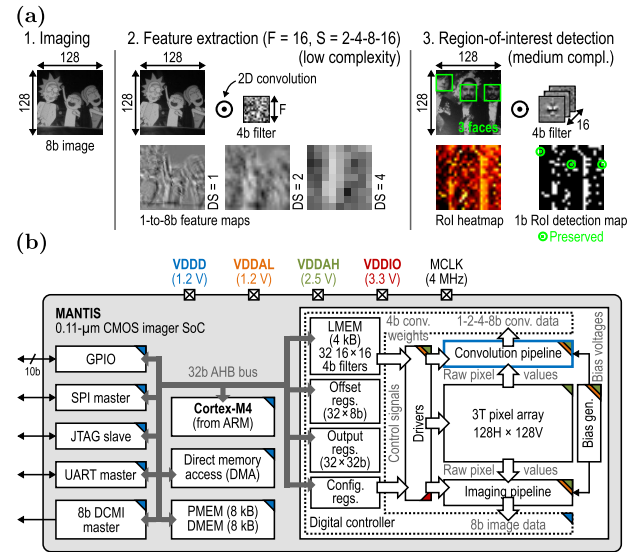


Fig. 2. MANTIS CMOS imager SoC (a) modes of operation and (b) architecture, detailing the different blocks in the digital core and image sensor analog core with their respective power domains.

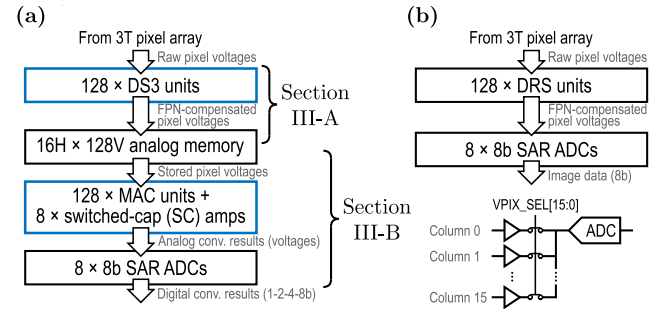


Fig. 3. Block diagram of the (a) convolution and (b) imaging pipelines.

operations, which can be abbreviated as DS, namely, double sampling, to mitigate the impact of FPN, voltage downshifting, to reduce the voltage level from the pixel array to the convolution processor, and image downsampling, to allow for multiscale operation. Second, a mixed-signal convolution processor implements 4b-weighted multiply-and-accumulate (MAC) operations in the charge domain, based on a modified switched-capacitor (SC) amplifier structure to compute the partial sum (psum) of a row of an image patch, and on a charge sharing operation in the capacitive digital-to-analog converter (CDAC) of the following successive-approximation (SAR) analog-to-digital converter (ADC) to aggregate psums of different rows. In our vision, MANTIS would be used as the first stage of a cascaded processing system [Fig. 1(c)] supporting low- to medium-complexity processing tasks, while a high-complexity processing based on convolutional or deep neural networks (CNNs or DNNs) would be executed by a digital processor. The major benefits of such a system are to limit the amount of I/O data transfers from the image sensor to the digital processor and to only dedicate energy to the processing of relevant data. This article extends our conference paper [11] by providing a more in-depth description of the circuits

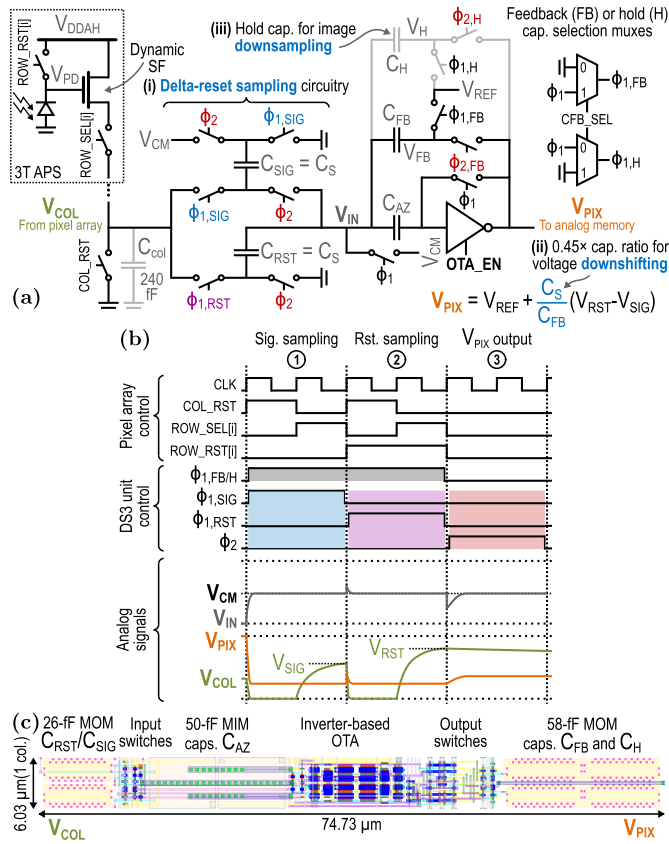


Fig. 4. (a) Schematic, (b) timing diagram, and (c) 90°-rotated layout of a single column-parallel DS3 unit.  $V_{CM} = 1.2$  V and  $V_{REF} = 0.6$  V in (a).

constituting the convolution pipeline, highlighted in Fig. 2, as well as additional experimental results. The remainder of this article is organized as follows. First, Section II describes the architecture of the SoC. Then, Section III discusses the design and implementation of the proposed mixed-signal convolution pipeline, while Section IV presents the measurement results of the SoC. Finally, Section V compares this work with the state of the art, and Section VI offers some concluding remarks.

## II. SYSTEM-ON-CHIP DESCRIPTION

### A. Architecture

This section describes the modes of operation and architecture of MANTIS CMOS imager SoC, respectively, depicted in Fig. 2(a) and (b). Three modes of operations are supported. First, the imaging mode produces 8b  $128 \times 128$  images, which are necessary to thoroughly compare the mixed-signal on-chip execution of the convolution operations, subject to analog nonidealities, to an ideal software baseline. Second, FE can be performed using 2-D convolution operations between the image and 4b-weighted filters of fixed size  $F = 16$ . All parameters are programmable, with the filter stride  $S$  and the DS factor  $DS$ , respectively, taking any power-of-two value between 2 and 16, and between 1 and 4, and the number of filters ranging from 1 to 32. This mode generates feature maps (fmaps) with a programmable power-of-two resolution

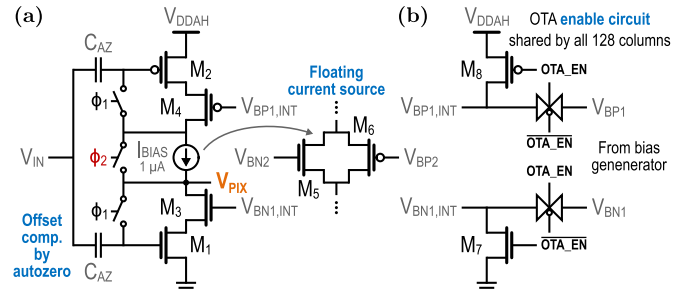


Fig. 5. Schematic of (a) the inverter-based OTA proposed in [12] and (b) the enable circuit shared by all 128 column-parallel DS3 units.

between 1 and 8 bits. Finally, an RoI detection mode supports the comparison of fmap values with a different threshold for each filter directly in the SAR ADCs. These thresholds are implemented as offsets modifying the fmap values. In this last mode of operation, 1b fmaps are created by the imager, which are subsequently combined to yield an RoI heatmap and 1b detection map [Fig. 2(a)] for the detection of faces.

Furthermore, the SoC architecture revolves around a Cortex-M4 central processing unit (CPU) from ARM, embedding a mixed-signal image sensor macro. First, regarding the *digital part* of the SoC, efficient data transfers are supported by a direct memory access (DMA) peripheral allowing to move data through the advanced high-performance (AHB) bus from the imager output registers to a master digital camera interface (DCMI), which then transmits this data off chip in an 8b-parallel fashion. Moving on to the image sensor macro [Fig. 2(b) right], it includes several configuration registers for the parameters of the convolution operations discussed here-above, among others, a 4-kB local SRAM memory, denoted as LMEM, which can store up to 32 4b  $16 \times 16$  filters and 32 8b registers for the corresponding thresholds or offsets in RoI mode. These registers impact the behavior of the digital controller piloting the analog core of the imager. Next, the *analog part* of the SoC relies on a 3T  $128 \times 128$  pixel array with two readout pipelines: (i) a convolution pipeline supporting both FE and RoI detection modes and (ii) an imaging pipeline used in imaging mode. It also includes a bias generation circuit used by both pipelines. The digital core is supplied at 1.2 V, while the analog circuitry relies on two supplies at 1.2 and 2.5 V, respectively.

### B. Convolution Pipeline

In the convolution pipeline [Fig. 3(a)], raw pixel voltages go through 128 column-parallel DS3 units, which zero out the FPN by a double sampling technique known as delta-reset sampling (DRS). It consists in subtracting the signal voltage of a pixel from its reset voltage, thereby suppressing the impact of local mismatch on its output. In addition, DS3 units also perform image DS to support multiscale operation. The output voltage of DS3 units is then stored in an analog memory with a capacity of 16 rows. Next, the stored pixel values are employed as inputs to 128 MAC units, connected to eight SC amplifiers computing partial convolution results or psums in



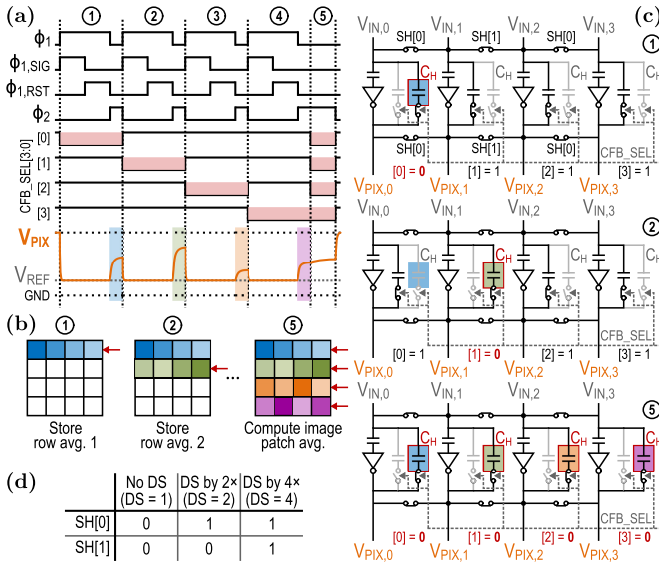


Fig. 6. Illustration of image DS by 4x with (a) the timing diagram, (b) the operation principle, and (c) the schematic of four DS3 units in four neighboring columns. (d) Connection of the switches shorting the inputs and outputs of inverter-based OTAs for different DS factors.

the analog domain, under the form of voltages. These psums are stored in the CDAC of the SAR ADCs following the SC amplifiers, before being aggregated by a charge sharing operation in the CDAC and digitized to produce convolution results, with the resolution of the produced fmaps being a power of two between 1 and 8 bits.

### C. Imaging Pipeline

In the imaging pipeline [Fig. 3(b)], DRS is used to mitigate FPN, as is done in DS3 units in the convolution pipeline. DRS units also implement voltage downshifting to adapt the 2.5-V signals from the pixel array to the 1.2-V input of the 8b SAR ADCs. Note that the outputs of column-parallel DRS units in 16 adjacent columns are multiplexed to a single ADC, leading to a total of eight ADCs to digitize a complete row.

