



**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2022

FPGA Implementation of An Event-driven Saliency-based Selective Attention Model

Asgari, Hajar ; Risi, Nicoletta ; Indiveri, Giacomo

DOI: <https://doi.org/10.1109/biocas54905.2022.9948687>

Posted at the Zurich Open Repository and Archive, University of Zurich
ZORA URL: <https://doi.org/10.5167/uzh-231230>
Conference or Workshop Item

Originally published at:

Asgari, Hajar; Risi, Nicoletta; Indiveri, Giacomo (2022). FPGA Implementation of An Event-driven Saliency-based Selective Attention Model. In: IEEE Biomedical Circuits and Systems Conference (BioCAS), Taipei, Taiwan, 13 October 2022 - 15 October 2022.

DOI: <https://doi.org/10.1109/biocas54905.2022.9948687>

FPGA Implementation of An Event-driven Saliency-based Selective Attention Model

Hajar Asgari

Institute of Neuroinformatics
University of Zurich and ETH Zurich
Zurich, Switzerland
hajar@ini.uzh.ch

Nicoletta Risi

Bio-Inspired Circuits and Systems Lab
University of Groningen
Groningen, Netherlands
n.risi@rug.nl

Giacomo Indiveri

Institute of Neuroinformatics
University of Zurich and ETH Zurich
Zurich, Switzerland
giacomo@ini.uzh.ch

Abstract—Artificial vision systems of autonomous agents face very difficult challenges, as their vision sensors are required to transmit vast amounts of information to the processing stages, and to process it in real time. One first approach to reduce the data transmission is to use event-based vision sensors, whose pixels produce events only when there are changes in the input. However, even for event-based vision, transmission and processing of visual data can be quite onerous. Currently these challenges are solved by using high-speed communication links and powerful machine vision processing hardware. But if resources are limited, instead of processing all the sensory information in parallel, an effective strategy is to divide the visual field into several small sub-regions, choose the region of highest saliency, process it and shift serially the focus of attention to regions of decreasing saliency. This strategy, commonly used also by the visual system of many animals, is typically referred to as “selective attention”. Here we present a digital architecture implementing a saliency-based selective visual attention model for processing asynchronous event-based sensory information received from a Dynamic Vision Sensor (DVS). For ease of prototyping, we use a standard digital design flow and map the architecture on a Field Programmable Gate Array (FPGA). We describe the architecture block diagram highlighting the efficient use of the available hardware resources demonstrated through experimental results exploiting a hardware setup where the FPGA interfaced with the DVS camera.

Index Terms—Neuromorphic engineering, selective visual attention, event-driven, foveal vision, FPGA.

I. INTRODUCTION

Event-based vision sensors are becoming increasingly popular for a wide range of vision processing applications that require low latency and low power consumption [1]. As these constraints are the same ones faced by biological vision systems, it can be useful to study the computational strategies used by animal vision systems to optimize the performance of artificial event-based ones. In particular, a prominent strategy that biology uses to carry out efficient vision processing with limited resources is to use *selective attention* [2]–[4]. Instead of simultaneously processing all the information that can be provided by the input sensors, this strategy serially selects *salient* sub-regions of the input space shifting from sub-region

to sub-region, and processing them sequentially. These sub-regions are selected by choosing the region of the visual space that has the highest value in the corresponding “saliency map”. In Itti’s bottom-up selective attention model [4], saliency maps are computed by combining multiple visual features, such as color contrast, texture, motion, etc. [4].

As one of the most relevant features in these saliency maps is given by motion cues and dynamic changes in the visual scene (e.g., produced by stationary but flashing LEDs), here we focus on the activity of the DVS sensor [5], which naturally correlates with these salient features [6]. The use of this sensor reduces the amount of data transferred to the processing areas and, consequently, the amount of power required to process it. Based on this, we present an event-based visual processing system that uses a biologically plausible computational model of a saliency-based form of bottom-up attention, to select and sequentially process sub-regions of the visual field. While early attempts to implement selective attention models in neuromorphic hardware have focused on mixed-signal analog/digital implementations [7]–[11], few have been done using a standard digital design flow [12], [13]. Here we implement a saliency-based event-driven attention mechanism at pixel level which depend on the exponential decay of the pixel states. We use digital FPGA devices for fast prototyping, while still allowing real-time processing of events streamed from the DVS. The active-vision system that we present carries out computation in parallel using an event-based asynchronous communication infrastructure that employs the Address-Event Representation (AER) [14]–[16].

