

A Neuromorphic VLSI Device for Implementing 2-D Selective Attention Systems

Giacomo Indiveri

Abstract—Selective attention is a mechanism used to sequentially select and process salient subregions of the input space, while suppressing inputs arriving from nonsalient regions. By processing small amounts of sensory information in a serial fashion, rather than attempting to process all the sensory data in parallel, this mechanism overcomes the problem of flooding limited processing capacity systems with sensory inputs. It is found in many biological systems and can be a useful engineering tool for developing artificial systems that need to process in real-time sensory data. In this paper we present a neuromorphic hardware model of a selective attention mechanism implemented on a very large scale integration (VLSI) chip, using analog circuits. The chip makes use of a spike-based representation for receiving input signals, transmitting output signals and for shifting the selection of the attended input stimulus over time. It can be interfaced to neuromorphic sensors and actuators, for implementing multichip selective attention systems. We describe the characteristics of the circuits used in the architecture and present experimental data measured from the system.

Index Terms—Active vision, analog very large scale integration (VLSI), neuromorphic, selective attention, winner-take-all.

I. INTRODUCTION

SELECTIVE attention is a mechanism used by a wide variety of biological systems to optimize their limited parallel-processing resources by identifying relevant subregions of the sensory input space and processing them in a serial fashion, shifting sequentially from one subregion to the other. This mechanism acts as a dynamic filter that allows the system to determine what information is relevant for the task at hand and to process it, while suppressing the irrelevant information that the system is not able to analyze simultaneously. It can be a very effective engineering tool for designing artificial systems that need to process in real-time sensory information and that have limited computational resources.

Biological selective attention mechanisms are believed to be modulated by stimulus-driven and goal-driven factors to facilitate the emergence of a “winner” from several potential targets [1]. While stimulus-driven selective attention appears to be a rapid bottom-up task-independent mechanism, goal-driven selective attention appears to act in a slower top-down volition controlled manner. The hardware architecture we present in this paper implements a real-time model of the stimulus-driven form of selective attention, based on the *saliency map* concept, originally put forth by Koch and Ullman [2]. Saliency

map-based models of selective attention account for many of the observed behaviors in neurophysiological and psycho-physical experiments [3] and have interesting computational properties that led to several software implementations, applied to machine vision and robotic tasks [4]–[7]. Similarly, several very large scale integration (VLSI) systems for implementing selective attention mechanisms have also been presented [8]–[11]. These VLSI systems though have all the common feature of implementing single-chip *visual* selective attention mechanisms: they all contain photo-sensing elements and processing elements on the same focal plane and typically apply the competitive selection process to *visual* stimuli sensed and processed by the focal plane processor itself.

Unlike these systems, the architecture described in this paper has been designed as the central part of a multichip system, able to receive input signals from many different types of sensory devices. Input signals need not arrive only from visual sensors, but could represent a wide variety of sensory stimuli obtained from different sources. In this framework, sensory signals are sent to (and from) the selective attention chip in the form of asynchronous binary pulses of fixed height, but with variable inter-pulse intervals (similar to neural spike trains), conforming to the address-event representation (AER) [12], [13]. Decoupling the sensing stage from the processing stage and using the AER to receive and transmit signals has several advantages: a multichip AER attention system could use multiple sensors to construct the input saliency map for the selective attention chip; the visual sensors used for generating the saliency map could be relatively high-resolution silicon retinas and would not have the small fill factors that single-chip two-dimensional (2-D) attention systems are troubled with; top-down modulating signals could be fused with the bottom-up generated saliency map to bias the selection process; multiple instances of the same selective attention chip could be used to construct hierarchical selective attention architectures.

In the following sections we describe the architecture of the selective attention chip and present experimental results that demonstrate the expected functionality of the chip and suggest possible applications.

II. THE SELECTIVE ATTENTION CHIP

The selective attention chip described in this paper was fabricated using a standard 2 μm CMOS technology. Its size is approximately 2 mm \times 2 mm and it contains of an array of 8 \times 8 cells. The chip’s architecture, easily expandable to arrays of arbitrary size, is laid out on a square grid, with input and output AER interfacing circuits.

Manuscript received November 29, 2000; revised May 14, 2001. This work was supported by the Swiss National Science Foundation SPP Grant.

The author is with the Institute of Neuroinformatics, University of Zurich and ETH Zurich, CH 8057 Zurich, Switzerland.

Publisher Item Identifier S 1045-9227(01)09507-8.