## III. CIRCUIT IMPLEMENTATION AND DESIGN OF THE MIXED-SIGNAL CONVOLUTION PIPELINE

The objective of this section is to present the circuit implementation of the mixed-signal convolution pipeline and to provide insights regarding the design of the circuits it contains. To do so, this section is divided in two parts. Section III-A deals with the calculation of FPN-compensated downsampled pixel voltages and their storage in the analog memory and, consequently, covers the DS3 units and the analog memory. Section III-B discusses the charge-domain MAC operations and the digitization of the convolution results, and examines the MAC units, the SC amplifiers, and the SAR ADCs.

### A. Image Readout, Downsampling, and Storage

1) *DS3 Units for Image Readout and Downsampling*: Fig. 4 illustrates the operation of a column-parallel DS3 unit to read a single pixel value while performing DRS and voltage

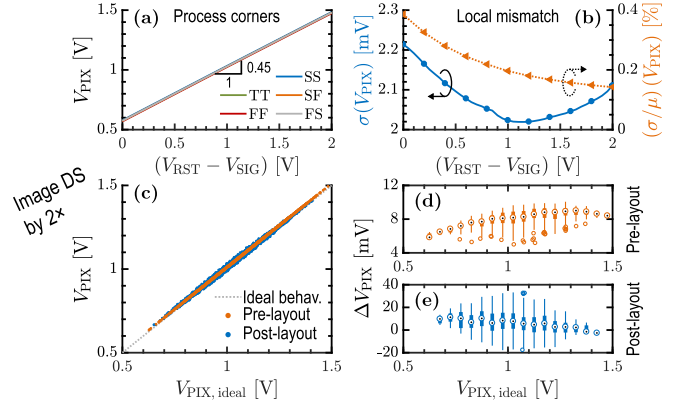


Fig. 7. (a)  $V_{PIX}$  in process corners and (b) variability of  $V_{PIX}$  for  $10^3$  MC simulations with local mismatch. Image DS by 2x for  $2 \times 10^3$  random combinations of inputs drawn from a uniform distribution, with (c) comparison between ideal and simulated results, and distributions of the error  $\Delta V_{PIX} = (V_{PIX} - V_{PIX,ideal})$  in (d) pre- and (e) post-layout simulations. All figures correspond to the TT 25 °C corner, except for (a), which covers all five process corners.

downshifting. This operation consists of three steps [Fig. 4(b)]. In step ①, the signal coming from the 3T APS, resulting from the discharge of the internal pixel node  $V_{PD}$  by the photocurrent during the exposure time, is read on the column voltage  $V_{COL}$ . To do so, we rely on the partial settling or dynamic SF readout from [13], which consists in resetting  $V_{COL}$  to ground using the  $COL\_RST$  switch, before enabling the SF during a finite amount of time ( $ROW\_SEL[i] = 1$ ). This readout is more energy-efficient than the conventional one using a current source at the bottom of the column, as it eliminates any static current consumption. It also presents an optimal settling time [13], which minimizes the variability of  $(V_{RST} - V_{SIG})$ , that we find to be  $0.5 \mu s$  in our design. At the end of step ①, the signal value has been sampled on the 26-fF MOM cap  $C_{SIG}$ . Then, in step ②, the pixel is reset, and the resulting value is sampled on  $C_{RST}$ , whose capacitance is the same as  $C_{SIG}$ . Finally, during step ③, these two caps are connected with opposite polarities, and their charges are dumped on a 58-fF MOM feedback cap  $C_{FB}$ , resulting in a voltage  $V_{PIX}$  in which  $(V_{RST} - V_{SIG})$  is multiplied by the capacitance ratio  $C_S/C_{FB} = 0.45$ . Thus, the operations of (i) DRS and (ii) voltage downshifting are realized. Moreover, in the schematic depicted in Fig. 4(a), switches connected to ground or  $V_{DD}$  are, respectively, implemented with single nMOS or pMOS, while other switches are transmission gates (TGs). They rely on 3.3-V I/O transistors to withstand the 2.5-V supply, with  $W = 0.25 \mu m$  and  $L = L_{min} = 0.34 \mu m$ , except for the TGs connected to  $V_{REF}$  and  $V_{CM}$  for which  $L = 0.68 \mu m$  to reduce leakage. Regarding the caps, they are chosen to ensure that local mismatch, noise, and layout parasitics have a minimal impact on the circuit behavior, but they could be downsized as long as the uncertainty and voltage attenuation remain within the specifications of the target algorithm. Further reduction could be achieved by accounting for these nonidealities in the training algorithm, as is done in [14] for in-memory computing (IMC). A similar design choice is made for other mixed-signal circuits in this work.

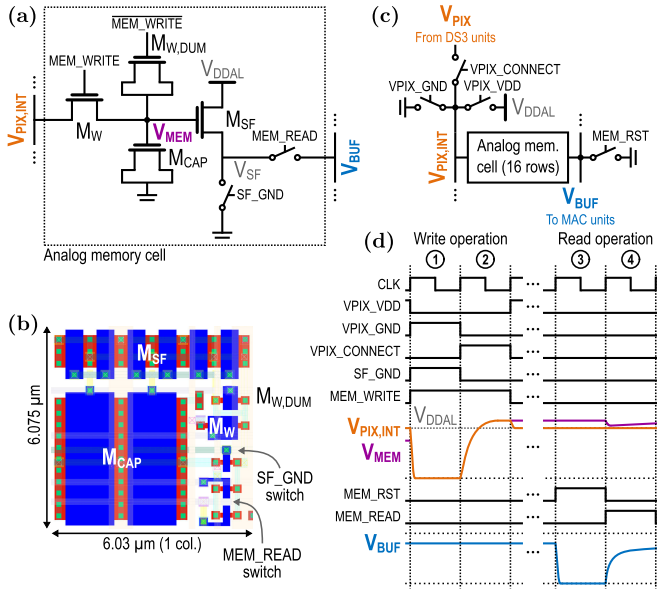


Fig. 8. (a) Schematic and (b) layout of an analog memory cell. (c) Connections at the input and output of a column of 16 row of memory cells and (d) timing diagram for write and read operations.

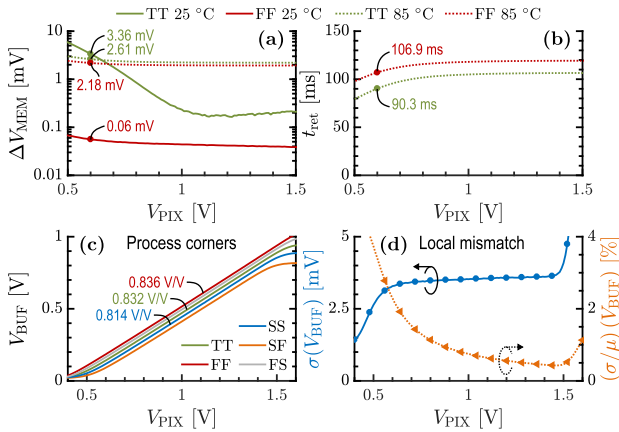


Fig. 9. (a) Voltage change of  $V_{MEM}$  after 100 ms in TT and FF, and at 25 °C and 85 °C. (b) Retention time  $t_{ret}$  in TT 85 °C and FF 85 °C (worst case).  $t_{ret}$  is defined as the time at which the initial voltage has changed by more than 2.35 mV, corresponding to half an LSB for a 1.2-V 8b ADC. At 25 °C, (c) transfer function from  $V_{PIX}$  to  $V_{BUF}$  in process corners and (d) variability of  $V_{BUF}$  for  $10^3$  MC simulations with local mismatch.

Besides, to compensate for the offset of the inverter-based operational transconductance amplifier (OTA) [12], it is put in autozero (AZ) during steps ① and ②. This corresponds to sampling on the two 50-fF MIM caps  $C_{AZ}$  the difference between the common-mode voltage  $V_{CM}$  and  $V_{GS}$  of  $M_{1-2}$  with a fixed 1- $\mu A$  bias current, imposed by the floating current source formed by  $M_{5-6}$  [Fig. 5(a)]. Moreover, a key feature to reach a high EE is the enable circuit shared by all amplifiers [Fig. 5(b)], implemented by, respectively, clamping bias voltages  $V_{BN1,INT}$  and  $V_{BP1,INT}$  to ground and  $V_{DDAH}$  and, thus, allowing to duty cycle DS3 units to save power.

Next, the image DS relies on a principle proposed in [6] and represented for a DS by 4 $\times$  in Fig. 6. As illustrated in

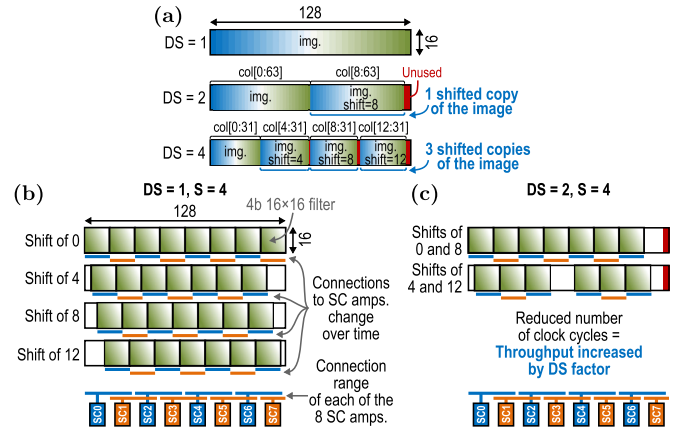


Fig. 10. (a) Storage pattern of the image into the analog memory for different DS factors. Filter striding for a convolution operation with a stride  $S=4$  (b) without image DS ( $DS=1$ ) and (c) with image DS by 2 $\times$  ( $DS=2$ ).

Fig. 6(c), the inputs and outputs of the OTAs in four adjacent columns are shorted together by switches whose configuration depends on the DS factor [Fig. 6(d)]. The average of each row of a 4 $\times$ 4 image patch is computed and stored in the hold cap  $C_H$  of one of the four columns, as shown in steps ①–④ [Fig. 4(a)–(c)]. Once all row averages have been computed, all  $C_H$  caps are simultaneously connected during step ⑤, and the resulting voltage is the average of row averages or, in other words, the average of the image patch. The proposed DS3 unit fits into the 6.03- $\mu m$  pixel pitch and occupies 74.73  $\mu m$  in height [Fig. 4(c)], a dimension that could be further reduced if transistors could be placed below MIM caps  $C_{AZ}$ .

This circuit is robust to process, voltage, and temperature (PVT) variations due to its ratiometric nature, as long as the OTA is designed to operate in the relevant corners. Thus, this article does not aim at providing an exhaustive characterization of the proposed circuits in PVT corners, but focuses on their main performance and limitations. Post-layout simulations in Fig. 7 confirm the independence with respect to process [Fig. 7(a)] and show that the  $\sigma$  and  $\sigma/\mu$  of  $V_{PIX}$  due to local mismatch for a single DS3 unit are respectively below 2.2 mV and 0.4% across the input range. Regarding the output voltage noise, a theoretical expression is given by  $\bar{v}_n = (C_S/C_{FB})(2kT/C_S)^{1/2}$ , where  $k$  is Boltzmann's constant and  $T$  the absolute temperature. At 25 °C,  $\bar{v}_n = 0.25$  mV, and its impact is significantly lower than that of mismatch. In addition, the performance of DS is evaluated for DS by 2 $\times$  in Fig. 7(c)–(e), for  $2 \times 10^3$  combinations of input voltages drawn from a uniform distribution. The  $\sigma$  of the error  $\Delta V_{PIX}$  increases from less than 1 mV pre-layout to approximately 10 mV post-layout, as highlighted in Fig. 7(d) and (e), due to capacitive coupling among nodes  $V_{IN}$ ,  $V_{PIX}$ , and  $V_H$  [Fig. 4(a)], which could be reduced by improving the layout.