In Section II we present the details of the selective attention model. Section III describes the proposed event-based attention mechanism. In Section IV we introduce the hardware architectures, and in Section V we presents the simulation and experimental results of the proposed system. Finally, Section VI concludes the paper.

II. BACKGROUND

Selective attention is a powerful computational strategy used by biological systems to optimize the use of the limited computing resources and minimize latency and reaction times. It acts by selecting the most salient location of the visual scene, ignoring distractors and irrelevant information, and allocating

the system’s computational resources for fast and accurate processing of the information in that region. Subsequently, the system chooses other regions of the visual scene with decreasing saliency, processing them in a sequential manner. The choice of the regions to process can be influenced by “bottom-up” stimuli (e.g., such as high contrast or moving objects) or “top-down” preferences (such regions of the scene that ignore the sky, or that include faces). Several computational models of selective attention under both top-down and bottom-up influences have been proposed [17]–[20]. The first explicit, biologically plausible computational model of bottom-up saliency-based selective visual attention was proposed by Koch and Ullman in 1985 [3].

Inspired by biological vision systems, the first processing stage in any bottom-up attention model is computing different features such as intensity, contrast, color opponency, orientation, direction, or motion [4]. The different feature maps contribute with different weights to create a unique saliency map. The strength of the feature contributions can be influenced by top-down modulation or training [4]. A Winner-Take-All (WTA) network selects the most salient location in the saliency map and focuses the system’s attention to it. As the region is being attended, an inhibition of return mechanism is used to suppress it, so that less salient regions can win the WTA competition. Depending on the parameters governing its dynamics, this mechanism can allow the system to switch only between the two most salient regions, the top three, and so on.

To test selective attention in real-time visual processing tasks, several event-driven saliency-based models have been proposed [20], [21], as well as their hardware implementations, using neuromorphic circuits [11], [22], [23], and using FPGAs [12], [13]. However, almost all the FPGA implementations made use of frame-based cameras [13], and very few (only one to our knowledge) interfaced to silicon retina devices [12]. In this paper we proposed a digital event-driven saliency-based attention model implemented on FPGA for prototyping to process event-streams of the DVS camera.

III. EVENT-BASED SALIENCY ATTENTION MODEL WITH TOP-DOWN MODULATION

Here we present our FPGA-compatible event-based selective attention architecture. As bottom-up selective attention models are especially appealing for hardware implementations, due to their modular and easily expandable nature [10], we focus the bulk of our work on these sub-sets of models. First, we describe how the model processes visual inputs from event-cameras and exploits the event-based nature of these cameras to compute saliency from high-contrast changing visual stimuli; then we extend the bottom-up saliency model to incorporate a top-down modulation mechanism, which allows steering the focus of attention (FOA) according to the target task.

A. Event-based Bottom-up saliency

As opposed to the original computational model [17], the system we designed uses only one feature channel that encodes the changes present in the visual scene, sensed by the DVS.

Event-camera pixels, such as those present in the DVS, operate independently and respond only to local changes in the light intensity, which makes them inherently act as local edge detectors of moving objects. In this respect, saliency is defined as the region of the visual scene that produces the highest number of events, which in turn corresponds to the area that has the highest number of high contrast moving objects.

Let the FOA be defined as a window of size $[m_x, m_y]$ centered around the most *salient* pixel. The *event* $e(x, y, t, p)$ represents the brightness change detected by the pixel with coordinates (x, y) at time t and with polarity p (with $p=1$ for contrast increases, on ON-event, and $p=-1$, or OFF-event, for contrast decreases). Then, the most *salient* pixel at time t can be defined as follows:

$$P^*(t) = \max_p s_P(t)$$

with s_P being the pixel state at time t . See IV-C for a detailed description of the algorithm, which updates the pixels state in an event-based manner.

B. Top-down Biasing

After receiving a pixel event and extracting information about the pixel’s location, we can have a top-down bias ahead of the bottom-up processing. The top-down biasing can be executed in two methods: top-down gating and top-down modulation. In the “top-down gating” approach, all the areas out of a predefined region of interest are cut out. In the “top-down modulation” approach, all the arriving pixels are processed but the changing gain of the state for each pixel is modulated according to the predefined conditions for the top-down biasing.

