Chapter 5

NEUROMORPHIC ENGINEERING

Giacomo Indiveri
Institute of Neuroinformatics, UNI- ETH Zurich, Switzerland, giacomo@ini.phys.ethz.ch

Abstract: Neuromorphic circuits are analog circuits that implement models of biological systems for sensory processing. Because the neuromorphic circuits share the same physical constraints as their biological counterparts, they have similar organizational structures, and use similar strategies for optimizing robustness to noise, and power consumption. Silicon retinas and other neuromorphic vision chips, built as hardware models of biological vision systems, represent efficient artificial sensory pre-processors. General processing networks which use detailed models of neurons are also being investigated. Simple neuromorphic systems have been built, using networks of silicon neurons and neuromorphic chips as front-ends for pre-processing incoming sensory signals. In this chapter we present two representative case studies describing a single-chip vision sensor and a multi-chip processing network, that contain most of the characteristic elements found in today's neuromorphic systems.

Key words: Neuromorphic, analog VLSI, silicon retina, winner-take-all, WTA, Address-Event Representation, AER, Integrate-and-Fire neuron.

5.1. INTRODUCTION

Neural network theories, used as an additional methodology for solving pattern recognition and constraint minimization problems, have emerged in recent years as a practical technology and represent a well established research field. Neural network algorithms, the type of non-linearities present in the transfer functions of their computational elements and the architectures that implement them are often loosely inspired by biological systems.

An emerging new technology which tries to establish even closer links to biology, capitalizing on the advantages of interdisciplinary research, is the one of *neuromorphic engineering*. Specifically, neuromorphic engineering

applies the computational principles discovered in biological organisms to those tasks that biological systems perform easily, but which have proved difficult to do using traditional engineering techniques. For example, biological neural systems for sensory perception and motor control are compact, energy efficient and robust to noise both in the input data and in the internal state variables. They typically have a relatively simple organization, consisting of arrays of similar processing elements that interact in nonlinear ways mainly with nearest neighbours. Neuromorphic systems, rather than implementing abstract neural networks remotely related to these types of systems, are hardware devices, containing analog circuits, that attempt to model in detail, (up to the device-physics level) their properties and the physical processes in them embedded that underlie neural computation [1]. The closest medium, widely accessible to the research community, that allows researchers to implement detailed hardware models of neural systems is *silicon*. Using analog, continuous time circuits implemented with a standard CMOS VLSI technology it is possible to build low-cost, compact implementations of such models. The greatest successes of neuromorphic analog VLSI (aVLSI) to date have been in the emulation of peripheral sensory transduction: Silicon retinas and silicon cochleas have been successfully implemented and used in a wide variety of applications [2-4]. In these analog devices, as in their biological counterparts, it is the structure of the architecture, the morphology of the system, that determines their functionality. This constraint is added to the ones that come from the fact that neuromorphic systems have to cope with issues such as minimizing power consumption, maximizing robustness to noise and optimizing reliability in their performance, while interacting in real-time with the environment. It is by trying to satisfy these very constraints that researchers are hoping to obtain more insight into the workings of biological neural systems.

Neuromorphic engineering is thus mainly concerned with *hardware* correlates of biological systems. The nature of the research carried out by neuromorphic engineers is twofold: on one side there is the desire to learn more about the computational properties of the brain by tackling the same problems that nature and evolution solved in the course of 600 million years, on the other there is the desire to design and develop efficient neuromorphic engineered systems that can be used to solve real world problems and that can eventually lead to successful industrial applications.

Rather than giving a brief overview of the progressively increasing number of neuromorphic devices, systems and applications present in the literature, in the next sections I will describe two specific examples of neuromorphic systems that capture main essence of this field: a single chip neuromorphic vision system and a multi-chip network of integrate-and-fire

neurons. The former is a small compact low-power focal-plane processor that implements at an abstract level some of the principal visual processing stages observed in the primary visual cortex of mammals for visual tracking applications, with potential commercial valence. The latter on the other hand, is the basis for implementing reconfigurable multi-chip neural network architectures, applied to basic scientific investigation and eventually to signal processing tasks.

5.2. SINGLE-CHIP NEUROMORPHIC VISION SENSORS

Neuromorphic vision sensors are typically analog VLSI devices that implement hardware models of biological visual systems and that can be used for machine vision tasks [2,3]. It is only recently that these hardware models have become elaborate enough for use in a variety of engineering applications [5]. These types of devices and systems offer an attractive, low cost alternative to special purpose DSPs for machine vision tasks. They can be used either for reducing the computational load on the digital system in which they are embedded or, ideally, for carrying out all of the necessary computation without the need of any additional hardware. They process images directly at the focal plane level. Typically each pixel contains local circuitry that performs in real time different types of spatio-temporal computations on the continuous analog brightness signal. In contrast CCD cameras or conventional CMOS imagers merely *measure* the brightness at the pixel level, eventually adjusting their gain to the average brightness level of the whole scene. In neuromorphic vision chips, photoreceptors, memory elements and computational nodes share the same physical space on the silicon surface. The specific computational function of a neuromorphic sensor is determined by the structure of its architecture and by the way its pixels are interconnected. Since each pixel processes information based on locally sensed signals and on data arriving from its neighbours, the type of computation being performed is fully parallel and distributed. Another important feature is the asynchronous operation of neuromorphic sensors, which is preferable to clocked operation for sensory processing, given the continuous nature of sensory signals. Clocked systems introduce temporal aliasing artifacts that can significantly compromise the time-dependent computations performed in real-time sensory processing systems.