2) *Analog Memory for Image Storage*: The schematic and operation of the analog memory are described in Fig. 8. A memory cell with a structure close to [4] and [5] [Fig. 8(a)] consists of a 32-fF MOS cap  $M_{CAP}$ , an access transistor  $M_W$  with a dummy transistor  $M_{W,DUM}$  with half the length,

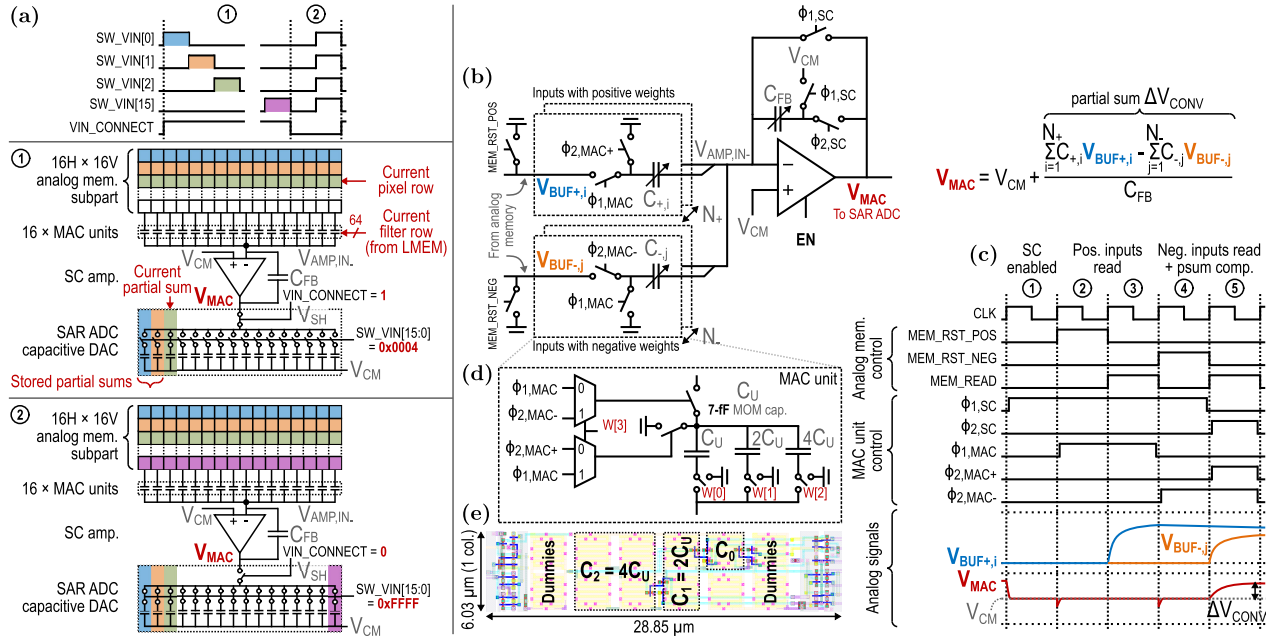


Fig. 11. (a) Principle of the convolution operation between a  $16 \times 16$  image patch and a 4b  $16 \times 16$  filter. (b) Schematic of the SC amplifier realizing the MAC operations between a single row of the image patch and filter, with (c) the corresponding timing diagram. The common-mode voltage  $V_{CM}$  is equal to  $V_{DDAL}/2 = 0.6$  V. (d) Detailed schematic and (e)  $90^\circ$ -rotated layout of one of the 16 MAC units connected to the SC amplifier.

to compensate for the charge injection of  $M_W$ , and an SF  $M_{SF}$  employed in a dynamic fashion for its reduced mismatch of  $V_{BUF}$  and decreased static power consumption, similar to the pixel readout in Section III-A. This memory cell occupies a silicon area of  $6.03 \times 6.075 \mu m$  [Fig. 8(b)], close to that of a pixel. To read or write a cell located within a column of the analog memory [Fig. 8(c) and (d)], several switches are used to connect the column internal voltage  $V_{PIX,INT}$  to the output of the DS3 units, or to ground/ $V_{DDAL}$ . During a write operation, in step ①,  $V_{PIX,INT}$ ,  $V_{MEM}$  and  $V_{SF}$  are grounded to overwrite the memory cell content without any impact from previously stored values. Then, in step ②,  $V_{MEM}$  is driven to  $V_{PIX}$  by the DS3 unit connected to the column, before disconnecting  $V_{MEM}$  from  $V_{PIX,INT}$ . During a read operation,  $V_{BUF}$  is first reset to ground in step ③, before reading the memory cell by partial settling in step ④.

When the memory is not written or read, the retention of the memory cells needs to be maximized in the worst case corner, here FF 85  $^\circ C$ . To do so, we minimize the leakage of the access transistor  $M_W$  by implementing it with a 3.3-V I/O nMOS with  $W = 0.18 \mu m$  and  $L = 1 \mu m$  and, additionally, by driving  $V_{PIX,INT}$  to  $V_{DDAL} = 1.2$  V in retention to limit the  $V_{DS}$  of  $M_W$  and further reduce the leakage, given that  $V_{PIX}$  approximately ranges from 0.6 to 1.5 V [Fig. 7(a)]. Continuing with the post-layout characterization of the memory, Fig. 9(a) highlights that the typical voltage change of the stored voltage after 100 ms is, respectively, 2.61 and 2.18 mV in the TT and FF process corners at 85  $^\circ C$ , while Fig. 9(b) indicates the retention times of 90.3 and 106.9 ms in the same conditions, the retention time being defined as a change of  $\pm LSB/2$  with respect to the initially stored voltage, i.e.,  $\pm 2.35$  mV for a 1.2-V supply and an 8b resolution. In addition, in Fig. 9(a), the linear increase of  $\Delta V_{MEM}$  in TT 25  $^\circ C$  for  $V_{PIX} < 1$  V can

be explained by a slow transient of  $V_{SF}$  lightly affecting  $V_{MEM}$  through capacitive coupling. In Fig. 9(c), we observe that the transfer function of the SF has a slope  $A_{SF}$  below 1 V/V due to the body effect resulting from  $M_{SF}$ 's body being grounded, and that it is impacted by variations of  $M_{SF}$ 's threshold voltage in process corners even though the slope remains around 0.83 V/V. Finally,  $V_{BUF}$  has a  $\sigma$  around 3.5 mV in the usable part of the input range, corresponding to a maximum  $\sigma/\mu$  of 2.3% for  $V_{PIX} = 0.6$  V, while in comparison, the output noise  $\bar{v}_n = A_{SF}(kT/C_{MEM})^{1/2}$  equal to 0.3 mV at 25  $^\circ C$  is negligible. Future designs could compensate the SF mismatch by making use of an OTA-based feedback loop to write the analog memory, as proposed by Seo et al. [15].

## B. Charge-Domain Multiply-and-Accumulate Operations

1) Operation Principle: When no DS is applied to the image, the columns of the pixel array match that of the analog memory in a one-to-one fashion, and the whole width of the analog memory is used to store the image [Fig. 10(a)]. The convolution operation is, thus, performed between several replicas of the 4b  $16 \times 16$  filter and different image patches without overlap [Fig. 10(b)]. As the filter is shifted to the right, the connections between the analog memory and the eight SC amplifiers are modified over time to follow the movement of the filter. However, when a DS by  $2 \times$  is applied, the image is only 64-columns wide, so the first half of the analog memory stores the downsampled image, while the second half stores a version of the image shifted by eight columns to the left [Fig. 10(a)]. This routing from the outputs of the DS3 units to the analog memory is ensured by switches changing the connections depending on the DS factor. When computing the convolution operation, this storage pattern of



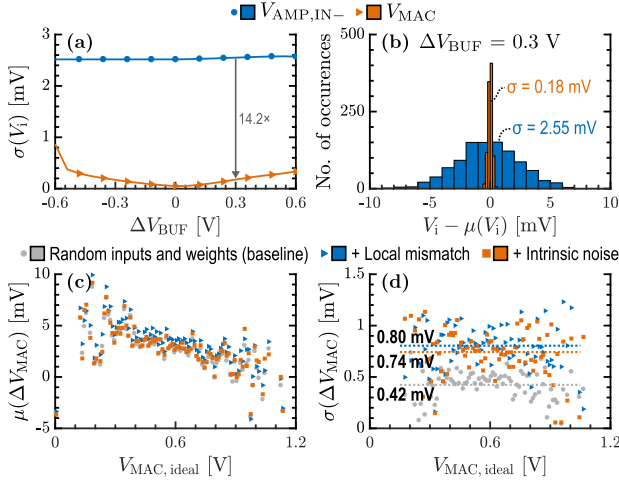


Fig. 12. (a) Standard deviation and (b) distribution of  $V_{AMP,IN-}$  and  $V_{OUT}$ , for  $10^5$  MC simulations with local mismatch. (b) Histograms for  $\Delta V_{BUF} = 0.3$  V. (c) Mean and (d) standard deviation of the error  $\Delta V_{MAC} = (V_{MAC} - V_{MAC, ideal})$  for  $10^4$  random combinations of inputs and weights without mismatch and noise, with only local mismatch, and with only intrinsic noise. All figures correspond to the TT 25 °C corner.

the image into the analog memory allows to improve throughput, by executing in parallel operations corresponding to two different shifts of the image in the execution without DS [Fig. 10(c)]. The throughput is thereby increased by the DS factor, here  $2\times$ . The same reasoning holds for a DS by  $4\times$  for which three shifted versions of the image are used, as shown in Fig. 10(a).

2) *Switched-Cap Amplifiers for Multiplication:* We now zoom in on the convolution operation between a  $16\times 16$  image patch and a 4b  $16\times 16$  filter, computed by the process depicted in Fig. 11 using an SC amplifier. Phase ① of this process, presented in Fig. 11(a), consists in successively computing the psums resulting from the convolution of a row of pixels stored in the analog memory with the corresponding row of 4b filter weights stored in the LMEM. Each psum is stored in a 16th of the SAR ADC CDAC, until all psums have been computed. In phase ②, the CDAC is disconnected from the SC amplifier ( $V_{IN\_CONNECT} = 0$ ), and all caps storing psums are shorted together to compute the final convolution result by charge sharing on node  $V_{SH}$ .