IV. HARDWARE DESIGN AND IMPLEMENTATION

Fig. 1 shows a schematic of the proposed architecture, which consists of seven blocks. The Handshake Receiver (HSR) handles the input events. HSR implements the standard 4-phase handshake AER protocol between the sender and the receiver block. The output events are monitored through a USB connection but can also be sent to the other event-based devices using a Handshake Sender (faded HSS block in Fig. 1). To prevent meta-stable states at the interface between different clock domains, a double flip-flop synchronizer (FMS) is applied before the input handshake. The core blocks of the architecture are described in more detail in the following sections.

A. Data Processing Element(DPE)

The event-camera included in our architecture is the DAVIS240C which employs a serial data format, i.e. x and y coordinates of the events are sent one after the other [5], [24]. Thus, the Data Processing Element block (or DPE) takes care of generating the 2D event coordinate by merging x and y input words and forwards the resulting pixel_ID to the receiver block. Two instances of this block are used, which differ in the spatial resolution of the output events.

- **DPE Exploring**, or DPE (Exp), which operates on the full input resolution, extracts the pixel_ID from the input events and forwards it to the Saliency block (SAL); SAL

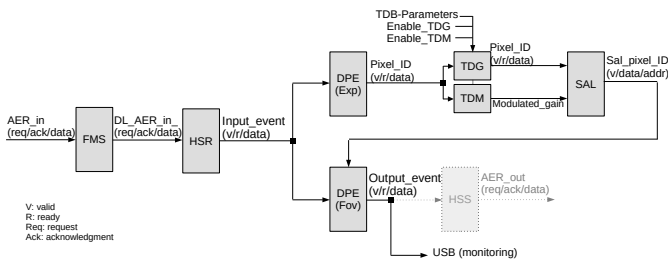


Fig. 1: High-level block diagram of the network architecture. FMS: Flip-flop Meta-stability Synchronizer. HSR: Handshake Receiver. DPE (Exp): Data Processing Element in Exploring Mode. DPE (Fov): Data Processing Element in Fovea Mode. TDG: Top-Down Gating. TDM: Top-Down Modulation. SAL: Saliency block. HSS (faded block): Handshake Sender.

defines the most salient pixel according to the input events and a history of the previous events.

- DPE Fovea, or DPE (Fov), which receives the coordinate of the most salient pixel by the SAL block and extracts only the events inside the FOA.

B. Top-Down Biasing Blocks

After exploring the input events and extracting the pixel_ID, the top-down mechanism works on top of the bottom-up saliency. In the proposed system, the top-down biasing mechanism is implemented using two different methods. Both are explained in more detail in III-B.

1) *Top-Down Gating (TDG)*: This block receives the extracted pixel_ID from DPE (Exp) block. When this block is enabled, via the input parameter Enable_TDG, input events cross the TDG only if they meet the input top-down biasing conditions (TDB_parameters).

2) *Top-Down Modulation (TDM)*: TDM also receives the pixel_IDs from the DPE (Exp). This block handles all the input events with a modulated state-changing gain. The amplitude of the modulation depends on the top-down biasing conditions. The highest modulation happens when the input events meet the TDM requirements.

C. Saliency block (SAL)

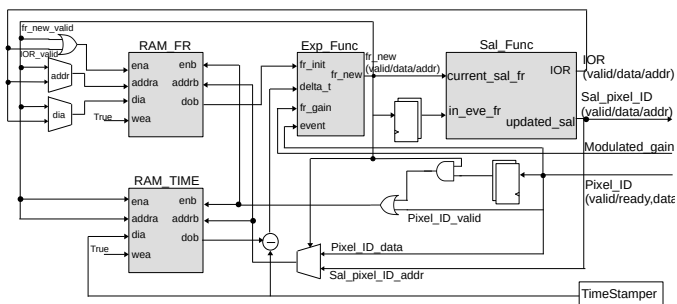


Fig. 2: High-level block diagram of the saliency block (SAL). This block receives the input pixel_ID and determines the most salient pixel.