Visual tracking

Tracking features of interest as they move in the environment is a computationally demanding task for machine vision systems. The control loop of active vision systems, comprising motors that steer the visual sensor, relies on the speed of the specific computation carried out. The stability of the system depends on the latency of the sensory-motor control loop itself. To reduce this latency and improve the performance of the active vision system several custom VLSI sensors that pre-process the input image and extract the position of the target, have been proposed [6-8]. Here we describe a tracking system that, as previously proposed solutions, reduces the computational cost of the processing stages interfaced to it by carrying out an extensive amount of computation at the focal plane itself, and transmitting only the result of this computation, rather than extensive amounts of data representing the raw input image. Despite the principle of the approach here followed is very similar in nature to the one followed by the authors cited above, the tracking architecture here described differs from previously proposed ones in two key features: it selects high-contrast edges independent of the absolute brightness of the scene (as opposed to simply selecting the scene's brightest [6-8]); it uses a hysteretic winner-take-all (WTA) network, with positive feedback and lateral coupling [9], to lock-onto and smoothly track the selected targets (different from WTA networks used in other tracking device [6,8]).
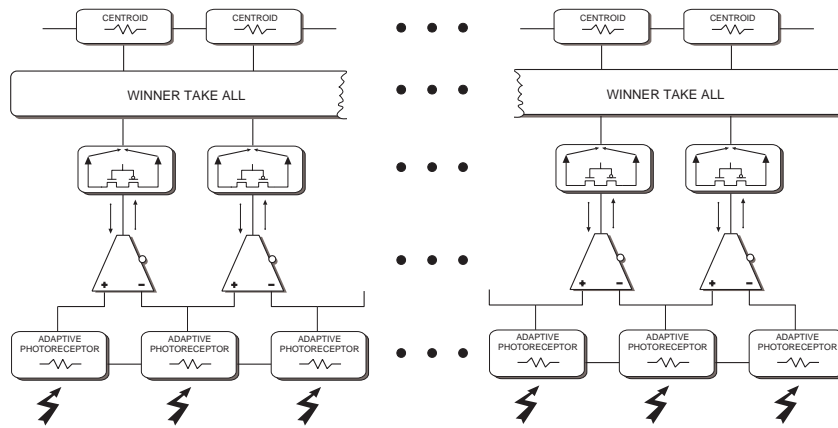

Figure 5-1. Block diagram of single-chip tracking system. Spatial edges are detected at the first computational stages by adaptive photoreceptors connected to spatial-derivative circuits. A winner-take-all network selects the edge with strongest contrast and   a position-to-voltage circuit encodes its position with a single continuous analog voltage.
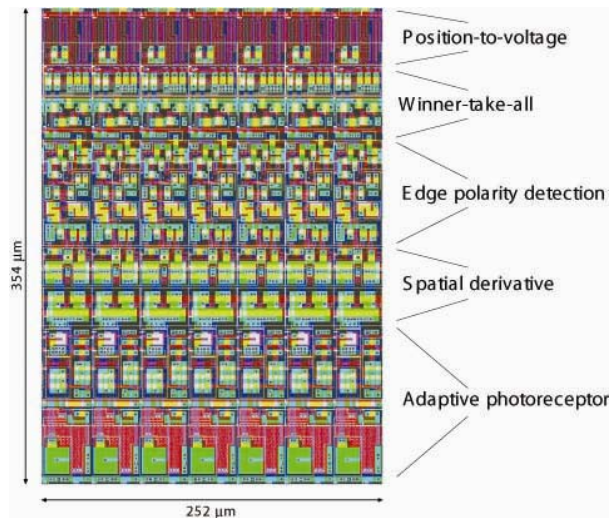
Figure 5-2. Portion of layout of containing 7 processing columns. The size of each computational stage is evidenced on the right.

These features allow systems that use this architecture to reliably track edges with the highest contrast present in the sensor's field of view in a wide variety of illumination conditions.

The tracking chip

The tracking architecture here described is a one-dimensional array of "intelligent pixels"', structured in a hierarchical way that can be implemented on a single chip device. As the architecture is one-dimensional, each pixel is a long column of circuits. This allows the design to minimize area usage and to maximize the number of pixels on the device.

Image brightness data is processed in parallel through five main computational stages. A block diagram of the device's architecture is depicted in Fig. 5.1. The first stage is an array of adaptive photoreceptors [10] that map logarithmically image intensity into their output voltages. The second stage is composed of circuital blocks that perform spatio-temporal processing on the input signal to extract high-contrast moving edges. At the third computational stage the polarity of each edge is gated so that the sensor selectively responds either to ON edges (dark to bright transitions), or to OFF edges (bright to dark transitions) or to both. The fourth stage uses a hysteretic winner-take-all (WTA) network [9] which selects and locks onto the feature with strongest spatial contrast moving at the speed that best matches the photoreceptor's velocity tuning. Finally in the last stage there is a position-to-voltage circuit [11], that encodes the spatial position of the

WTA network's output with a single analog value. An example of a tracker chip layout is shown in Fig. 5.2.