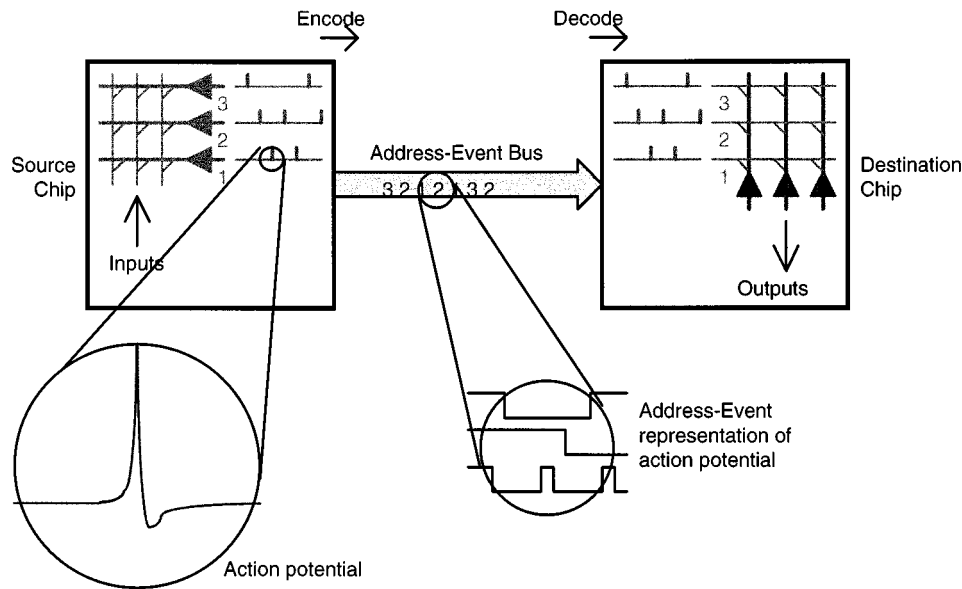


Fig. 1. Schematic diagram of an AER chip to chip communication example: “slow” action potentials generated by the “many” (N) neurons of the source chip are transmitted as “fast” address events, on the “few” ($\log N$) wires of the AE bus, to the synapses of the destination chip neurons.




Fig. 2. Block diagram of a basic cell of the 8×8 selective attention architecture.

A. The AER

The AER communication protocol allows the chip to exchange data while processing signals in parallel, in real time. In this protocol input and output signals are transmitted as asynchronous binary data streams which carry the analog information in their temporal structure (see Fig. 1). Each event is represented by a binary word encoding the address of the sending node. The address of the sending element is conveyed in parallel along with two handshaking control signals [12], [14]. Systems containing more than two AER chips can be constructed implementing additional special purpose off-chip arbitration schemes [13], [15]. These schemes employ lookup-tables (stored in EPROMs or computed by microcontrollers) to remap address-events. Specifically, address-events can be remapped from multiple sending nodes to a single one, thus generating the receiving node’s *receptive*

field, or from a single sending node to multiple receiving ones, thus generating a sending node’s *projective field*. In multichip systems, AER devices as the one proposed in this work can achieve peak throughputs of 2.5 MS/s (million spikes per second), with handshaking cycle times of approximately of 700 ns, if implemented on 64×64 arrays using a $2 \mu\text{m}$ CMOS technology [16].

B. The Chip’s Architecture

In a system containing AER sensors interfaced to the selective attention chip, address events would reach, at the input stage of each cell of the 8×8 array, excitatory synaptic (integrator) circuits that convert the digital voltage pulse streams into analog input currents. Fig. 2 shows the block diagram of one of the architecture’s cells. The input current integrated by the excitatory synapse (see I_{ex} in Fig. 2) is sourced into a neuromorphic

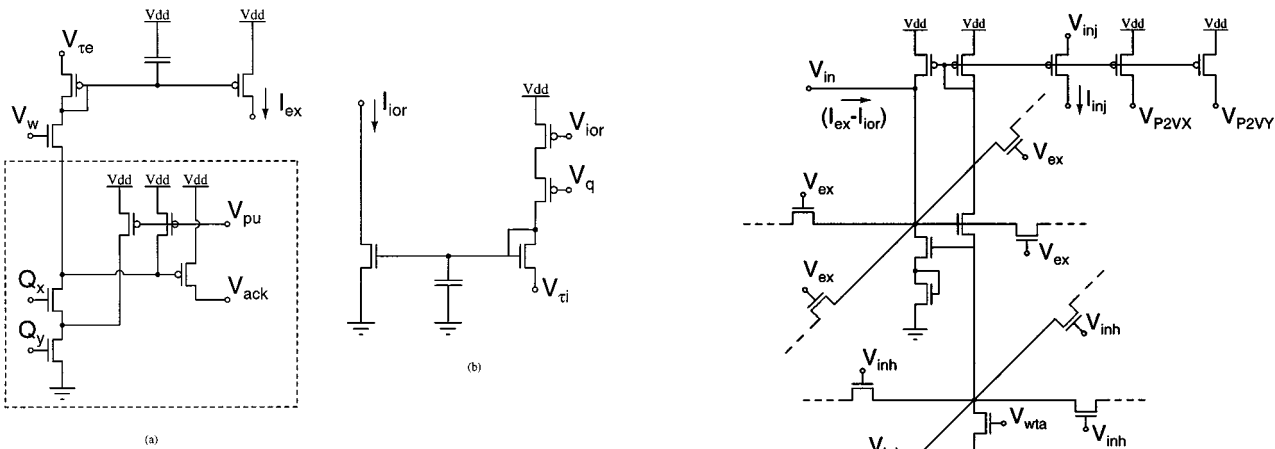


Fig. 3. Synaptic circuits. (a) Input excitatory synapse. Address events are converted into pulses by the circuit in the dashed box. Pulses are integrated into the excitatory current I_{ex} by the p-type current-mirror integrator. The integrator's gain and time constant are modulated by the control voltages V_w and V_{te} . (b) Inhibitory synapse. On-chip pulses (V_{ior}) are integrated into the inhibitory current I_{ior} by the n-type current-mirror integrator. The time constant and gain of this integrator are modulated by the voltages V_q and V_{τ_i} .