Fig. 2 illustrates the core blocks of the Saliency Block (or SAL) architecture. SAL extracts the location of the most salient pixel. In the SAL architecture two separate Dual Port BRAM are used: RAM_FR stores the pixel state $s_P(t)$ as 21-bits registers (with fixed-point representation: $s=1$ bit for the sign, $i=12$ integer bits, and $f=8$ fractional bits); RAM_TIME stores timestamp of the last event generated at each pixel location.

Let P^* be the current most salient pixel location with state s^* , which is selected based on all the previous input events. At each incoming event from pixel P at time point t , the SAL block performs the following steps:

- 1) Read the last pixel event time stamp t_{old} from the corresponding entry of the RAM_TIME.
- 2) Read the latest pixel state $s_{P(t_{old})}$ at the last pixel event timestamp t_{old} from the corresponding entry of the RAM_FR.
- 3) Compute the updated pixel state $s_P(t)$ as follows:

$$s_P(t) = 1 + s_{P(t_{old})} \cdot \exp[(t_{old} - t)/\tau] \quad (1)$$

with τ stored as input parameter. This operation is handled by Exp_Func block (Fig. 2), where the exponential decay term is implemented using the piece-wise linear approximation, to reduce the hardware resource utilization on the FPGA.

- 4) Store the new pixel state and timestamp in the corresponding BRAM entries.
- 5) Update the state of the current most salient pixel by repeating steps 1, 2, and 3 for P^* .
- 6) Compare the new pixel state to the state of the current most salient pixel P^* . If $s_P(t) > s^*$:

- Update the most salient pixel location, which is forwarded to the DPE(Fov) block.
- Excite the state of the updated most salient pixel as the center of FOV or P^* :

$$s[x, y] = s[x, y] + S_{plus}, \quad \forall (x, y) \in FOA \quad (2)$$

- Inhibit the state of all pixels in the previous FOA:

$$s[x, y] = s[x, y] - S_{minus}, \quad \forall (x, y) \in FOA \quad (3)$$

with S_{plus} and S_{minus} stored as input parameters for the highest excitation and inhibition values, respectively.

By penalizing the pixels in the previously selected salient region, the last step adapts the Inhibition of Return (IOR) mechanism proposed in [4] to our event-based implementation.

V. RESULTS AND DISCUSSION

The proposed system implements a digital electronic event-driven foveal vision. All parts of the design are described using the standard top-bottom digital design flow.

As a proof of concept we implemented our electronic Foveal vision for processing event-camera inputs on a Kintex-7 XC7KT160T FPGA, which is hosted on the Opal Kelly XEM7360 board. Based on the results from Table I, the entire

proposed network uses only less than 1 percent of the available Slice FFs and LUTRAMs and DSPs, 2.31 percent of available Slice LUTs and 38 percent of available BRAMs. Furthermore, since the whole pipeline operates directly on the hardware, the system’s latency is in the order of a few microseconds.

TABLE I: Resource utilization of the proposed digital electronic Foveal vision on Kintex-7(XC7KT160T).

	Proposed	Available	Utilization
Slice FFs	1646	202800	0.81%
Slice LUTs	3240	102400	2.31%
LUTRAMs	33	35000	0.09%
BRAMs	123.5	325	38%
DSPs	2	600	0.33%
Device	Kintex.7		

To validate the architecture in more real-world scenarios, the design was tested with samples from the DVS128 Gesture Dataset [25], which contains DVS recordings from different hand gestures and illumination conditions. Fig. 3 shows the trajectory of the most salient locations of a sample from the *left hand clockwise* gesture. Fig. 3-(Left) shows the time surface plot of the input events collected within a time window ($t=[0;43.1]$ ms). The darker events are the most recent ones. The pixels shown in red are the most salient pixels selected across time. Fig. 3-(Right) shows the output of the pipeline. Events are confined within a shifting 16×16 window, which corresponds to the most salient region at a given time point.




Fig. 3: Time surface plots of the input (Left) and output (Right) of the system.

To examine the behavior of the proposed architecture in terms of the ability to attend to different salient locations, a controlled data stream was used to assess the system behavior. The data consists of 3 active pixels with different firing rates (depicted in gray scale levels in Fig. 4). The architecture attends first to the pixel with the highest firing rate (dark gray point). Then it progressively shifts to the other active locations with decreasing firing rate (from gray to light gray regions). Thus, this result highlights how the proposed system can effectively shift the focus of attention towards different salient locations according to the local instantaneous firing rate.