   Adaptive Photoreceptor Circuits. These types of photoreceptor circuits [10] have been used extensively in many neuromorphic sensors.   The response of the circuit is invariant to absolute light intensity, and changes logarithmically with image brightness. Fig. 5.3(a) shows the circuit diagram of such type of photoreceptors. It consists of a photodiode D in series with a transistor M_1 in source-follower configuration, and a negative feedback loop from the source to the gate of M_1.   The feedback loop consists of a high-gain inverting amplifier in common-source configuration (M_2,M_3), [10] and a resistive element R_1.   The transient output voltage change dV_{out} is amplified with respect to the DC output voltage change dV_{fb} by the capacitive divider ratio $A_C \equiv (C_1+C_2)/C_2$, as long as the adaptation effect can be neglected.
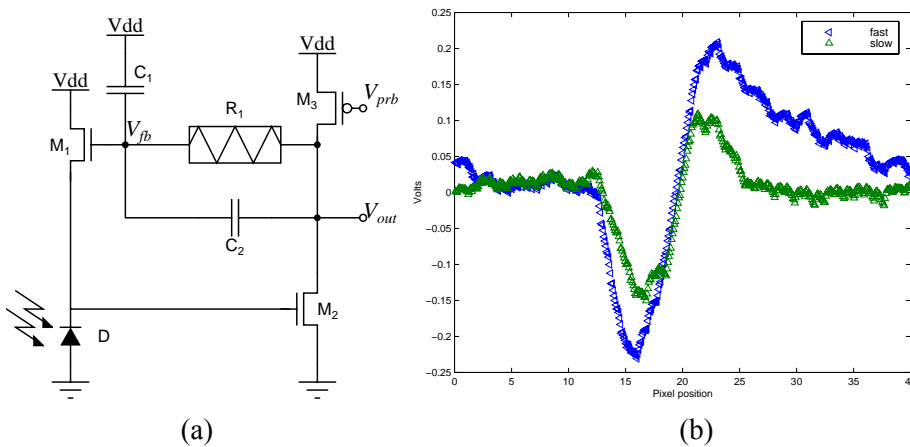

Figure 5-3. (a) Adaptive photoreceptor circuit. (b) Response of an array   of adaptive photoreceptors, with fast adaptation rate, to a dark   bar on a white background moving from right to left with an on-chip   speed of 31mm/s (left pointing triangles) and at a slightly slower speed (upward pointing triangles)

   The photo sensor adapts to variations in the photo current on a long time scale, which usually reflect slow changes in the background illumination that are typical of natural lighting conditions. The adaptation state is represented by the charge Q_{fb} stored on the capacitor plates of the feedback node. The output voltage V_{out} depends on this adaptation state and on the input signal represented by V_{fb}.   The adaptation dynamics are determined by the characteristics of the resistive element. A large resistance results in slow adaptation: a small resistance in fast adaptation. A linear resistor with a resistance R makes V_{out} adapt exponentially after an irradiance step. However, with the specific capacitance and sheet resistance values provided

by typical semiconductor technology the decay time constants achieved with linear resistors would be too small for practical applications. The resistive elements are therefore typically built with active elements, which can be operated at low currents that are matched to typical capacitance values [12].

Fig. 5.3(b) shows the response of an array of adaptive photoreceptors, containing a ``tobi-element'' [10] as a resistive element, to a moving bar. The photoreceptor's adaptation rate was set to be very high, such that the photoreceptors adapt quickly to brightness transients. Because of its adaptation property, this particular type of photoreceptor has a response with both contrast and speed dependence.

Spatial derivative circuits and edge-polarity detector circuits. Spatial derivative can be implemented using simple transconductance amplifiers operated in the subthreshold regime. The amplifiers receive input voltages from neighboring photoreceptors and provide a bidirectional output current that is proportional to the hyperbolic tangent of their differential input [12]. The output current saturates smoothly as the differential voltage increases (in absolute value) beyond 200-300mV. The possibility of electronically smoothing the input image (at the adaptive-photoreceptors stage) allows the user to operate the spatial derivative circuit always in its linear range, for a stimulus with fixed spatial frequencies.
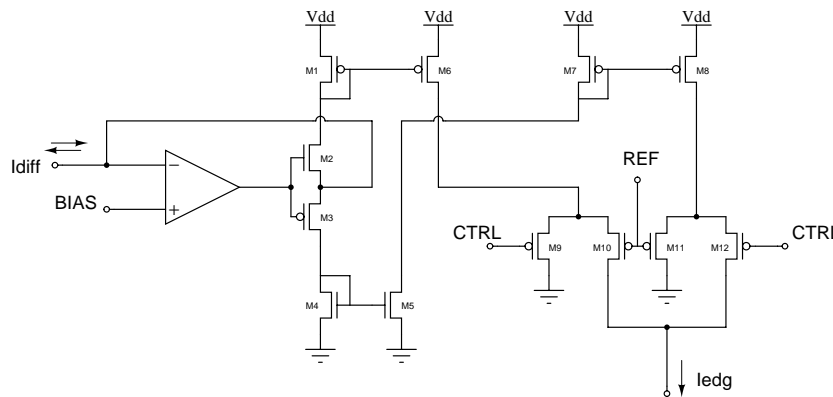

Figure 5-4. Circuit diagram of the current polarity detector. Positive I_{diff} currents are conveyed to the n-type current mirror M4,M5. Negative I_{diff} currents are conveyed to M6 through the p-type current mirror M1,M6. Depending on the values of the control voltage signals V_{CTRL} and V_{REF}, the output current I_{edg} represents a copy of only one of the two polarities of I_{diff}, or of both polarities of I_{diff} (see text for details).