Going one step further, the psum of a row is computed by the SC amplifier circuit drawn in Fig. 11(b), whose timing diagram is detailed in Fig. 11(c). In step ①, the OTA, based on a two-stage Miller architecture, is enabled, and its feedback is activated. As for the DS3 units, power gating the OTA is a key feature to save energy and improve the EE of the accelerator. Then, steps ② and ③, respectively, consist in resetting the columns of the analog memory corresponding to positive-weighted inputs and in reading these inputs from the analog memory. Next, in step ④, the columns of the analog memory corresponding to a negative weight are reset, while connecting them to the input of the corresponding MAC units. Finally, during step ⑤, the negative-weighted inputs are read, and charges are dumped on node  $V_{AMP,IN-}$  by caps  $C_{+,i}$  and  $C_{-,j}$ , yielding an output voltage  $V_{MAC}$  containing the psum  $\Delta V_{CONV}$  referred to  $V_{CM}$ . The formula given in Fig. 11(b)




Fig. 13. Leakage current through TGs in the MAC unit can lead to variability of  $V_{MAC}$  due to global process variations. (a) Simplified schematic of the MAC unit with transistor-level switch implementation, illustrating the origin of this leakage, and (b) standard deviation of the error  $\Delta V_{MAC}$  for  $10^4$  random combinations of inputs and weights with local mismatch and global process variations, for TGs realized with LVT core devices with  $L = 120$  and  $240$  nm, or HVT core ones with  $L = 120$  nm.

can be intuitively understood by noticing that the inputs of caps  $C_{+,i}$  are applied when  $\phi_{1,SC} = 1$ , while the inputs of caps  $C_{-,j}$  are applied when  $\phi_{2,SC} = 1$ . The inputs associated with  $C_{+,i}$  follow the behavior of a non-inverting SC amplifier, while those associated with  $C_{-,j}$  follow that of an inverting one, thus explaining the formula for  $V_{MAC}$ . Despite the fact that the computation is performed in the mixed-signal domain and suffers from analog nonidealities, the proposed structure features several properties ensuring the robustness of the computation. (i) It has a single-ended output, which does not rely on intermediate differential voltages, avoiding an incorrect result when the differential voltage is small, but the common mode is large and potentially subject to saturation. This is an issue encountered in previous charge-domain near-sensor architectures [5]. (ii) The proposed structure is ratiometric and robust to PVT variations. (iii) It is not impacted by the statistical offset of the OTA because of the offset-insensitive switching scheme. Indeed, the charges at node  $V_{AMP,IN-}$  are

$$Q_1 = - \sum_{i=1}^{N_+} C_{+,i} (V_{BUF+,i} - V_{AMP,IN-}) - \sum_{j=1}^{N_-} C_{-,j} (-V_{AMP,IN-}) + C_{FB} (V_{AMP,IN-} - V_{CM}) \quad (1)$$

for  $\phi_{1,SC} = 1$ , with  $N_+$  and  $N_-$  the number of positive and negative-weighted inputs, and

$$Q_2 = - \sum_{j=1}^{N_-} C_{-,j} (V_{BUF-,j} - V_{AMP,IN-}) - \sum_{i=1}^{N_+} C_{+,i} (-V_{AMP,IN-}) + C_{FB} (V_{AMP,IN-} - V_{MAC}) \quad (2)$$

for  $\phi_{2,SC} = 1$ . Interestingly, the resulting expression for  $V_{MAC}$  based on the conservation of charge at node  $V_{AMP,IN-}$ , i.e.,  $Q_1 = Q_2$ , does not depend on the value of  $V_{AMP,IN-}$  when  $\phi_{1,SC} = 1$  and, therefore, is independent of the OTA's offset.

Furthermore, the implementation of the 4b weights is given in Fig. 11(d), with the most-significant bit (MSB)  $W[3]$  corresponding to the sign bit and least-significant bits (LSBs)  $W[2:0]$  to the magnitude bits. The sign bit determines which signals control the connections at the input of the MAC unit, while the magnitude bits determine the number of 7-fF unitary

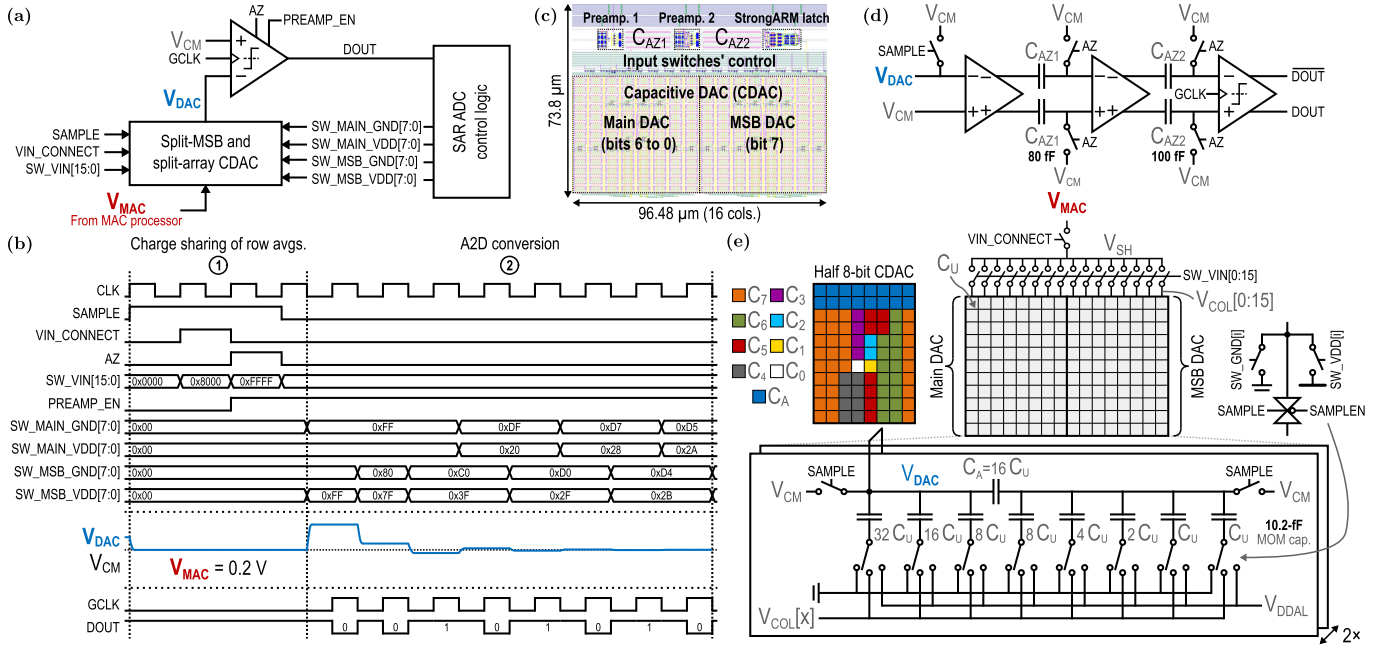


Fig. 14. (a) Schematic, (b) timing diagram, and (c) layout of the 8b SAR ADC spanning over 16 columns of the analog memory. Detailed schematic of (d) the dynamic comparator, consisting of two preamplification stages and a latch, and (e) the CDAC, employing the split-MSB and -array techniques.

MOM caps  $C_U$  connected in parallel in each MAC unit. This circuit thus implements integer weights ranging from  $-7$  to  $7$ , multiplied by a factor  $0.25\times$ , originating from the fact that each column includes a part of the feedback cap  $C_{FB}$  equal to  $4C_U$ . The MAC unit fits inside the pixel pitch and occupies a height of  $28.85\ \mu\text{m}$  [Fig. 11(e)].

The performance of these MAC units is characterized with post-layout simulations in Fig. 12. First, Fig. 12(a) and (b) corresponds to a setup in which eight MAC units have their weight set to  $+7$  with a shared input voltage  $V_{\text{BUF},+}$ , while the other eight units have their weight set to  $-7$  with input voltage  $V_{\text{BUF},-}$ . The impact of local mismatch is studied in this context with  $10^3$  Monte Carlo (MC) simulations. Across the input range  $\Delta V_{\text{BUF}} = (V_{\text{BUF},+} - V_{\text{BUF},-})$ , the  $\sigma$  of  $V_{\text{MAC}}$  remains below  $1\ \text{mV}$ , while that of  $V_{\text{AMP},\text{IN}-}$  is affected by the statistical offset of the OTA, leading to a  $2.5\text{-mV}$   $\sigma$ . More specifically, for  $\Delta V_{\text{BUF}} = 0.3\ \text{V}$ , the proposed offset-insensitive structure reduces the  $\sigma$  by  $14.2\times$  from  $2.55$  to  $0.18\ \text{mV}$ . However, this first setup only accurately describes a specific realization of input voltages and weights. Therefore, Fig. 12(c) and (d) extends this analysis to a baseline of  $10^4$  random combinations of inputs and weights drawn from uniform distributions, with local mismatch and intrinsic noise subsequently applied on top of it. In the  $V_{\text{MAC}}$  output range between  $0.15$  and  $1.05\ \text{V}$ , avoiding transistors to be biased outside of saturation, the average error in Fig. 12(c) shows a deterministic behavior hinting at a slope error, which can be related to parasitic capacitances or charge injection, while the average standard deviation of the error over the considered output range increases from  $0.42\ \text{mV}$  to, respectively,  $0.80$  and  $0.74\ \text{mV}$  with local mismatch and intrinsic noise. These variabilities are predominantly due to the mismatch of MOM caps and the thermal  $kT/C$  noise due to the sampling of input voltages on the caps in the MAC units. These results demonstrate the robustness of the

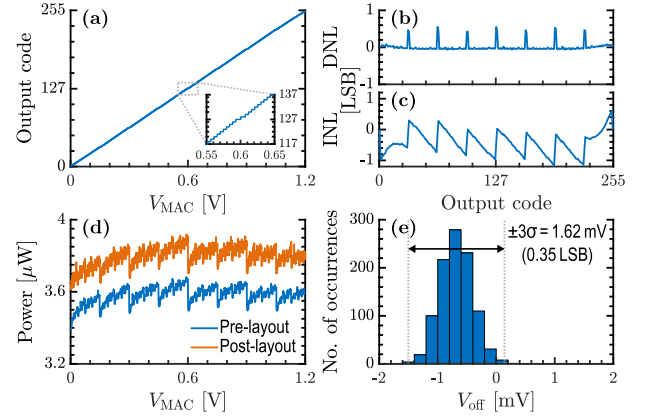


Fig. 15. (a) Transfer function of the SAR ADC with the corresponding (b) DNL and (c) INL. (d) Power consumption on  $V_{\text{DDAL}} = 1.2\ \text{V}$  (CDAC, comparator, and drivers) as a function of the input voltage. (e) Statistical offset of the StrongARM latch referred to the input of the first preamplifier. All figures correspond to the TT  $25^\circ\text{C}$  corner.

proposed architecture to these analog nonidealities. However, Fig. 13 highlights one of the architectural details which could be further improved. As emphasized in Fig. 13(a), when one of the magnitude bits is equal to zero, a leakage current flows through the TG, implemented with high-speed low- $V_{\text{th}}$  (LVT) core devices with a minimum length of  $120\ \text{nm}$ . Hence, this introduces a stochastic error on  $V_{\text{MAC}}$  when both local mismatch and global process variations are considered, with an average  $\sigma$  around  $7.46\ \text{mV}$  [Fig. 13(b)] compared to  $0.80\ \text{mV}$  with mismatch only. To mitigate this problem, the transistors constituting the TG can either be made longer or rely on low-leakage high- $V_{\text{th}}$  (HVT) core devices, respectively, reducing the average  $\sigma$  to  $2.12$  and  $0.40\ \text{mV}$ . Interestingly, the  $0.40\text{-mV}$   $\sigma$  obtained with an HVT TG is even lower than the one obtained with local mismatch only with the LVT TG,

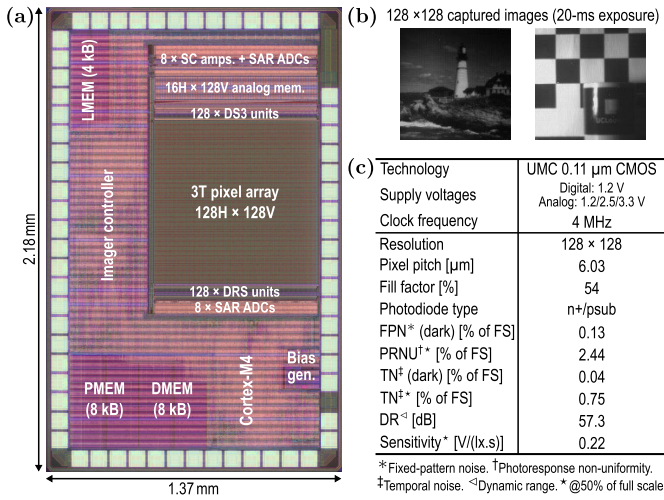


Fig. 16. (a) Chip microphotograph with overlaid layout. (b) Example of captured images with an exposure time of 20 ms and (c) imaging characteristics of the proposed imager.