Fig. 5 illustrates the results with a top-down biasing condition. Exploring the events of half of the screen based on the




Fig. 4: Examining the forgetting ability of the proposed system by applying a controlling input data stream. Darker colors show higher firing rates.

application demands is an example of top-down biasing conditions. In this experiment, we have used input recordings from the DAVIS240. Unlike in the previous scenarios, here the pipeline is only attending to the most salient locations on the upper-half (Fig. 5a) and left-half (Fig. 5b) parts of the scene. The fovea window stays at one location for a limited time window (indicated on top of each plot), until the next most salient location is selected.




Fig. 5: Bottom-up attention mechanism with two different top-down biasing parameters. The region of interest for finding the most salient location is: (a) upper-half. (b) left-half.

VI. CONCLUSION

We presented an implementation of a digital bio-inspired event-driven foveal vision on an FPGA. The system implements a simplified model of the bottom-up stimulus-driven saliency-based selective-attention mechanism. The proposed network has been synthesized and as a proof of concept implemented on a Xilinx Kintex-7 FPGA device in a setup containing a DAVIS240 camera connected to a dedicated FPGA. We demonstrated how the system can sequentially select the spatial locations of the most salient inputs present in the camera’s field of view with experimental results. The system receives input signals in the form of address events, selects the spatial locations of the most salient inputs present in the camera’s field of view, and only transmits events from the selected sub-region. This reduces the amount of data that needs to be processed, e.g., to extract features or detect specific targets, and enables the construction of a complex saliency-based vision processing system. In particular, it enables further neuromorphic event-based processing chips to process high-resolution images by using their limited resources (e.g., cores of 256 neurons [26], [27]) sequentially, as it happens in biology.