analog circuit that, connected with its neighbors, implements a hysteretic winner-take-all (WTA) network [17]. The output current of each WTA cell is used to activate both an integrate and fire (I&F) neuron and two position to voltage (P2V) circuits [18]. The P2V circuits encode the x and y coordinates of the winning WTA cell with two analog voltages, while the I&F neurons generate pulses that are used by the AER interfacing circuits to encode the position of the winning WTA cell with address-events. The neuron's spikes are also integrated by the local inhibitory synapse connected to it, to generate a current I_{ior} that is subtracted from the current I_{ex} (see Fig. 2). Fig. 3 shows the circuit diagram of both excitatory and inhibitory synapses. Both circuits use compact, nonlinear current-mirror integrators [16], [19] to integrate their input spikes. The transistors in the dashed box of Fig. 3(a) implement the AER input interfacing circuits and can operate correctly over a wide range of input pulse widths, ranging from a few hundred nanoseconds to milliseconds [16]. The gain and time constants of the two current-mirror integrators are set by two pairs of control voltages (V_w and V_{te} for the excitatory synapse and V_q and V_{τ_i} for the inhibitory synapse).

The sum of the currents ($I_{ex} - I_{ior}$) is sourced into the input node of the hysteretic WTA cell [node V_{in} in Fig. 4(a)]. Each cell is connected to its four nearest neighbors, both with lateral excitatory connections and lateral inhibitory connections [see Fig. 4(a)]. The inhibitory connections are modulated by the bias voltage V_{inh} and control the spatial extent over which competition takes place. If lateral inhibition is maximally turned on ($V_{inh} = V_{dd}$), all WTA cells of the architecture are connected together and only one winner can be selected at a time (global inhibition). If V_{inh} is low, the WTA network allows multiple winners to be selected, as long as they are sufficiently distant from each other (local inhibition). Similarly, lateral excitatory connections, modulated by the bias voltage V_{ex} , control the amount of lateral facilitatory coupling between cells. If lateral coupling is enabled, the system tends to select new winners in the immediate neighborhood of the currently selected cell.

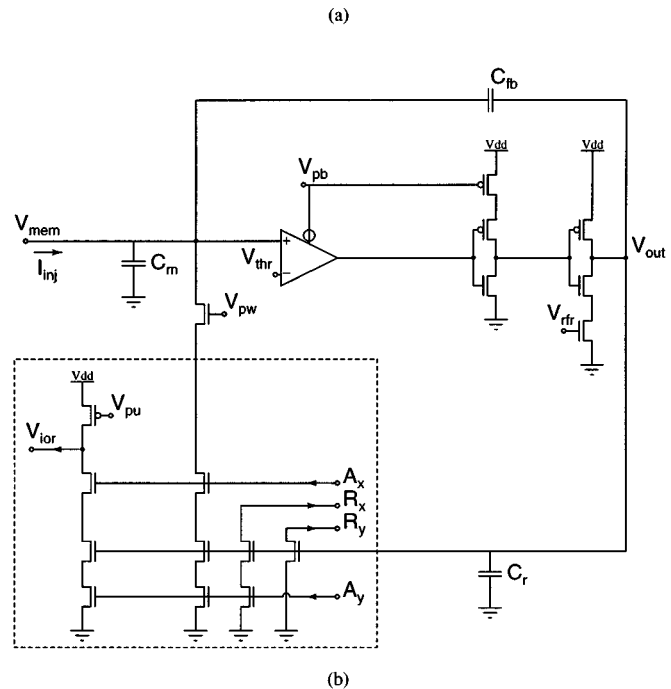


Fig. 4. (a) Hysteretic WTA cell. Input currents are sourced into node V_{in} and three copies of the output current are sent to the two P2V circuits and to the I&F neuron. (b) Local output integrate and fire neuron. When the membrane voltage V_{mem} increases above V_{thr} the output voltage V_{out} is driven to V_{dd} and an address event is generated. The transistors in the dashed box are part of the output AER circuitry.