Fig. 17. (a) Output code, (b) PRNU and TN, and (c) SNR as a function of illuminance, computed for ten images for each illuminance. (d) Power breakdown in imaging mode. All figures correspond to an exposure of 20 ms.

suggesting that the sensitivity to local mismatch not only stems from MOM caps, but also from the TG leakage.

3) *SAR ADCs for Aggregation and Digitization*: The SAR ADCs employed in this work follow the topology presented in Fig. 14(a) with the associated timing diagram shown in Fig. 14(b). Similar to the SC amplifiers, each SAR ADC spans 16 pixel columns [Fig. 14(c)]. The two main functions of the SAR ADC employed in this work are as follows: (1) the aggregation of psums by charge sharing, following the calculation of the psums of rows by the SC amplifiers and (2) the A2D conversion of the convolution result, following the SAR principle. In terms of circuit implementation, the detailed comparator architecture in Fig. 14(d) features two differential preamplification stages based on a differential pair driving a load of diode-connected transistors and providing a total gain of approximately 10 V/V. These stages both embed AZ capabilities based on caps at their output and are followed by a dynamic StrongARM latch. In addition, the CDAC combines two existing techniques to reduce its power and area over-

heads. First, it relies on a split-MSB array [16], based on two identical DACs for the MSB and all the LSBs, respectively, called MSB DAC and main DAC in Fig. 14(c)–(e). More specifically, as illustrated in Fig. 14(b), the MSB DAC is initially switched to  $V_{DD}$ . Then, when the output data  $D_{OUT}$  of the comparator is equal to zero, the current bit of the MSB DAC is switched to ground, whereas when  $D_{OUT}$  is equal to one, the current bit of the main DAC is switched to  $V_{DD}$ . This allows to harmonize power consumption across the input voltage range by optimizing the switching of the DAC. Then, it also makes use of a split-capacitive array [17], employing an attenuation cap  $C_A$  in Fig. 14(e) to reduce the impact of the LSBs, ultimately allowing to diminish the capacitance of the MSBs. Besides, possible further improvements in terms of EE and silicon area can be obtained by employing parasitic caps instead of explicit MOM ones, as outlined by Harpe [18]. Interestingly, in the RoI detection mode generating 1b fmaps, the offset associated with each filter is also implemented with the CDAC, by switching up (resp. down) bits of the main (resp. MSB) DAC to implement a positive (resp. negative) offset.

Regarding the post-layout simulation results in typical conditions (TT 25  $^{\circ}\text{C}$ ) in Fig. 15, a relatively good linearity is achieved in Fig. 15(a), with the differential nonlinearity (DNL) comprised between  $-0.07$  and  $0.55$  LSB, and the integral one (INL) between  $-1.17$  and  $0.62$  LSB [Fig. 15(b) and (c)]. Regarding power consumption, it is relatively constant across the input voltage range because of the split-MSB CDAC, with a mean value of  $3.59$  and  $3.78$   $\mu\text{W}$  in pre- and post-layout simulations [Fig. 15(d)]. At last, the input-referred statistical offset of the comparator features a  $\pm 3\sigma$  value of  $1.62$  mV corresponding to  $0.35$  LSB, hence ensuring that the comparator operation is robust to local mismatch [Fig. 15(e)].

#### IV. EXPERIMENTAL RESULTS

MANTIS imager SoC has been fabricated in UMC  $0.11\text{-}\mu\text{m}$  bulk CMOS technology. The chip microphotograph is shown in Fig. 16(a), together with the examples of captured images in Fig. 16(b). This section presents the experimental results obtained with the fabricated chip, respectively, focusing on the characterization of the image sensor in Section IV-A and of the mixed-signal near-sensor convolution processor in Section IV-B. Finally, the applicative performance is evaluated on a face RoI detection task in Section IV-C.

##### A. Imaging Performance

We first characterize the performance of MANTIS in imaging mode with the results summarized in Figs. 16(c) and 17. This characterization is performed by exposing the imager to an uniform light flux, generated by an Olympus KL 1500 halogen light source going through an integrating sphere and by capturing ten images for each light flux level. The transfer function between the imager 8b output code and the light flux per unit area, i.e., the illuminance expressed in  $\text{lm}/\text{m}^2$  or lx, is depicted in Fig. 17(a). It highlights that, for a 20-ms exposure time, the usable illuminance ranges from 120 to 1500 lx. In addition, the transfer function levels off at low illuminance



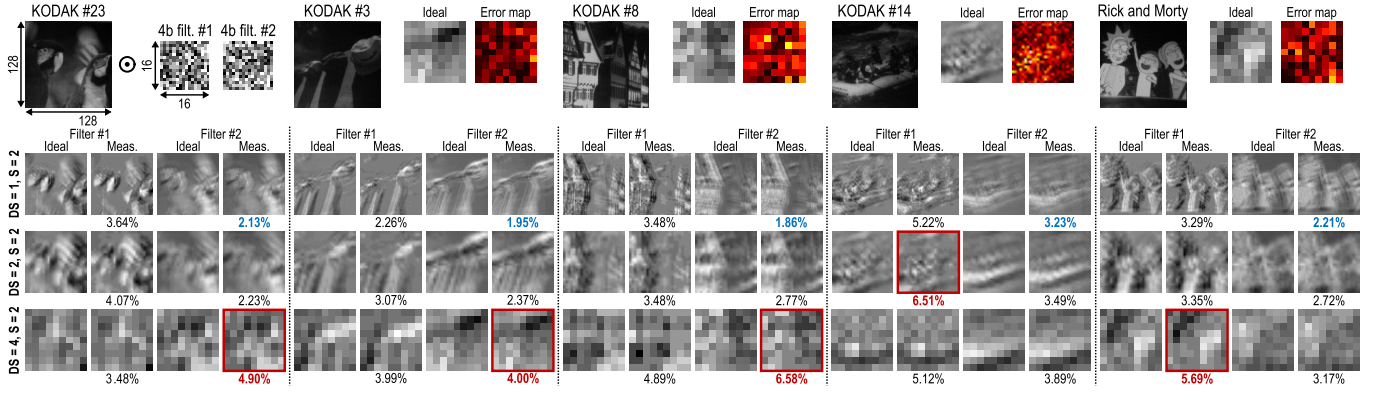


Fig. 18. Comparison between ideal and measured fmaps, respectively computed in software (MATLAB) and on chip, obtained by a convolution operation between the image and two random 4b 16×16 filters. The parameters used for this operation are  $S = 2$ ,  $DS = 1, 2$ , or  $4$ , and an exposure time of 12.5 ms. The RMSE for each fmap is written below it, and the error map corresponds to the fmap with the worst RMSE among the displayed ones.

TABLE I

SUMMARY OF THE MEASURED PERFORMANCE OF MANTIS IMAGER SoC IN CONVOLUTION MODE FOR FOUR FILTERS (12.5-MS EXPOSURE TIME)

Image downsampling (DS)	1				2				4			
Filter stride (S)	2	4	8	16	2	4	8	16	2	4	8	16
Frame rate* [fps]	18.2	79.7	79.7	79.7	79.7	79.7	79.7	79.7	79.7	79.7	79.7	79.7
Throughput† [MOPS]	121	137.3	36.7	<b>10.5</b>	<b>408.3</b>	110.4	32.0	<b>10.5</b>	211.7	65.3	23.5	<b>10.5</b>
Feature map RMSE° [%]	<b>3.01</b>	3.25	4.00	4.69	3.40	3.98	6.30	8.68	4.88	<b>11.34</b>	9.19	8.45
Power <sup>‡</sup> (accelerator) [ $\mu$ W]	66.84	<b>76.20</b>	22.36	8.40	58.74	17.40	6.60	4.03	10.07	4.42	3.29	<b>2.70</b>
EE <sup>‡</sup> (accelerator) [TOPS/W]	7.24	7.31	6.57	<b>4.98</b>	27.80	25.38	19.40	10.37	<b>84.09</b>	59.17	28.61	15.48
Energy/OP <sup>‡</sup> (accelerator) [fJ/op]	138.1	138.7	152.1	<b>200.9</b>	36.0	39.4	51.6	96.4	<b>11.9</b>	16.9	35.0	64.6
Power <sup>‡</sup> (SoC) [ $\mu$ W]	338.5	<b>384.7</b>	297.4	268.9	357.0	288.0	264.7	256.3	271.9	258.3	253.3	<b>250.9</b>
EE <sup>‡</sup> (SoC) [TOPS/W]	1.43	1.43	0.49	<b>0.16</b>	<b>4.57</b>	1.53	0.48	<b>0.16</b>	3.11	1.01	0.37	0.17
Energy/OP <sup>‡</sup> (SoC) [pJ/op]	0.70	0.70	2.02	<b>6.43</b>	<b>0.22</b>	0.65	2.07	<b>6.13</b>	0.32	0.99	2.69	6.00
Processing energy (SoC) [pJ/(pix-frame-filt)]	<b>284.1</b>	73.6	56.9	51.5	68.3	55.1	50.7	49.0	52.0	49.4	48.5	<b>48.0</b>

\* Frame rate is limited by the 12.5-ms exposure time. † Expressed in operations with analog inputs and 4b weights. ° Computed over 10 images with 10 random filters. ‡ Includes the analog memory, SC amplifiers, SAR ADCs and drivers on  $V_{DDAL}$ . ‡ Normalized to 1b operations. ‡ Includes the imager analog macro ( $V_{DDAL}$  and  $V_{DDAH}$ ), and the digital core, i.e., the Cortex-M4 CPU, the imager controller, and the SRAM macros ( $V_{DDD}$ ).