REFERENCES

- [1] G. Gallego, T. Delbrück, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. J. Davison, J. Conradt, K. Daniilidis, *et al.*, “Event-based vision: A survey,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 1, pp. 154–180, 2020.
- [2] G. Rizzolatti, *Mechanisms of Selective Attention in Mammals*, pp. 261–297. Boston, MA: Springer US, 1983.
- [3] C. Koch and S. Ullman, “Shifts in selective visual-attention – towards the underlying neural circuitry,” *Human Neurobiology*, vol. 4, no. 4, pp. 219–227, 1985.
- [4] L. Itti and C. Koch, “Computational modeling of visual attention,” *Nature Reviews Neuroscience*, vol. 2, no. 3, pp. 194–203, 2001.
- [5] “Product of inilabs: Dynamic vision sensor.” <http://www.inilabs.com/products/dynamic-vision-sensors>, 2009.
- [6] J. Thompson, “Attention, biological motion, and action recognition,” *Neuroimage*, vol. 59, no. 1, pp. 4–13, 2012.
- [7] T. Horiuchi, T. Morris, C. Koch, and S. DeWeerth, “Analog VLSI circuits for attention-based, visual tracking,” in *Advances in Neural Information Processing Systems (NIPS)* (M. Mozer, M. Jordan, and T. Petsche, eds.), vol. 9, pp. 706–712, MIT Press, 1997.
- [8] T. Morris, T. Horiuchi, and S. DeWeerth, “Object-based selection within an analog VLSI visual attention system,” *IEEE Transactions on Circuits and Systems II*, vol. 45, no. 12, pp. 1564–1572, 1998.
- [9] G. Indiveri, “Modeling selective attention using a neuromorphic analog VLSI device,” *Neural Computation*, vol. 12, no. 12, pp. 2857–2880, 2000.
- [10] G. Indiveri, “A neuromorphic VLSI device for implementing 2-D selective attention systems,” *IEEE Transactions on Neural Networks*, vol. 12, pp. 1455–1463, Nov. 2001.
- [11] C. Bartolozzi and G. Indiveri, “Selective attention in multi-chip address-event systems,” *Sensors*, vol. 9, no. 7, pp. 5076–5098, 2009.
- [12] B. Kim, H. Okuno, T. Yagi, and M. Lee, “Implementation of visual attention system using artificial retina chip and bottom-up saliency map model,” in *Neural Information Processing* (B.-L. Lu, L. Zhang, and J. Kwok, eds.), (Berlin, Heidelberg), pp. 416–423, Springer Berlin Heidelberg, 2011.
- [13] J. L. Molin, C. S. Thakur, E. Niebur, and R. Etienne-Cummings, “A neuromorphic proto-object based dynamic visual saliency model with a hybrid fpga implementation,” *IEEE Transactions on Biomedical Circuits and Systems*, vol. 15, no. 3, pp. 580–594, 2021.
- [14] J. Lazzaro, J. Wawrzynek, M. Mahowald, M. Sivilotti, and D. Gillespie, “Silicon auditory processors as computer peripherals,” *IEEE Transactions on Neural Networks*, vol. 4, pp. 523–528, 1993.
- [15] K. Boahen, “Multiple pathways: Retinomorph chips that see quadruple images,” in *Proceedings of the Seventh International Conference on Microelectronics for Neural, Fuzzy and Bio-inspired Systems; Microneuro’99*, (Los Alamitos, CA), pp. 12–20, IEEE Computer Society, Apr. 1999.
- [16] S. Deiss, R. Douglas, and A. Whatley, “A pulse-coded Communications infrastructure for neuromorphic systems,” in *Pulsed Neural Networks* (W. Maass and C. Bishop, eds.), ch. 6, pp. 157–78, MIT Press, 1998.
- [17] L. Itti, E. Niebur, and C. Koch, “A model of saliency-based visual attention for rapid scene analysis,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, no. 11, pp. 1254–1259, 1998.
- [18] S. Park, S. Ban, S. J., L. Minho, O. Kaynak, E. Alpaydin, E. Oja, and X. Lei, “Implementation of visual attention system using bottom-up saliency map model,” in *Lecture notes in computer science, Artificial neural networks and neural information processing - ICANN / ICONIP* (B. Springer, ed.), vol. 2714, pp. 678–685, June 2003.
- [19] A. Borji and L. Itti, “State-of-the-art in visual attention modeling,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 1, pp. 185–207, 2013.
- [20] T. Uejima, E. Niebur, and R. Etienne-Cummings, “Proto-object based saliency model with texture detection channel,” *Frontiers in Computational Neuroscience*, vol. 14, 2020.
- [21] H. Akolkar, D. R. Valeiras, R. Benosman, and C. Bartolozzi, “Visual-auditory saliency detection using event-driven visual sensors,” in *International Conference on Event-based Control, Communication, and Signal Processing (EBCCSP)*, pp. 1–6, 2015.
- [22] C. S. Thakur, J. L. Molin, T. Xiong, J. Zhang, E. Niebur, and R. Etienne-Cummings, “Neuromorphic visual saliency implementation using stochastic computation,” in *International Symposium on Circuits and Systems, (ISCAS)*, pp. 1–4, IEEE, May 2017.
- [23] F. Galluppi, K. Brohan, S. Davidson, T. Serrano-Gotarredona, J. Carasco, B. Linares-Barranco, and S. Furber, “A real-time, event-driven neuromorphic system for goal-directed attentional selection,” in *Neural Information Processing*, pp. 226–233, Springer, 2012.
- [24] C. Brandli, R. Berner, M. Yang, S.-C. Liu, and T. Delbruck, “A 240×180 130 dB 3 μs latency global shutter spatiotemporal vision sensor,” *IEEE Journal of Solid-State Circuits*, vol. 49, no. 10, pp. 2333–2341, 2014.
- [25] “Dvs128 gesture dataset.” IBM Website.
- [26] P. A. Merolla, J. V. Arthur, R. Alvarez-Icaza, A. S. Cassidy, J. Sawada, F. Akopyan, B. L. Jackson, N. Imam, C. Guo, Y. Nakamura, B. Brezzo, I. Vo, S. K. Esser, R. Appuswamy, B. Taba, A. Amir, M. D. Flickner, W. P. Risk, R. Manohar, and D. S. Modha, “A million spiking-neuron integrated circuit with a scalable communication network and interface,” *Science*, vol. 345, pp. 668–673, Aug. 2014.
- [27] S. Moradi, N. Qiao, F. Stefanini, and G. Indiveri, “A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAPs),” *Biomedical Circuits and Systems, IEEE Transactions on*, vol. 12, pp. 106–122, Feb. 2018.