When a WTA cell is selected as a winner, its output transistors source dc currents into the two P2V row and column circuits. The winning WTA cell also sources a dc current I_{inj} into the input node V_{mem} of the local inhibitory neuron connected to it [see Fig. 4(b)]. The amplitude of the injection current I_{inj} is independent of the input current ($I_{ex} - I_{ior}$), but depends on the bias voltage V_{wta} and on the control voltage V_{inj} . This current, integrated onto the neuron's capacitor C_m of Fig. 4(b), allows the neuron's membrane voltage V_{mem} to increase linearly with time. As soon as V_{mem} reaches the threshold voltage V_{thr} , the neuron generates an action potential: the comparator and the

inverters of Fig. 4(b) drives V_{out} to the positive power supply rail. This activates the AER row and column request signals (R_x and R_y), which produce an address event. The output AER circuit's acknowledge signals (A_x and A_y) reset the pulse by allowing the neuron's membrane capacitance to discharge at a rate controlled by V_{pw} .

Next to transmitting their address events off chip, the output neurons, together with the local inhibitory synapse connected to them, implement the inhibition of return (IOR) mechanism (a key feature of many selective attention systems) [20], [21]. The spikes generated by the winning cell's output neuron are integrated by its corresponding inhibitory synapse and gradually increase the cell's inhibitory postsynaptic current I_{ior} . As the neuron keeps on firing, the net input current to that cell ($I_{ex} - I_{ior}$) decreases until a different cell is eventually selected as the winner. When the previous winning cell is deselected its corresponding local output neuron stops firing and its inhibitory synapse recovers, decreasing the inhibitory current I_{ior} back to zero.

The IOR mechanism forces the WTA network to switch from selecting the cell receiving the strongest input to selecting cells receiving inputs of decreasing strength, effectively enabling the system to "attend" sequentially the salient regions of the input space. Depending on the dynamics of the IOR mechanism, the WTA network will continuously switch the selection of the winner between the two strongest inputs, or between the three strongest, or between all inputs above a certain threshold, generating focus of attention *scanpaths* similar to the ones observed for human eye movements [22]. The dynamics of the IOR mechanism depend on the time constants of the excitatory and inhibitory synapses (set by $V_{\tau e}$ and $V_{\tau i}$ of Fig. 3, respectively), on their synaptic strength (set by V_w and V_q respectively), on the frequency of the input stimuli and on the frequency of the output inhibitory neuron [set by V_{inj} of Fig. 4(a)].

As the single circuits described in this Section have been characterized in detail in previous publications [17], [19], we will present, in the next Section, experimental data demonstrating the chip's performance at the system level.

III. EXPERIMENTAL RESULTS

To characterize the behavior of selective attention chip with well controlled input signals we interfaced it to a workstation, via a National Lab-PC+ I/O card and stimulated it using the AER communication protocol. With this setup we were able to stimulate all the 64 pixels of the network with voltage pulses (i.e., address-events) at a maximal rate of 500 Hz. As the input synapses were set to have time constants of the order of milliseconds, each cell appeared to receive input spikes virtually in parallel. The handshaking between the chip and the PC was carried out at run time by the hardware present in the National I/O card. The chip's input stimuli consisted of patterns of address-events being generated by the workstation at uniform rates of different frequencies. We performed two sets of experiments, to demonstrate the chip's response properties using both the analog P2V outputs and the digital AER output.





Fig. 5. (a) Output of the P2V circuits of the selective attention architecture measured over a period of 300 ms, in response to a test stimulus exciting four corners of the input array at a rate of 30 Hz and a central cell at a rate of 50 Hz. (b) Histogram of the chip's output address-events, captured over a period of 13.42 s in response to the same input stimulus.

A. Analog P2V Outputs

In the first set of experiments, we used a test stimulus that excited cells (2,2) (2,7) (7,2) and (7,7) of the selective attention chip with 30 Hz pulses and cell (5,5) with 50 Hz pulses. Fig. 5(a) shows the analog output of the P2V circuits in response to 300 ms of stimulation with the input "saliency map" described above. The system initially selects the central cell (5,5). But, as the IOR mechanism forces the WTA network to switch the selection of the winner, the system cycles through all other excited cells as well. The P2V circuits are actively driven when the WTA network is selecting a winner (i.e., when the output p-type transistors of Fig. 4(a) are sourcing current into the nodes $P2VX$ and $P2VY$). At the times in which no cell is winning (i.e., when all cells are inhibited), there is no active device