due to the relatively low photocurrent values compared with leakages inside the pixel, the photoresponse of the n+/psub diodes available in this CMOS logic process being far from optimal compared with a CMOS image sensor (CIS) process. Next, Fig. 17(b) evaluates two types of noise affecting the image quality, namely, photoresponse non-uniformity (PRNU), capturing the variability of pixel responses to light due to mismatch, and temporal noise (TN). PRNU and TN are worth 2.44 and 0.75% of the full scale (FS) at 50% of the FS and correspond to a signal-to-noise ratio (SNR) slightly above 20 dB in the usable illuminance range, dominated by the PRNU [Fig. 17(c)]. Taking a more theoretical perspective, the voltage noise due to thermal noise at the output of the DRS units can be expressed as

$$\bar{v}_n = \sqrt{2kT} \sqrt{\frac{A_{SF}^2}{C_{PD}} + \frac{1}{C_S}} \quad (3)$$

with  $C_{PD} = 12.2$  fF the pixel capacitance,  $C_S = 29$  fF the sampling capacitance, and  $A_{SF} = 0.69$  V/V the gain of the pixel's SF. This yields  $\bar{v}_n = 0.78$  mV at 25 °C corresponding to a 0.65% error with a 1.2-V dynamic range of  $V_{PIX}$ , which is in the same order of magnitude as the TN in Fig. 17(b). A 4T pixel array could easily be integrated to the proposed design by switching to pinned photodiodes and modifying the digital

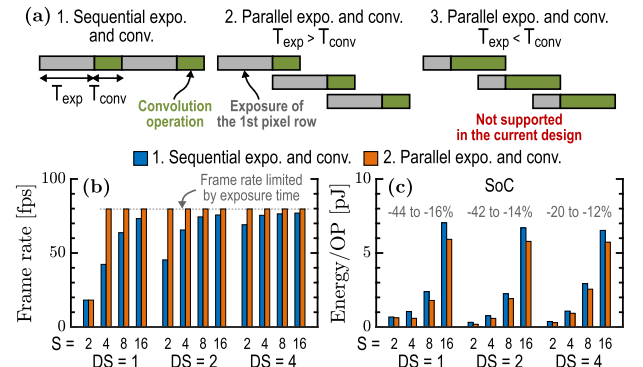


Fig. 19. (a) Illustration of sequential and parallel exposure and convolution. Measured (b) frame rate and (c) energy per 1b operation (SoC) for the sequential and parallel executions and for different DS and S configurations, four filters, and a 12.5-ms exposure time.

controller, and would reduce the contribution of TN because of correlated double sampling (CDS). The imaging characteristics are summarized in Fig. 16(c). Finally, Fig. 17(d) describes the power consumption of the imager for a frame rate of 29 fps. Power is dominated by the digital part, which represents 78% of the 335.6- $\mu$ W SoC power, with the following split: 38% for the imager controller, 25% for the CPU, and 13% for data transfers by the DMA. It could easily be reduced



by moderately scaling  $V_{DD}$  to 1 V or by making use of power-gating techniques. The power of the analog circuitry only amounts to 22% of the SoC power, with most of it (17%) being consumed by the pixel array and DRS units supplied at 2.5 V, while the remaining 5% correspond to SAR ADCs supplied at 1.2 V. The SoC power corresponds to an energy per pixel of 706.3 pJ/(pix-frame), which is larger than state-of-the-art values for low-power imagers, typically ranging from 100 to 300 pJ/(pix-frame) [13]. Nonetheless, this is perfectly normal, as our SoC is not optimized for imaging, and as imagers usually do not include a CPU and a DMA.

### B. Electrical Characterization of Near-Sensor Convolutions

To validate the proper operation of the mixed-signal convolution processor, two aspects need to be thoroughly quantified as follows: 1) the quality of the fmaps computed by the chip with respect to an ideal execution in software [Fig. 18] and 2) the throughput and EE of the MAC operations, at the accelerator and SoC levels [Figs. 19–21]. This analysis must cover the different configurations of the proposed convolution processor in terms of image DS factor and filter stride  $S$ . For the characterization of throughput and EE, we rely on the benchmarking outlined by Shanbhag and Roy [19] in the context of IMC, which shares striking similarities with convolutional imagers, except that IMC is weight-stationary, while convolutional imagers are input-stationary. The concept of *ADC column* originates from this article and corresponds in this work to a group of 16 columns of the analog memory feeding 16 MAC units, connected to an SC amplifier and an 8b SAR ADC.

We start by comparing ideal fmaps with measured ones based on two important steps. First, fmaps need to be normalized, as the pixel values of the 8b  $128 \times 128$  image used in the ideal software execution in MATLAB do not reflect the analog values employed on chip by the convolution processor. This difference stems from the various multiplicative factors applied to raw pixel voltages in the blocks constituting the convolution pipeline. For a given fmap denoted as  $f$ , the normalized fmap denoted as  $\hat{f}$  is computed as

$$\hat{f} = \frac{f - \mu(f)}{\sigma(f)} \quad (4)$$

thereby ensuring that the mean  $\mu$  and the standard deviation  $\sigma$  of the resulting fmap are, respectively, equal to zero and one. The second important step is the metric used to assess the quality of the computed fmap. Here, we rely on the root-mean-square error (RMSE) calculated as

$$\text{RMSE} = \frac{100\%}{2\max(|\hat{f}_{\text{meas}}|)} \sqrt{\frac{1}{N_f^2} \sum_{i=1}^{N_f} \sum_{j=1}^{N_f} (\hat{f}_{\text{ideal},ij} - \hat{f}_{\text{meas},ij})^2} \quad (5)$$

where  $\hat{f}_{\text{ideal}}$  and  $\hat{f}_{\text{meas}}$  are, respectively, the normalized ideal and measured fmaps, and  $N_f$  is the fmap size, obtained from

$$N_f = \left( \frac{128}{DS} - F \right) \frac{1}{S} + 1 \quad (6)$$

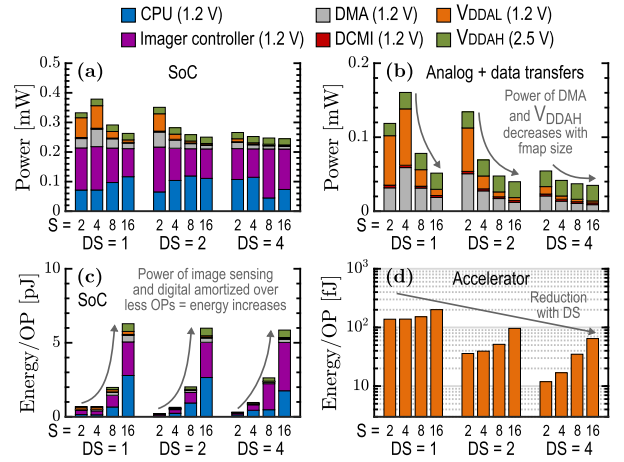


Fig. 20. Breakdown of (a) the total SoC power and (b) the power of the analog macro and digital data transfers. Energy per 1b operation for (c) the SoC and (d) the accelerator. All figures correspond to a **parallel** exposure and convolution, different DS and  $S$  configurations, four filters, and a 12.5-ms exposure time and present measurement results. The fraction of power due to the imager controller in (a) and (c) is estimated based on physical simulations.




Fig. 21. For **sequential** exposure and convolution,  $DS = 1$ ,  $S = 2$ , and a 12.5-ms exposure time, measured (a) frame rate, EE at the SoC and accelerator levels, and (b) energy per 1b operation with the number of filters.




Fig. 22. Training pipeline of the face ROI detector, based on a quantized CNN trained with QKeras and a TensorFlow backend (ideal software execution). The frame rate is 27 fps and is dominated by the duration of the convolution operation rather than by the exposure time.

with  $F = 16$  the filter size. The characterization is performed over ten images, among which nine are part of the KODAK dataset of natural images, and with ten 4b-weighted filters drawn from a uniform distribution. Table I details the RMSE results. It indicates that the RMSE is comprised between 3.01 and 11.34%, and that it tends to degrade for smaller fmaps with a larger DS factor and/or a larger filter stride. This is quite intuitive to understand, given that (5) relies on  $\max(|\hat{f}_{\text{meas}}|)$  to approximate the range of values contained in an fmap and is, hence, sensitive to errors in large values of  $\hat{f}_{\text{meas}}$ . However, we believe that the proposed metric provides both an intuition

TABLE II  
COMPARISON TABLE OF STATE-OF-THE-ART MIXED-SIGNAL VISION CHIPS

	In-sensor			Hybrid		Near-sensor				
	Jendernalik [1]	Carey [2]	Xu [3]	Lefebvre [9]	Song [10]	Kim / Bong [4] / [5]	Young [6]	Hsu [7]	Hsu [8]	Lefebvre This work
Publication Year	TCAS-I 2013	VLSI 2013	TCAS-I 2022	ISSCC 2021	VLSI 2021	ESSCIRC / JSSC 2017 / 2018	JSSC 2019	JSSC 2021	JSSC 2023	JSSC 2024
Technology	0.35 $\mu$ m CMOS	0.18 $\mu$ m CMOS	0.18 $\mu$ m CMOS	65nm CMOS	0.18 $\mu$ m CIS	65nm CMOS	0.13 $\mu$ m CIS	0.18 $\mu$ m CMOS	0.18 $\mu$ m CMOS	0.11 $\mu$ m CMOS
Area [mm <sup>2</sup> ]	9.8	10 $\times$ 10	N/A	2 $\times$ 2	4.1 $\times$ 5.2	3.3 $\times$ 3.6	4 $\times$ 4	2.46 $\times$ 2	2.46 $\times$ 2.18	1.37 $\times$ 2.18
Supply voltage [V]	3.3	1.8 (Digital) 1.5 (Analog)	0.8–1.8	0.8/1 (Digital) 0.95/1.05 (Analog)	1.8 (Digital) 2.5 (Analog)	0.5–0.8 (Digital) 2.5 (Analog)	0.9 (Digital) 1.5/2.5 (Analog)	0.5	0.8	1.2 (Digital) 1.2/2.5/3.3 (Analog)
Resolution	64 $\times$ 64	256 $\times$ 256	32 $\times$ 32	160 $\times$ 128	240 $\times$ 240	320 $\times$ 240	320 $\times$ 240	128 $\times$ 128	126 $\times$ 126	128 $\times$ 128
Shutter	Global	Global	Global	Rolling	Global	Rolling	Global	Rolling	Rolling	Rolling
Double sampling	No	No	DRS	No	CDS	No	CDS	No	No	DRS
Frame rate [fps]	100	100,000	156	24–268	120	1	30	480	50–250	18.2–79.7
Pixel pitch [ $\mu$ m]	35	32.3	35	9	9.8	7	4	7.6	7.6	6.03
Pixel complexity	18T APS + 2 MOS caps.	176T APS	61T APS	40T log(I) + 1 MIM cap.	14T APS + 4 caps.	3T APS	4T APS	4T PWM	4T PWM	3T APS
Fill factor [%]	23	6.2	9.1	12.9	20.1	N/A	60.4	36	36	54
DR [dB]	58	N/A	N/A	47.1	N/A	N/A	59.3	52.3	47.8	57.7
Feature type	- 3 $\times$ 3 kernels	- Edge detection - Median filtering	- 32 $\times$ 32 kernels	- 2 $\times$ 2 to 64 $\times$ 64 kernels - 16 $\times$ 16 lin. Haar filters - 6 scales (conv.) - 3 scales (Haar)	- Log. Haar filters	- 20 $\times$ 20 lin. Haar filters	- Log. gradients	- 3 $\times$ 3 kernels	- 3 $\times$ 3 kernels	- 16 $\times$ 16 kernels
Multiscale	No	Arbitrary	No	Arbitrary	Arbitrary	3 scales	Arbitrary	No	No	3 scales
Computation type	Current	Current	Current	Current	Charge	Charge	Charge	Current	Current	Charge
Weight resolution	Analog	N/A	1b	1.5b	1.5b	1.5b	N/A	4b	4b	4b
Feature resolution	Analog	1b or 8b	1b	1b or 8b	1b	1b	1.5b or 2.75b	1b to 8b	1b to 8b	1b, 2b, 4b, or 8b
Throughput <sup>‡</sup> [MOPS]	7.4	655,000	5.1	15.1–252.1	N/A	N/A	N/A	137.2	63.5	10.5–408.3
Throughput <sup>‡‡</sup> [MOPS]	N/A	N/A	5.1	22.7–378.2	N/A	N/A	N/A	548.7	254	42–1633.2
Power (accel.) [ $\mu$ W]	N/A	N/A	0.147–0.537	N/A	N/A	N/A	N/A	N/A	N/A	2.7–76.2
EE <sup>‡</sup> (accel.) [TOPS/W]	N/A	N/A	9.52–34.77	N/A	N/A	N/A	N/A	N/A	N/A	4.98–84.09
Power (SoC) [ $\mu$ W]	280	1,230,000	8.5	42–206	2,900	24–96	229–262	117	80.4–134.5	250.9–384.7
EE <sup>‡</sup> (SoC) [TOPS/W]	0.026	0.53	0.60	0.23–5.46	N/A	N/A	N/A	4.67	0.63–1.89	0.16–4.57
Processing energy (SoC) [pJ/(pix-frame-filt)]	683.6	187.7	3.3	2.5–103.9*	16.8 <sup>†</sup>	6.0–24.0 <sup>◊</sup>	49.7–56.9 <sup>◊</sup>	14.9	4.2–12.7	48.0–284.1*