The polarity of edges in the visual scene is encoded by the sign of the transconductance amplifiers' currents. Each of these currents is fed into a *current-conveyor* [13] of the type shown in Fig. 5.4. This circuit is used to

separate the positive component of the input current I_{diff} from the negative one, and to decouple the spatial derivative stage from the current-polarity selection stage. Negative input currents are conveyed to transistor M6, while positive ones are flipped through the current mirror M4,M5 and conveyed to M8. Transistors M6 and M8 source their currents to the polarity selection circuit (transistors M9-M12) [8]. Negative, positive, or combinations of both types of currents can be selected by properly setting the control voltages V_{CTRL} and V_{REF}. The output currents I_{edg} of all edge-polarity detector circuits are sourced, in parallel, to the elements of the next processing stage: a hysteretic winner-take-all network.

Hysteretic Winner-Take-All Network. CMOS implementations of winner-take-all (WTA) networks are an important class of circuits widely used in neural networks and pattern-recognition systems. They implement architectures that select one node, out of many, through a competition mechanism that depends on the amplitude of the architecture's input signals. Several types of WTA circuits have been proposed in the literature [9,14,15]. The WTA circuit proposed by Lazzaro *et al.* [15] optimizes power consumption and silicon area usage. It is ideal for applications that do not require high precision or high speed computation, such as sensory perception tasks. This circuit, proposed more than ten years ago, still remains one of the most compact and elegant designs of analog current-mode WTA circuits. It is asynchronous; it responds in real-time; and it processes all its input currents in parallel, using only two transistors per node, if the output signal is a voltage, and four transistors if the output signal is a current. Recently, some extensions to the basic design described in [15] have been proposed [8,9,16]. They endow the WTA circuit with local excitatory feedback [16] and with distributed hysteresis [8,9]. Local excitatory feedback enhances resolution and speed performance of the circuit, providing a hysteretic mechanism that withstands the selection of other potential winners unless they are stronger than the selected one by a set hysteretic current. Distributed hysteresis allows the winning input to shift between adjacent locations maintaining its winning status, without having to reset the network. These enhanced types of WTA networks are able to select and lock onto the input with strongest amplitude, and to track it as it shifts smoothly from one pixel to its neighbor [8,9].

Fig. 5.5 shows an example of the response of the tracker chip's hysteretic WTA network to a moving high-contrast bar. The top trace of the figure represents the net input current to the WTA network, and shows the effect of spatial smoothing of the sum of input currents with the hysteretic current from the winner's positive feedback loop. It is clear from this figure that the active winning cell is the one corresponding to pixel 26. The bottom trace

shows the instantaneous response of the adaptive photoreceptor array. The input stimulus was a 1 cm-wide black bar on a white background positioned at approximately 17 cm away from the focal plane and imaged onto the chip through a 4 mm lens moving from left to right with an on chip speed of 31 mm/s.
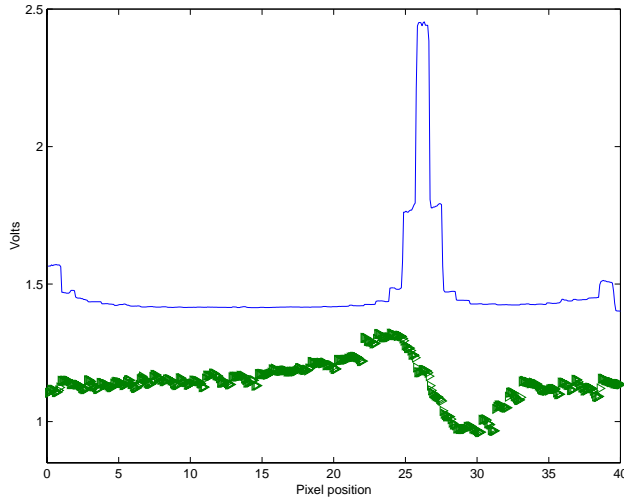

Figure 5-5. Response of the WTA network to the ON-edge of a bar moving from left to right at an on-chip speed of 31 mm/s. The top trace represents the currents I_{sum} of the WTA array while the bottom trace represents the voltage outputs of the array of adaptive photoreceptors.

Tracking robots

An application domain that is well suited for the visual tracking chip is that of vehicle-guidance and autonomous navigation. These types of tasks in fact require compact and power-efficient computing devices which should be robust to noise, tolerant to adverse conditions induced by the motion of the system (e.g. to jitter and camera calibration problems) and possibly able to adapt to the highly variable properties of the world. To test our tracking sensor within this framework, we successfully interfaced it several types of robotic platforms, ranging from *Koala* (K-Team, Switzerland) rovers to LEGO toys (see Fig. 5.6).
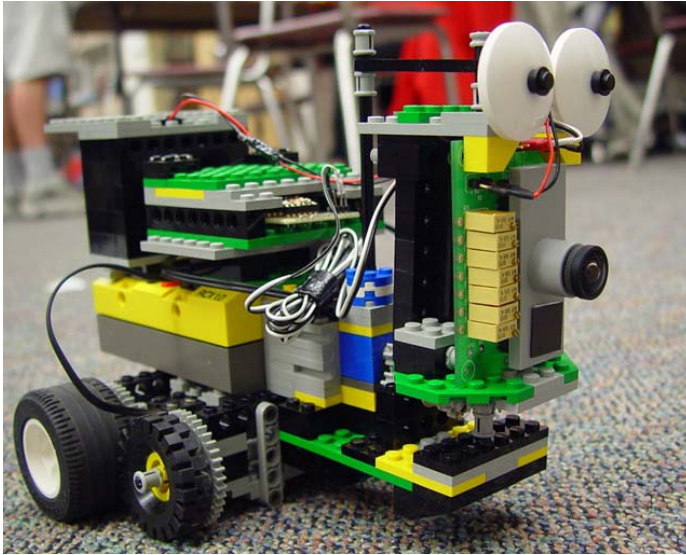
Figure 5-6. Tracker chip mounted on a LEGO robot performing a ``target exploration task''. Using very little CPU power, this robot is able to simultaneously *explore* (make random body/head movements), *attend* (orient the sensor toward high-contrast moving edges) and *pursuit* (drive towards the target).