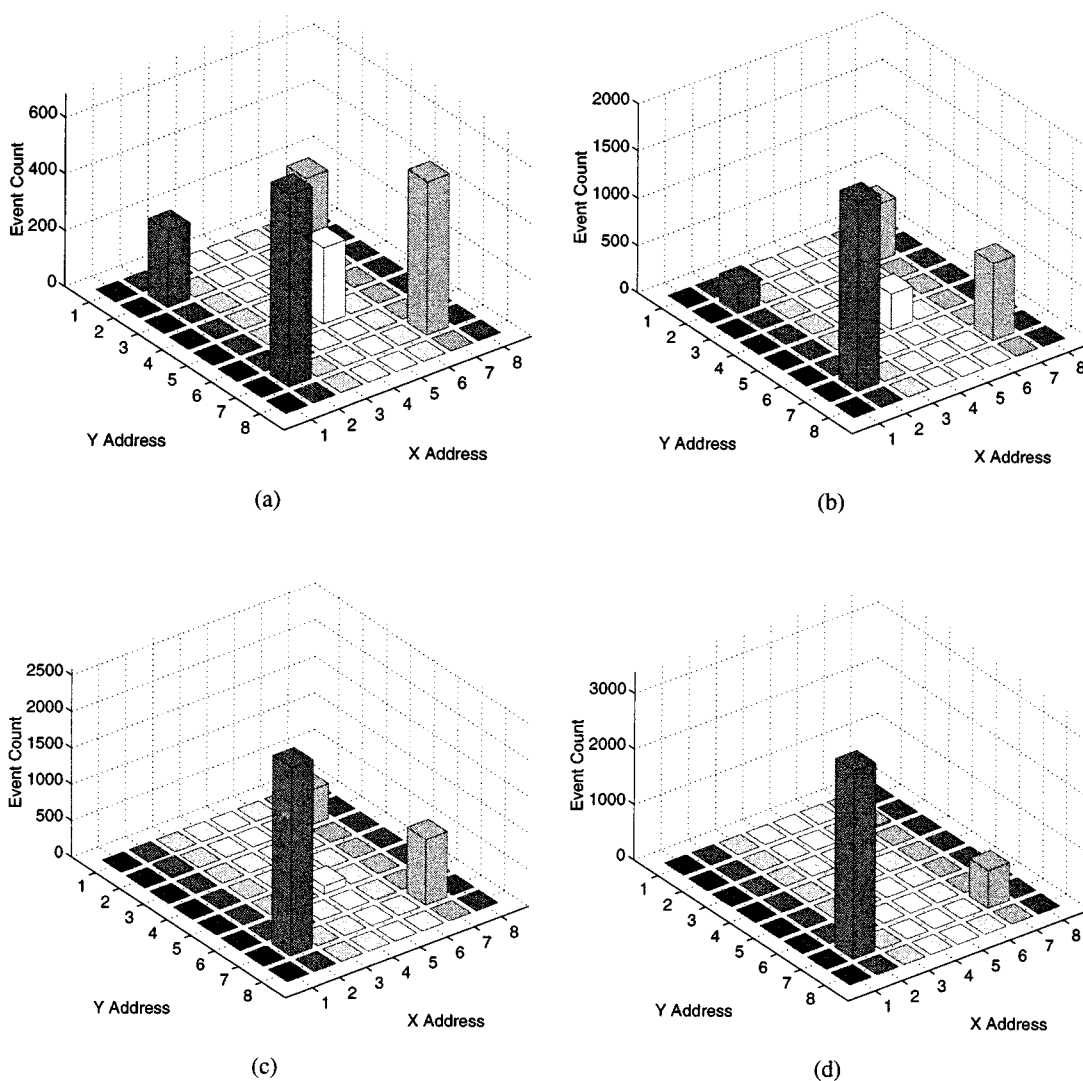


Fig. 6. Event histograms of addresses generated by the workstation sent to the chip (a) and output addresses generated by the selective attention chip (b), (c), and (d). All chip parameters are kept constant throughout the plots except for the bias parameter V_{τ_i} . The histogram in (b) was obtained with $V_{\tau_i} = 227$ mV, the one in (c) with $V_{\tau_i} = 207$ mV and the one in (d) with $V_{\tau_i} = 193$ mV.

driving the P2V circuits and their outputs tend to drift toward zero. This is evident in Fig. 5(a), for example, at the position corresponding to cell (7,2) in the lower right corner of the figure. When the network selects it as its eighth target, the horizontal P2V circuit outputs approximately 4.4 V and the vertical one outputs approximately 1.3 V. When the IOR mechanism forces the network to deselect the winner the outputs of the P2V circuits slowly drift toward zero. As soon as inhibition decreases, the network selects the cell (7,7) as the new (ninth) winner, the position to voltage circuits are actively driven again and their output quickly changes from approximately 3.6 V and 1.2 V to 4.2 V and 3.5 V (for the horizontal and vertical circuits, respectively).

B. Digital AER Outputs

To verify that the AER outputs are consistent with the analog P2V outputs, we stimulated the chip with the same pattern used for collecting the data of Fig. 5(a). We measured the address-events generated by the selective attention chip in response to

this input stimulus using a logic analyzer and plotted in Fig. 5(b) the histogram of such events. As shown, the chip's output address-events reflect, on average, the input stimulus and are consistent with the analog outputs of the P2V circuits.