\* For 4 filters. † For 25 filters. ◊ For 52 filters. ◊ Horizontal and vertical gradients are considered as two filters. ‡ Not normalized to the resolution of inputs and weights.

‡ Normalized to 1b operations.

and a quantification of the magnitude of the error, despite these inaccuracies. A few fmaps are displayed in Fig. 18 with the corresponding RMSE, as well as an error map for the fmap with the worst RMSE among the displayed ones. It reveals that the measured fmaps strongly resemble the ideal ones and properly capture the image features. Errors are barely noticeable with the naked eye and consist of slightly different values between fmaps.

Then, we turn to the assessment of the throughput and EE of the MAC operations. We first introduce the throughput as

$$\text{Throughput} = \text{fps} \cdot N_{\text{filt}} \cdot N_f^2 \cdot (2 \cdot F^2 \cdot \text{DS}^2) \quad (7)$$

where  $N_{\text{filt}}$  corresponds to the number of filters. This definition of the throughput does not account for the resolution of the inputs and weights involved in the MAC operations. Next, we can define the energy per 1b operation as

$$\text{Energy/OP} = \frac{\text{Power}}{\text{fps} \cdot N_{\text{filt}} \cdot N_f^2 \cdot (2 \cdot F^2 \cdot \text{DS}^2) \cdot (B_X \cdot B_W)} \quad (8)$$

where  $B_X$  and  $B_W$ , respectively, stand for the resolution of the inputs and weights. In the proposed SoC, MAC operations are based on analog inputs and 4b weights. Hence, we use  $B_X = 1$  and  $B_W = 4$ , even though using  $B_X$  equal to the effective number of bits (ENOB) at the input of the MAC units could be possible to compare the results with accelerators, such as IMC ones, for which the resolution of inputs is clearly defined. Throughput can also be normalized to 1b operations by multiplying its expression in (7) by  $B_X \cdot B_W$ .

Fig. 19(a) illustrates different cases regarding how the exposure and convolution operations intertwine. A sequential

execution is inefficient, as pixels can start being exposed as soon as they have been stored in the analog memory. Therefore, a parallel execution is preferable. The current version of the imager controller only supports the case in which the exposure time  $T_{\text{exp}}$  is longer than the duration of the convolution operation  $T_{\text{conv}}$  (case 2), but could easily be modified to support a parallel execution for  $T_{\text{exp}} < T_{\text{conv}}$  (case 3). This modification would be beneficial from an applicative standpoint, as it would allow to maximize the frame rate in all the configurations of the accelerator. Fig. 19(b) and (c) corresponds to an execution with four filters and a 12.5-ms exposure time for all possible configurations of DS and S. Fig. 19(b) reveals that a higher frame rate can be achieved with parallel execution, the limit being the exposure time. In terms of energy/OP at the SoC level, the parallel execution yields a reduction of 12 to 44% with the strongest reductions attained for small DS and S.

Besides, Fig. 20 presents the breakdowns of the power consumption and energy/OP at the levels of the accelerator and of the SoC. Regarding the SoC power [Fig. 20(a)], it ranges from 245 to 379  $\mu$ W. The CPU and imager controller have a relatively constant consumption around 0.2 mW across configurations, while the consumption related to the analog circuitry and data transfers is highly dependent on the configuration [Fig. 20(b)]. The power on  $V_{\text{DDAL}}$  and of the DMA declines for a larger DS and/or S, as  $N_f$  becomes smaller, while the power on  $V_{\text{DDAH}}$  and of the DCMI remains fairly constant. The former is indeed related to the pixel readout and DS3 units and does not change as the frame rate is the same for all configurations except for  $\text{DS} = 1$  and  $\text{S} = 2$ . This is the case, because a parallel execution is employed, and the frame rate is limited by the exposure time as in Fig. 19(a). The latter is

related to internal switching of the DCMI and does not account for the I/O power, which would, otherwise, scale with the amount of data, similar to the DMA. Therefore, the energy/OP at the SoC level [Fig. 20(c)] goes from 0.22 to 6.43 pJ and degrades for large strides, as the power is amortized over a smaller number of operations. Finally, the energy/OP at the accelerator level [Fig. 20(d)] is comprised between 12 and 201 fJ, corresponding to an EE of 84.09 and 4.98 TOPS/W. Fig. 20(d) and Table I further reveal that the energy/OP at the accelerator level improves with a larger DS, as the filter is applied to a larger number of pixels in the original image because of the DS operation. Interestingly, the two key features to achieve a high EE at the accelerator level are the power gating of OTAs and the amortization of the MAC operation energy over a large number of pixels because of image DS.

Finally, Fig. 21 studies the impact of the number of filters  $N_{\text{filt}}$  on the frame rate, EE, and energy/OP, for sequential exposure and convolution. Fig. 21(a) highlights that increasing  $N_{\text{filt}}$  causes the frame rate to drop, as  $T_{\text{conv}}$  becomes longer while  $T_{\text{exp}}$  remains constant, and that the accelerator EE remains relatively constant, while the SoC one slightly improves, as the fraction of time without convolution operations decreases, and the digital power is amortized over a larger number of operations. The same trend is reflected by the energy/OP at the SoC level in Fig. 21(b).

### C. Face Region-of-Interest Detection

This last experiment consists in demonstrating the operation of MANTIS in a face RoI detection use case. The structure and training pipeline of the RoI detector, implemented as a quantized CNN, are illustrated in Fig. 22. The first part of the RoI detector is a convolution layer executed on chip, using 16 4b  $16 \times 16$  filters and 8b offsets, and operating over the image downsampled by  $2 \times$ . It is followed by an off-chip fully connected (FC) layer with 8b weights, combining 1b fmaps to generate a 1b RoI detection map. Most of the workload is executed on chip, with 20.48 million operations in the convolution layer against 21.25 thousands in the FC one. Note that an ad hoc digital accelerator could realize the FC layer on chip. An interesting feature of this detector is that it reduces the data that need to be transmitted off chip to 7.63% of the raw 8b image, thus cutting down the I/O bandwidth by  $13.1 \times$ . As the EE at the SoC level for a sequential execution is 4.57 TOPS/W [Table I] and the difference between the two execution types does not exceed 44% [Fig. 19(c)], we expect the EE for a parallel execution to be above 2.56 TOPS/W. The network is trained with QKeras on a dataset consisting of background and face images used by Moons et al. [20] and achieves false and true negative rates (FNR and TNR) of 8.5 and 96.9% on the test set, respectively, for an ideal software execution. At last, Fig. 23(a) shows the detailed results of one of the test images and provides the overall performance over the ten test images. MANTIS achieves an 11.5% FNR while, respectively, discarding 81.9% and 81.3% of image patches for the ideal and measured executions. These results are in




Fig. 23. (a) Measured face RoI detection results, with details of a single test image and overall results for the ten test images. (b) Face RoI results over four additional test images.

line with the software execution but with a slight degradation coming from the fact that images generated by the imager are different from the ones in the dataset. Fig. 23(b) displays the face RoI results over four additional images, with a measured percentage of discarded image patches between 76.5 and 84.3%, and a single discarded face. Interestingly, the overall performance remains largely similar, whereas RoI detection maps are different for the ideal and measured executions, due to the adaptation of the biases of the convolution layer in measurement and an approximate modeling of raw pixel voltages' transformations inside the convolution pipeline.

## V. COMPARISON WITH THE STATE OF THE ART

In this section, we compare our work with the state of the art of mixed-signal vision chips in Table II and with other relevant accelerators. The proposed SoC relies on a 3T APS with DRS to compensate FPN. This a key enabler to achieve a  $6.03\text{-}\mu\text{m}$  pixel pitch and a 54% fill factor, which are superior to the existing vision chips, especially in-sensor ones for which the pixel pitch usually exceeds  $30\text{ }\mu\text{m}$ . In terms of functionality, MANTIS is the first work to combine large  $16 \times 16$  filters, a 4b weight resolution, and operation at multiple scales, making it suitable for medium-complexity vision tasks, such as RoI detection. The charge-domain MAC operations computed by the accelerator have an EE normalized to 1b operations (1b EE) between 4.98 and 84.09 TOPS/W, so  $2.4 \times$  better than [3]. This information is however missing from other works, which only report EE at the SoC level. The 1b EE at the SoC level ranges from 0.16 to 4.57 TOPS/W and is on



par or better than the existing works [7], [9], while, respectively, supporting a larger filter size and an increased weight resolution. Regarding the processing energy, it is larger than for other works, as this metric does not account for the filter characteristics. Next, let us compare near-sensor and hybrid vision chips. The latter usually present larger pixels, which contain registers to store weights [9], or large caps to store and read pixel values [10]. Regarding double sampling, it is possible when the readout is voltage-based [10] but not when photocurrents are used [9]. Finally, near-sensor vision chips are generally more flexible, as different configurations of the mixed-signal processor can easily be implemented. In contrast, hybrid architectures often involve hardwiring of connections between pixels and either require to move weights in the pixel array [9] or to combine row and column signals [10]. These architectural features thus limit the use of hybrid vision chips to specific applications.

Let us now consider other types of convolution accelerators, starting with those for which the digital control is implemented with an FPGA (not included in the power consumption). Two relevant works employ charge-domain computations, either using IMC with analog inputs and 1b weights [21] or in-column SC amplifiers with 3b weights, followed by ADCs performing a nonlinear quantization [22]. Their 1b EE at the accelerator level is, respectively, 1.25 and 0.017 TOPS/W, which is approximately  $4\times$  lower than the worst case EE obtained with MANTIS. Other works have their digital control embarked on chip. First, Abedin et al. [23] propose a current-domain in-sensor processor with 2b weights based on resistive RAM and attaining a 1b EE of 2.98 TOPS/W. Then, Wang et al. [24] feature a hybrid optical-electronic CNN processor with 1b weights, realizing the first convolution layer with a mask in the optical domain and reaching a 1b EE of 0.37 TOPS/W. Finally, Eki et al. [25] introduce a digital CNN processor stacked on an image sensor exhibiting a peak EE of 4.97 TOPS/W for 8b and 32b integer operations. MANTIS is on par with [23] and [24] in terms of EE, but offers more flexibility with its programmable convolution parameters. However, it is not as efficient as [25], leveraging the performance improvements due to scaling.