In these applications the computationally expensive part of the processing (involving visual preprocessing and target selection) is done in real-time by the neuromorphic sensor. Using simple control algorithms, in conjunction with these types of sensors, roving robots are able to reliably track lines randomly layed out on the floor, for a wide variety of conditions (e.g. floors with different texture, cables of different colors and sizes, extreme illumination conditions, etc.) Several robot-tracking examples, together with a detailed analysis of the tracking system and quantitative data, have been published in the literature [11].

As the visual processing circuits operate in a fully parallel way, and the hysteretic WTA circuit relies on a global competition mechanism that requires one single node for the whole array, tracking architectures of the type described above can easily be extended to two dimensions [6,7,17].

5.3. MULTI-CHIP NEUROMORPHIC SYSTEMS

The single chip neuromorphic systems of the type describe above have great advantages, such as size, fabrication cost and low power consumption, and extraordinary computational capabilities. However, to design systems

with greater computational power and higher flexibility one needs to resort to *multi-chip* systems.

Neuromorphic multi-chip systems generally consist of systems containing one or more *sensory* devices, such as silicon retinas, silicon cochleas or vision sensors, interfaced to one or more chips containing networks of spiking neuron circuits. These chips can process the sensory signals (e.g. detecting salient regions of the sensory space [18], learning correlations [19], etc.) and eventually transmit the processed signals to actuators, thus implementing complete neuromorphic sensory-motor systems.

Consistent with the neuromorphic engineering approach, the strategy used by neuromorphic devices to communicate analog signals across chip boundaries is inspired from the nervous system. Analog signals are converted into streams of stereotyped non-clocked digital pulses (spikes) and encoded using pulse-frequency modulation (spike rates). These digital pulses are transmitted using an asynchronous communication protocol based on the *Address-Event Representation (AER)* [20-22].
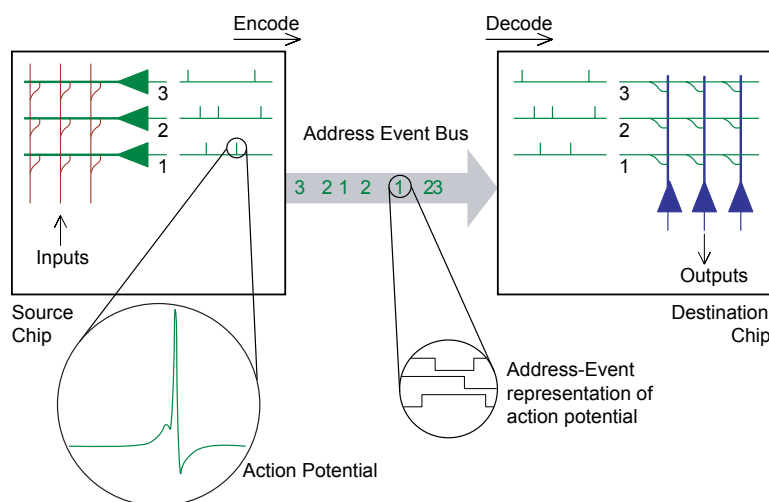
The Address-Event Representation

Figure 5-7. Schematic diagram of an AER chip to chip communication example. As soon as a sending node on the source chip generates an event its address is written on the Address-Event Bus. The destination chip decodes the address-events as they arrive and routes them to the corresponding receiving nodes.

In AER, each analog element on a sending device is assigned an address. When a spiking element generates a pulse its address is instantaneously put on a digital bus, using asynchronous logic (see Fig. 5.7). In this

asynchronous representation time represents itself, and analog signals are encoded by the inter-spike intervals between the addresses of their sending nodes. Address-events are the digital pulses written on the bus. In the case of single-sender/single-receiver communication, a simple handshaking mechanism ensures that all events generated at the sender side arrive at the receiver. The address of the sending element is conveyed as a parallel word of sufficient length, while the handshaking control signals require only two lines. Systems containing more than two AER chips are constructed by implementing additional special purpose off-chip arbitration schemes [22,23].

Address-Event Neuromorphic Sensors

The two most successful types of neuromorphic sensors developed in previous years are *silicon cochleas* [4,24] and *silicon retinas* [2,3,25]. The former implement detailed models of the human cochlea, producing outputs that could be useful for artificial speech recognizers, or for hearing aids. The silicon retinas on the other hand implement models of the retina's early processing stages and typically produce images that represent local changes in contrast (see Fig. 5.8 for an example of a silicon retina image).

Figure 5-8. Image captured from a 168×132 silicon designed by Jörg Kramer, (at the Institute of Neuroinformatics, Zurich), while the subject was moving.