The data of both Fig. 5(a) and (b) demonstrate how the IOR mechanism forces the network to switch the selection of the winner from one input to a different one, cycling through all sufficiently strong inputs. To demonstrate also how different IOR dynamics settings [modified for example by changing the bias voltage V_{τ_i} of Fig. 3(b)] affect the system's behavior, we performed a second experiment with a different input stimulus. The stimulation pattern used in this experiment excited cells (2,2), (5,5) and (7,2) with pulses at uniform frequency of 50 Hz, cell (7,7) with 100 Hz pulses and cell (2,7) with a 150 Hz pulses [see Fig. 6(a) for a histogram of the input address-events]. Fig. 6(b)–(d) show histograms of the chip's response for three different values of the bias voltage V_{τ_i} . The data of Fig. 6(b) was obtained by setting the time constant of the inhibitory synapse to a relatively high value ($V_{\tau_i} = 227$ mV). In this case once a cell is inhibited (after being selected as the winner), its input is sup-





Fig. 7. Output address events of the selective attention chip biased with $V_{\tau_i} = 207$ mV. The 2-D address space of the chip's architecture is mapped into the plot's one-dimensional ordinate vector by labeling each address successively, row by row.

pressed for an extensive period of time and the WTA network is forced to select all other (nonsuppressed) inputs. Conversely, the data of Fig. 6(d) was obtained by setting synapse time constant to a relatively low value ($V_{\tau_i} = 193$ mV). In this case the WTA network switches from selecting the cell receiving the strongest input to the cell receiving the second-strongest input and back. As the selected cells are not suppressed for sufficiently long periods of time, the remaining inputs never win the WTA competition. The histogram in Fig. 6(c) shows the data obtained for the intermediate case of $V_{\tau_i} = 207$ mV.


The same data used to compute the address-event histograms of Fig. 6 can be displayed using a different representation, to show the dynamics of the WTA competition stage. In Fig. 7, we plotted the address-events measured for the intermediate case of Fig. 6(c) over time. The addresses of the 8×8 cells are labeled successively row by row, such that labels 0 through 7 correspond to the addresses of the cells in the first row, labels 8 through 15 correspond to the addresses of cells in the second row and so on. Consistent with the histogram of Fig. 6(c), this plot shows how the system selects the cell (2,7) (labeled as 15 in Fig. 7) most frequently, switching occasionally to cell (5,5) (labeled as 37) and more often to cells (7,2) and (7,7) (labeled as 50 and 55). As mentioned in Section II, the details of the switching dynamics can be controlled by setting appropriately the bias voltages of the excitatory and inhibitory synaptic circuits [see (V_w, V_{τ_e}) and (V_q, V_{τ_i}) in Fig. 3] and the neuron's firing rate [controlled by V_{inj} of Fig. 4(b)]. These bias voltages, together with the other ones controlling the hysteretic WTA network's behavior [namely V_{wta} , V_{ex} and V_{inh} of Fig. 4(a)], endow the system with a sufficient amount of flexibility to be able to use the same chip in different types of selective attention tasks.

IV. SELECTIVE ATTENTION APPLICATIONS


The test stimuli used in the experiments of Section III were simple examples designed to demonstrate the expected behavior



(a)



(b)



(c)

Fig. 8. Image representations of saliency maps. (a) Saliency map corresponding to the input stimulus for the experiment of Fig. 5. (b) Saliency map used for the experiment of Fig. 6. (c) Fictitious example resembling a realistic saliency map.

of the selective attention chip. They do not resemble realistic saliency maps [see Fig. 8(a) and (b)]. In practical applications




Fig. 9. Selective attention active vision system. The selective attention chip processes sensory data coming from an AER imaging sensor and transmits its output to a workstation that drives the pan-tilt unit on which the sensor is mounted. A standard CCD camera is mounted next to the AER sensor to visualize the sensor's field of view.

saliency maps would more likely resemble the one shown in Fig. 8(c). More elaborate saliency maps could be processed by 2-D selective attention networks of greater size. The 8×8 architecture proposed in this paper can scale up to networks of arbitrary size: The performance of the hysteretic WTA circuits, which operate collectively in a massively parallel way, is not affected by the network's size. Similarly, given that in the selective attention system there is always one or a few winners at a time, the performance of the AER circuitry does not degrade with size (performance is affected only in architectures in which too many cells are trying to access the AER bus simultaneously).

We showed in a previous work [19], using a 32 pixel one-dimensional (1-D) version of the architecture presented in this paper, that these types of selective attention chips can operate reliably also on elaborate saliency maps, generated from high-resolution digitized images. In practical applications, the images could come for example from a camera connected to the workstation and the selective attention chip could be used to allocate in real-time CPU (image processing) resources only to the first n most salient regions of the image, or to scan the whole image in an intelligent way, sorting the scanning process by region saliency. Depending on the chip's bias settings, the system could also be tuned to visit each region only once, switching from region to the other slowly, or to revisit each region over and over again, switching from one region to the other quickly. Systems of this type would already benefit from the real-time response properties of the selective attention chip. But the most effective way of exploiting the computational properties of this chip would be to use it in conjunction with neuromorphic sensors that employ the AER communication protocol, such as silicon retinas or silicon cochleas [23]–[25]. These types of systems could be used as research tool for testing, in real-time, with real stimuli, different hypotheses on biological selective attention mechanisms [1], [26], [27]. Or they could be used