## VI. CONCLUSION

In this work, we presented MANTIS, a mixed-signal near-sensor convolutional imager SoC intended for FE and RoI detection. It is the first mixed-signal vision chip to combine large  $16\times 16$  filters with 4b weight resolution, operation at three different scales, and DRS to remove FPN and improve computational accuracy. MANTIS is enabled by two main circuit innovations. First, DS3 units combine DRS, voltage downshifting, and image DS. Next, near-sensor MAC operations are computed in the charge domain, using SC amplifiers to compute psums and charge sharing in the CDAC of the SAR ADCs to aggregate psums and compute the convolution result. MANTIS, respectively, reaches peak EEs normalized to 1b operations of 4.57 and 84.09 TOPS/W at the accelerator and SoC levels, while producing fmaps with an RMSE between 3.01 and 11.34%. Finally, face RoI detection was

demonstrated with an FNR of 11.5%, while discarding 81.3% of image patches and reducing the imager output data to 7.63% of the raw image. This work demonstrates that near-sensor vision chips can successfully tackle tasks requiring a higher resolution of inputs and weights, as opposed to in-sensor vision chips, which are currently limited to noisy inputs and low-resolution weights. Further works should focus on the digital part of the SoC, albeit some analog blocks could also benefit from the utilization of more advanced techniques. New opportunities for the implementation of mixed-signal vision chips also arise from 2.5-D/3-D packaging.

## ACKNOWLEDGMENT

The authors would like to thank Prof. Marian Verhelst and Dr. Bert Moons for granting us access to their face detection dataset, Dr. Rémi Dekimpe for his help with the DMA, Prof. Charlotte Frenkel and Dr. Adrian Kneip for fruitful discussions, and Eléonore Masarweh for the microphotograph.

## REFERENCES

- [1] W. Jendernalik, G. Blakiewicz, J. Jakusz, S. Szczepanski, and R. Piotrowski, "An analog sub-milliwatt CMOS image sensor with pixel-level convolution processing," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 60, no. 2, pp. 279–289, Feb. 2013.
- [2] S. J. Carey, A. Lopich, D. R. Barr, B. Wang, and P. Dudek, "A 100,000 fps vision sensor with embedded 535 GOPS/W  $256\times 256$  SIMD processor array," in *Proc. Symp. VLSI Circuits*, 2013, pp. C182–C183.
- [3] H. Xu et al., "Senputing: An ultra-low-power always-on vision perception chip featuring the deep fusion of sensing and computing," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 69, no. 1, pp. 232–243, Jan. 2022.
- [4] C. Kim, K. Bong, I. Hong, K. Lee, S. Choi, and H. Yoo, "An ultra-low-power and mixed-mode event-driven face detection SoC for always-on mobile applications," in *Proc. 43rd IEEE Eur. Solid State Circuits Conf. (ESSCIRC)*, Sep. 2017, pp. 255–258.
- [5] K. Bong, S. Choi, C. Kim, D. Han, and H. Yoo, "A low-power convolutional neural network face recognition processor and a CIS integrated with always-on face detector," *IEEE J. Solid-State Circuits*, vol. 53, no. 1, pp. 115–123, Jan. 2018.
- [6] C. Young, A. Omid-Zohoor, P. Lajevardi, and B. Murmann, "A data-compressive 1.5/2.75-bit log-gradient QVGA image sensor with multi-scale readout for always-on object detection," *IEEE J. Solid-State Circuits*, vol. 54, no. 11, pp. 2932–2946, Nov. 2019.
- [7] T.-H. Hsu et al., "A 0.5-V real-time computational CMOS image sensor with programmable kernel for feature extraction," *IEEE J. Solid-State Circuits*, vol. 56, no. 5, pp. 1588–1596, May 2021.
- [8] T.-H. Hsu et al., "A 0.8 V intelligent vision sensor with tiny convolutional neural network and programmable weights using mixed-mode processing-in-sensor technique for image classification," *IEEE J. Solid-State Circuits*, vol. 58, no. 11, pp. 3266–3274, Nov. 2023.
- [9] M. Lefebvre, L. Moreau, R. Dekimpe, and D. Bol, "7.7 A 0.2-to-3.6 TOPS/W programmable convolutional imager SoC with in-sensor current-domain ternary-weighted MAC operations for feature extraction and region-of-interest detection," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2021, pp. 118–120.
- [10] H. Song, S. Oh, J. Salinas, S. Park, and E. Yoon, "A 5.1 ms low-latency face detection imager with in-memory charge-domain computing of machine-learning classifiers," in *Proc. Symp. VLSI Circuits*, Jun. 2021, pp. 1–2.
- [11] M. Lefebvre and D. Bol, "A mixed-signal near-sensor convolutional imager SoC with charge-based 4b-weighted 5-to-84-TOPS/W MAC operations for feature extraction and region-of-interest detection," in *Proc. IEEE Custom Integr. Circuits Conf. (CICC)*, Apr. 2024, pp. 1–2.



- [12] B. Gönen, F. Sebastiano, R. Quan, R. van Veldhoven, and K. A. A. Makinwa, "A dynamic zoom ADC with 109-dB DR for audio applications," *IEEE J. Solid-State Circuits*, vol. 52, no. 6, pp. 1542–1550, Jun. 2017.
- [13] I. Park, W. Jo, C. Park, B. Park, J. Cheon, and Y. Chae, "A  $640 \times 640$  fully dynamic CMOS image sensor for always-on operation," *IEEE J. Solid-State Circuits*, vol. 55, no. 4, pp. 898–907, Apr. 2020.
- [14] A. Kneip, M. Lefebvre, J. Verecken, and D. Bol, "IMPACT: A 1-to-4b 813-TOPS/W 22-nm FD-SOI compute-in-memory CNN accelerator featuring a 4.2-POPS/W 146-TOPS/mm<sup>2</sup> CIM-SRAM with multi-bit analog batch-normalization," *IEEE J. Solid-State Circuits*, vol. 58, no. 7, pp. 1871–1884, Jul. 2023.
- [15] J.-O. Seo, M. Seok, and S. Cho, "A 44.2-TOPS/W CNN processor with variation-tolerant analog datapath and variation compensating circuit," *IEEE J. Solid-State Circuits*, vol. 59, no. 5, pp. 1603–1611, May 2024.
- [16] B. P. Ginsburg and A. P. Chandrakasan, "An energy-efficient charge recycling approach for a SAR converter with capacitive DAC," in *Proc. IEEE Int. Symp. Circuits Syst.*, Jul. 2005, pp. 184–187.
- [17] Y. Li and Y. Lian, "Improved binary-weighted split-capacitive-array DAC for high-resolution SAR ADCs," *Electron. Lett.*, vol. 50, no. 17, pp. 1194–1195, Aug. 2014.
- [18] P. Harpe, "A compact 10-b SAR ADC with unit-length capacitors and a passive FIR filter," *IEEE J. Solid-State Circuits*, vol. 54, no. 3, pp. 636–645, Mar. 2019.
- [19] N. R. Shanbhag and S. K. Roy, "Benchmarking in-memory computing architectures," *IEEE Open J. Solid-State Circuits Soc.*, vol. 2, pp. 288–300, 2022.
- [20] B. Moons, D. Bankman, L. Yang, B. Murmann, and M. Verhelst, "Binar-Eye: An always-on energy-accuracy-scalable binary CNN processor with all memory on chip in 28 nm CMOS," in *Proc. IEEE Custom Integr. Circuits Conf. (CICC)*, Apr. 2018, pp. 1–4.
- [21] H. Valavi, P. J. Ramadge, E. Nestler, and N. Verma, "A 64-tile 2.4-Mb in-memory-computing CNN accelerator employing charge-domain compute," *IEEE J. Solid-State Circuits*, vol. 54, no. 6, pp. 1789–1799, Jun. 2019.
- [22] B. Jeong, J. Lee, J. Choi, M. Song, Y. Son, and S. Y. Kim, "A 0.57 mW@1 FPS in-column analog CNN processor integrated into CMOS image sensor," *IEEE Access*, vol. 11, pp. 61082–61090, 2023.
- [23] M. Abedin, A. Roohi, M. Liehr, N. Cady, and S. Angizi, "MR-PIPA: An integrated multilevel RRAM (HfO<sub>x</sub>)-based processing-in-pixel accelerator," *IEEE J. Explor. Solid-State Comput. Devices Circuits*, vol. 8, pp. 59–67, 2022.
- [24] X. Wang, Z. Huang, T. Liu, W. Shi, H. Chen, and M. Zhang, "6.9 A 0.35 V 0.367 TOPS/W image sensor with 3-layer optical-electronic hybrid convolutional neural network," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2024, pp. 116–118.
- [25] R. Eki et al., "9.6 A 1/2.3 inch 12.3 Mpixel with on-chip 4.97 TOPS/W CNN processor back-illuminated stacked CMOS image sensor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, vol. 64, Feb. 2021, pp. 154–156.



**Martin Lefebvre** (Member, IEEE) received the M.Sc. and Ph.D. degrees in engineering sciences from the Université catholique de Louvain, Louvain-la-Neuve, Belgium, in 2017 and 2024, respectively.

He is currently a post-doctoral researcher with the cognitive sensor nodes and systems (CogSys) laboratory led by Prof. C. Frenkel at Delft University of Technology, Delft, The Netherlands, working on neuromorphic hardware/software co-design for efficient on-chip learning. His research interests also include mixed-signal vision chips for embedded

image processing and low-power current references.

Dr. Lefebvre serves as a reviewer for various IEEE journals and conferences, including IEEE JOURNAL OF SOLID-STATE CIRCUITS and IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS I AND II.



**David Bol** (Senior Member, IEEE) received the Ph.D. degree in engineering science from the Université catholique de Louvain, Louvain-la-Neuve, Belgium, in 2008, in the field of ultra-low-power digital nanoelectronics.

In 2005, he was a visiting Ph.D. student with CNM, Seville, Spain. In 2009, he was a post-doctoral researcher with intoPIX, Louvain-la-Neuve. In 2010, he was a visiting post-doctoral researcher with the laboratory for manufacturing and sustainability, UC Berkeley, Berkeley, CA, USA. In 2015,

he participated to the creation of e-peas semiconductors spin-off company. He currently leads the electronic circuits and systems (ECS) group, UCLouvain, focused on ultra-low-power design of integrated circuits for environmental and biomedical IoT applications, including computing, power management, sensing, and wireless communications, where he is currently an associate professor. He has authored more than 150 articles and conference contributions and holds three delivered patents. He is actively engaged in a social-ecological transition in the field of information and communication technologies (ICT) research with a post-growth approach.

Prof. Bol (co-)received five awards in IEEE conferences, such as ICCD 2008, SOI Conference 2008, FTFC 2014, ISCAS 2020, and ESSCIRC 2022, and supervised the Ph.D. thesis of Charlotte Frenkel who received the 2021 Nokia Bell Scientific Award and the 2021 IBM Innovation Award for her Ph.D. degree. He serves as a reviewer for various IEEE journals and conferences and presented several keynotes in international conferences.