Until recently these sensory devices transmitted their information off-chip using conventional techniques, such as multiplexers or scanners. With the advent of the Address-Event Representation we now have also AER silicon retinas and cochleas that produce streams of address-events representing the activity of each individual pixel. With these AER sensors the bandwidth used for signal transmission is allocated optimally only for those pixels that are active (as opposed for example to scanning techniques, that allocate the same bandwidth for all the pixels, independent of their activity). The

address-events (spikes) generated by these sensors can then be processed by synapses and networks of spiking neurons implemented on one or more receiving AER chips.

Address-Event Networks of Integrate-and-Fire Neurons

A spiking neuron model that allows us to implement large, massively parallel networks of neurons is the Integrate-and-Fire (I&F) model.

Networks of I&F neurons have been shown to exhibit a wide range of useful computational properties, including feature binding, segmentation, pattern recognition, onset detection, input prediction, etc. [26]. The recent and growing interest in pulse--based neural networks has lead to the design and fabrication of an increasing number of VLSI networks of Integrate—and—Fire (I&F) neurons. These types of devices have great potential, allowing researchers to implement simulations of large networks of spiking neurons with complex dynamics in real time, possibly solving computationally demanding tasks. This is especially true as continuous improvements in VLSI technology allow for the fabrication of AER devices containing thousands of elements, operating in parallel. For these devices to be be practically realizable, it is crucial to have pulse generating elements with minimal power consumption (locally) and with pulse-frequency saturation and adaptation mechanisms to limit and reduce the power consumption globally and to optimize communication bandwidth for the transmission of address-events. The I&F neuron circuit depicted in Fig. 5.9 implements these saturation and adaptation mechanisms, and has been shown to be one of the lowest-power circuits of its kind [27].
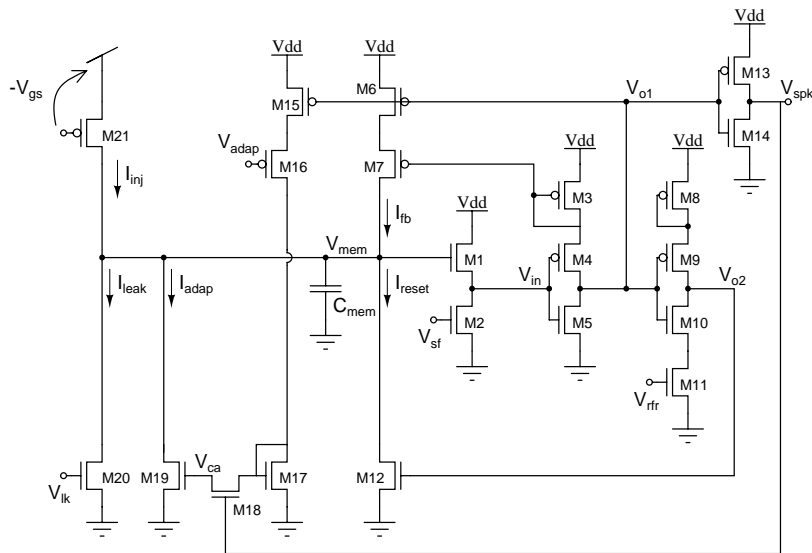
Figure 5-9. Circuit diagram of an ultra low-power integrate-and-fire neuron.

It comprises twenty transistors and one (explicit) capacitor. The circuit can be subdivided in six main blocks: a source follower M1-M2, for increasing the linear integration range and for modulating the neuron's threshold voltage; an inverter with positive feedback M3-M7, for reducing power dissipation due to the switching short-circuit currents at the input; an inverter with controllable slew-rate M8-M11, for setting arbitrary refractory periods; a digital inverter M13-M14, for generating the fast digital pulse that signals the occurrence of a spike; a transient current-mirror integrator M15-M19, for implementing a spike-frequency adaptation mechanism, and a minimum size transistor M20 for implementing a constant current leak.

Networks of I&F neurons consist of arrays of these types of neurons connected to synaptic circuits that generate currents with biologically plausible dynamics. The synaptic circuits used in neuromorphic devices can exhibit simple non-linear integration properties [20], short-term depression Properties [28,29], or plasticity/learning properties [19,30].

Researchers in the neuromorphic engineering community are starting just now to put all these components together to form AER neural network devices. In Fig. 5.10 we show the activity of one of these types of devices, containing a network of 32 I&F neurons and of 32×8 plastic synapses in response to constant currents being injected into each neuron. Each dot in Fig. 5.10 (a) represents an address-event. The address of the spiking neuron is on the ordinates, while time is on the abscissae. In Fig. 5.10 (b) we plotted the neurons' mean firing rates as a function of their input current (set by changing the V_{gs} of Fig. 5.9), for different refractory period settings (V_{rfr}

of Fig. 5.9), on a semi-logarithmic scale. Given the exponential relationship between V_{gs} and the current of a MOSFET working in weak-inversion [12], Fig. 5.10 (b) shows how the firing rate is linear with the input current, saturating at higher asymptotic values, for increasing values of V_{rfr} (decreasing refractory period duration).