as low-cost alternatives to implement visual/auditory tracking or monitoring systems. For example, rather than using several fixed high-resolution (high-cost) cameras to monitor an environment, one could use a single motorized high-resolution camera driven by a selective attention system, comprising an AER silicon retina with a wide-field of view lens interfaced to the selective attention chip. An active vision system of this type, described in detail in [28], is shown in Fig. 9: the selective attention chip receives input from an AER imaging sensor [24] and transmits the address of the winning pixel to a workstation, that is used to drive the pan-tilt unit on which the sensor is mounted. A standard CCD camera is mounted next to the sensor, to visualize the sensor's field of view. The AER sensor responds to contrast transients and its address events report the position of moving objects. The selective attention chip selects the locations with highest contrast moving objects and cycles through them, while the workstation drives the pan-tilt unit centering the selected locations with the sensor's imaging array. An example of behavior of such a system, in response to real-world stimuli is shown in Fig. 10. The selective attention chip selects a moving target in the top-left part of the image and the pan-tilt unit is driven in a way to center the AER sensor's imaging array with the location selected by the attention chip.

V. CONCLUSION

We presented an analog VLSI device that implements a 2-D neuromorphic model of selective attention, for sequentially selecting the most salient locations of its inputs space. The device accepts input signals in the form of address-events and transmits its output using the same AER representation. It contains also analog position-to-voltage output circuits that can be used for quickly assessing the position of the selected input or for driving actuators, such as dc motors or pan-tilt units. We showed

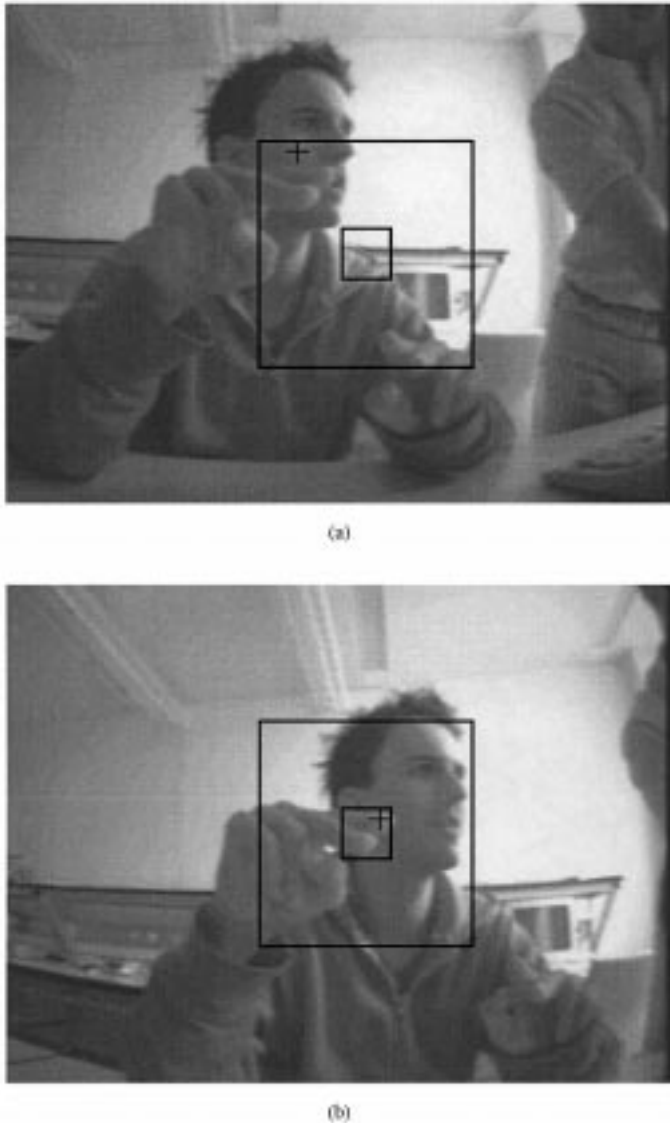


Fig. 10. Active vision system behavior, in response to moving stimuli. The image represents the CCD camera's output. The outer frame superimposed on the image represents the AER sensor's field of view. The inner frame represents the center of the sensor's imaging array and the cross represents the output of the selective attention chip. (a) Snapshot taken before the camera movement (the selective attention chip has just selected the moving stimulus in the top left part of the image). (b) Snapshot taken just after the camera movement (the salient stimulus is now centered in the AER sensor's imaging array).

experimental data that validates the functionality of the system, using control inputs generated on a workstation. We demonstrated how the chip's bias parameters can be used to impose different behaviors of the system and suggested possible applications.

The possibility to interface various types of AER sensory devices to the chip and to transmit the result of the selective attention competition to further processing stages using the same representation, allows the design of modular multichip selective attention systems. In these multichip systems, the properties of neuromorphic sensors, of the selective attention chip and of the AER ensure that the transduction of sensory signals, the selection of the most salient region and the implementation of the complex dynamics that characterize the selective attention

mechanism would all take place virtually in parallel and run in continuous and in real time. These features, together with the possibility to bias the system to exhibit different types of selective attention behaviors, provide obvious advantages both for scientific investigation of selective attention system properties and for engineering applications.