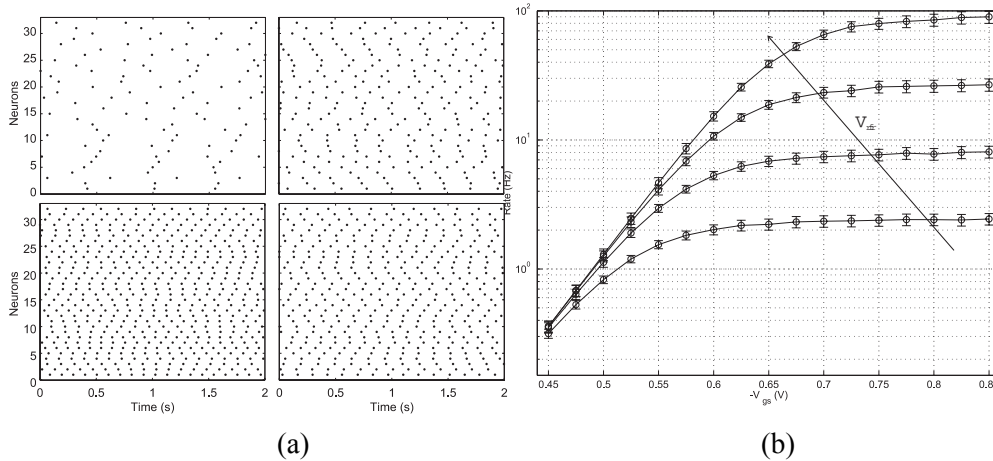(a)   (b)

Figure 5-10. (a) Raster plots showing the activity of an AER array of 32 I&F neurons in response to a constant input current, for four decreasing values of the refractory period (clockwise from the top left quadrant). (b) Mean response of all neurons in the array to increasing values of a global input current, for the same refractory period settings. The error bars represent the responses standard deviation throughout the array.

These AER networks of I&F neurons act as *transceivers*: they receive address-events in input and generate events in output. The topology of the network together with the weights of the synapses interconnecting the neurons determine the network's functionality. Address-event systems allow us to arbitrarily configure network topologies by re-mapping the (digital) address events (e.g.using lookup tables, micro controllers, or dedicated PCI boards [23]). There are currently different approaches for controlling synaptic weights. These include the use of Floating-Gate devices [31], binary synapses [15], or spike-timing based weight update rules [30, 32].

Multi-chip AER systems are still in a developmental stage, however all individual components have been designed, and the AER communication infrastructure is starting to be well established. This technology is likely to become instrumental both for basic research on computational neuroscience, and for practical applications that involve sensory signal processing, adaptation to changes in the input signals, recognition, etc.

5.4. CONCLUSIONS

During the past decade, complementary metal oxide semiconductor (CMOS) very large scale integration (VLSI) technology has been used to construct a wide range of neural analogs, from single synapses to networks of spiking neurons, and simple vision processing devices. These typically analog circuits exploit the inherent physics of transistors to produce an efficient computation of a particular task. The analog circuits have the advantage of emulating biological systems in real time. To the extent that the physics of the transistors matches well the computation to be performed, the analog VLSI circuits use less power and silicon area than would an equivalent digital system.   This is an important advantage because any serious attempt to replicate the computational power of brains must use resources as effectively as possible. The brain performs about 10^{16} operations per second.   Using the best digital technology that we can envisage, this performance would dissipate over 10 MW [33], by comparison with the brain's consumption of only a few Watts. Neuromorphic analog VLSI circuits are also no match for neuronal circuits, but they can be a factor of 10^4 more power efficient than heir digital counterparts.

The examples of neuromorphic systems we described in this chapter show how it is possible to build single-chip VLSI systems able to perform elaborate visual tasks, and multi-chip systems able to receive and process signals represented as spikes, or address-events. Our results with the single-chip example indicate that these types of systems can be used as efficient compact and low-cost solutions for real-world applications and can be considered as a viable alternative to conventional digital machine vision systems. Similarly our results with multi-chip AER systems indicate that those types of circuits can be reliably used in massively parallel VLSI networks of I&F neurons, to simulate in real--time complex spike-based learning algorithms, and possibly to implement complete sensory-motor systems that learn to interact with their environment and to adapt to its changing conditions.

ACKNOWLEDGMENTS

chip systems, and Rodney Douglas and Kevan Martin for providing an outstanding intellectual environment at the Institute of Neuroinformatics in Zurich, and for their support.

REFERENCES

1. R. Douglas, M. Mahowald, and C. Mead. Neuromorphic analogue VLSI. *Annu. Rev. Neurosci.*, 18:255–281, 1995.
2. M. Mahowald and C. Mead. *Analog VLSI and Neural Systems*, chapter Silicon Retina, pages 257–278. Addison-Wesley, Reading, MA, 1989.
3. K.A. Boahen and A.G. Andreou. A contrast sensitive silicon retina with reciprocal synapses. In D.S. Touretzky, M.C. Mozer, and M.E. Hasselmo, editors, *Advances in neural information processing systems*, volume 4. IEEE, MIT Press, 1992.
4. R. Sarpeshkar, R.F. Lyon, and C.A. Mead. An analog VLSI cochlea with new transconductance amplifiers and nonlinear gain control. In *Proc. IEEE Int. Symp. on Circuits and Systems*, volume 3, pages 292–296. IEEE, May 1996.
5. C. Koch and B. Mathur. Neuromorphic vision chips. *IEEE Spectrum*, 33 (5):38–46, May 1996.
6. V. Brajovic and T. Kanade. Computational sensor for visual tracking with attention. *IEEE Journal of Solid State Circuits*, 33(8):1199–1207, August 1998.
7. R. Etienne-Cummings, J. Van der Spiegel, and P. Mueller. A visual smoot pursuit tracking chip. In D. S. Touretzky, M. C. Mozer, and Hasselmo M. E., editors, *Advances in Neural Information Processing Systems*, volume 8. MIT Press, 1996.
8. T. G. Morris, T. K. Horiuchi, and S. P. DeWeerth. Object-based selection within an analog VLSI visual attention system. *IEEE Trans. on Circuits and Systems II*, 45(12):1564–1572, 1998.
9. G. Indiveri. A current-mode hysteretic winner-take-all network, with excitatory and inhibitory coupling. *Analog Integrated Circuits and Signal Processing*, 28(3):279–291, September 2001.
10. T. Delbrück and C.A. Mead. Analog VLSI phototransduction by continuous-time, adaptive, logarithmic photoreceptor circuits. In C. Koch and H. Li, editors, *Vision Chips: Implementing vision algorithms with analog VLSI circuits*, pages 139–161. IEEE Computer Society Press, 1995.
11. G. Indiveri. Neuromorphic analog VLSI sensor for visual tracking: Circuits and application examples. *IEEE Trans. on Circuits and Systems II*, 46(11):1337–1347, November 1999.
12. S.-C. Liu, J. Kramer, G. Indiveri, T. Delbrück, and R. Douglas. *Analog VLSI:Circuits and Principles*. MIT Press, 2002.
13. C. Tomazou, J. Lidgey, F., and G. Haigh, D., editors. *Analogue IC design: the current-mode approach*. Peter Peregrinus Ltd, 1990.
14. T. Serrano and B. Linares-Barranco. A modular current-mode high precision winner-take-all circuit. *IEEE Trans. on Circuits and Systems II*, 42(2):132–134, Feb 1995.
15. J. Lazzaro, S. Ryckebusch, M.A. Mahowald, and C.A. Mead. Winnertake-all networks of $O(n)$ complexity. In D.S. Touretzky, editor, *Advances in neural information processing systems*, volume 2, pages 703–711, San Mateo - CA, 1989. Morgan Kaufmann.
16. A. Starzyk, J. and X. Fang. CMOS current mode winner-take-all circuit with both excitatory and inhibitory feedback. *Electronic Letters*, 29(10):908–910, May 1993.

17. G. Indiveri, P. Oswald, and J. Kramer. An adaptive visual tracking sensor with a hysteretic winner-take-all network. In *Proc. IEEE International Symposium on Circuits and Systems*, pages 324–327. IEEE, May 2002.

18. G. Indiveri. A neuromorphic VLSI device for implementing 2-D selective attention systems. *IEEE Trans. on Neural Networks*, 12(6):1455–1463, November 2001.

19. E. Chicca, D. Badoni, V. Dante, M. D'Andreagiovanni, G. Salina, S. Fusi, and P. Del Giudice. A VLSI recurrent network of integrate–and–fire neurons connected by plastic synapses with long term memory. *IEEE Trans. Neural Net.*, 14(5):1297–1307, September 2003.

20. K.A. Boahen. Communicating neuronal ensembles between neuromorphic chips. In T. S. Lande, editor, *Neuromorphic Systems Engineering*, pages 229–259. Kluwer Academic, Norwell, MA, 1998.

21. J. Lazzaro, J. Wawrzynek, M. Mahowald, M. Sivilotti, and D. Gillespie. Silicon auditory processors as computer peripherals. *IEEE Trans. on Neural Networks*, 4:523–528, 1993.

22. S. R. Deiss, R. J. Douglas, and A. M. Whatley. A pulse-coded communications infrastructure for neuromorphic systems. In W. Maass and C. M.Bishop, editors, *Pulsed Neural Networks*, chapter 6, pages 157–178. MITPress, 1998.

23. V. Dante and P. Del Giudice. The PCI-AER interface board. In A. Cohen, R. Douglas, T. Horiuchi, G. Indiveri, C. Koch, T. Sejnowski, and S. Shamma, editors, *2001 Telluride Workshop on Neuromorphic Engineering Report*, pages 99–103, 2001. http://www.ini.unizh.ch/telluride/previous/report01.pdf.

24. E. Fragnière, A. van Schaik, and E. Vittoz. Design of an analogue VLSI model of an active cochlea. *Jour. of Analog Integrated Circuits and Signal Processing*, 13(1/2):19–35, May 1997.

25. J. Kramer. An integrated optical transient sensor. *IEEE Trans. on Circuits and Systems II*, 49(9):612–628, Sep 2002.

26. W. Maass and C. M. Bishop. *Pulsed Neural Networks*. MIT Press, 1998.

27. G. Indiveri. A low-power adaptive integrate-and-fire neuron circuit. In *Proc. IEEE International Symposium on Circuits and Systems*. IEEE, May 2003.

28. E. Chicca, G. Indiveri, and R.J. Douglas. An adaptive silicon synapse. In *Proc. IEEE International Symposium on Circuits and Systems*. IEEE, May 2003.

29. C. Rasche and R. Hahnloser. Silicon synaptic depression. *Biological Cybernetics*, 84(1):57–62, 2001.

30. G. Indiveri. Neuromorphic bistable VLSI synapses with spike-timing dependent plasticity. In *Advances in Neural Information Processing Systems*, volume 15, Cambridge, MA, December 2002. MIT Press.

31. C. Diorio, P. Hasler, B.A. Minch, and C. Mead. A single-transistor silicon synapse. *IEEE Trans. Electron Devices*, 43(11):1972–1980, 1996.

32. A. Bofill and A.F. Murray. Circuits for VLSI implementation of temporally asymmetric Hebbian learning. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information processing systems*, volume 14. MIT Press, Cambridge, MA, 2001.

33. C.A. Mead. Neuromorphic electronic systems. *Proceedings of the IEEE*, 78(10):1629–1636, 1990.