ACKNOWLEDGMENT

The author would like to thank R. Douglas and K. Martin for their support, the MOSIS foundation for fabricating the VLSI chip, and A. Whatley and G. Chang for their comments on the manuscript.

REFERENCES

- [1] E. Niebur and C. Koch, "Computational architectures for attention," in *The Attentive Brain*, R. Parasuraman, Ed. Cambridge, MA: MIT Press, 1998, pp. 163–186.
- [2] C. Koch and S. Ullman, "Shifts in selective visual-attention—toward the underlying neural circuitry," *Human Neurobiol.*, vol. 4, no. 4, pp. 219–227, 1985.
- [3] L. Itti and C. Koch, "Computational modeling of visual attention," *Nature Neurosci. Rev.*, vol. 2, pp. 194–204, 2001.
- [4] L. Itti, C. Koch, and E. Niebur, "A model of saliency-based visual attention for rapid scene analysis," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 20, pp. 1254–1259, Nov. 1998.
- [5] H. Bosch, R. Milanese, and A. Labbi, "Object segmentation by attention-induced oscillations," in *Proc. IEEE Int. Joint Conf. Neural Networks*, vol. 2, 1998, pp. 1167–1171.
- [6] P. E. Trahanias, S. Velissaris, and T. Garavelos, "Visual landmark extraction and recognition for autonomous robot navigation," in *Proc. IEEE Int. Conf. Intell. Robots Syst. IROS '97*, vol. 2, 1997, pp. 1036–1043.
- [7] S. Baluja and D. A. Pomerleau, "Expectation-based selective attention for the visual monitoring and control of a robot vehicle," *Robot. Autonomous Syst. J.*, vol. 22, pp. 329–344, 1997.
- [8] V. Brajovic and T. Kanade, "Computational sensor for visual tracking with attention," *IEEE J. Solid-State Circuits*, vol. 33, pp. 1199–1207, Aug. 1998.
- [9] T. K. Horiuchi and C. Koch, "Analog VLSI-based modeling of the primate oculomotor system," *Neural Comput.*, vol. 11, pp. 243–265, 1999.
- [10] G. Indiveri, "Neuromorphic analog VLSI sensor for visual tracking: Circuits and application examples," *IEEE Trans. Circuits Syst. II*, vol. 46, pp. 1337–1347, Nov. 1999.
- [11] T. G. Morris, T. K. Horiuchi, and S. P. DeWeerth, "Object-based selection within an analog VLSI visual attention system," *IEEE Trans. Circuits Syst. II*, vol. 45, pp. 1564–1572, Dec. 1998.
- [12] J. Lazzaro, J. Wawrzynnek, M. Mahowald, M. Sivilotti, and D. Gillespie, "Silicon auditory processors as computer peripherals," *IEEE Trans. Neural Networks*, vol. 4, pp. 523–528, 1993.
- [13] S. R. Deiss, R. J. Douglas, and A. M. Whatley, "A pulse-coded communications infrastructure for neuromorphic systems," in *Pulsed Neural Networks*, W. Maass and C. M. Bishop, Eds. Cambridge, MA: MIT Press, 1998, ch. 6, pp. 157–178.
- [14] K. A. Boahen, "Multiple pathways: Retinomorph chips that see quadruple images," in *Proc. 7th Int. Conf. Microelectron. Neural, Fuzzy, Bio-Inspired Syst.; Microneuro'99*, Los Alamitos, CA, Apr. 1999, pp. 12–20.
- [15] C. Higgins and C. Koch, "Multichip motion processing," in *Proc. Conf. Advanced Res. VLSI*, Atlanta, GA, 1998.
- [16] K. A. Boahen, "Communicating neuronal ensembles between neuromorphic chips," in *Neuromorphic Systems Engineering*, T. S. Lande, Ed. Norwell, MA: Kluwer, 1998, pp. 229–259.
- [17] G. Indiveri, "A current-mode analog hysteretic winner-take-all network, with excitatory and inhibitory coupling," *J. Analog Integrated Circuits Signal Processing*, 2001.
- [18] S. P. DeWeerth, "Analog VLSI circuits for stimulus localization and centroid computation," *Int. J. Comput. Vision*, vol. 8, no. 3, pp. 191–202, 1992.
- [19] G. Indiveri, "Modeling selective attention using a neuromorphic analog VLSI device," *Neural Comput.*, vol. 12, no. 12, pp. 2857–2880, Dec. 2000.

- [20] B. Gibson and H. Egeth, "Inhibition of return to object-based and environment-based locations," in *Percept. Psychophys.*, 1994, vol. 55, pp. 323–339.
- [21] Y. Tanaka and S. Shimojo, "Location vs feature: Reaction time reveals dissociation between two visual functions," *Vision Res.*, vol. 36, no. 14, pp. 2125–2140, July 1996.
- [22] A. L. Yarbus, *Eye Movements and Vision*. New York: Plenum, 1967.
- [23] K. A. Boahen, "A retinomorphic vision system," *IEEE Micro*, vol. 16, no. 5, pp. 30–39, Oct. 1996.
- [24] J. Kramer, An Integrated Optical Transient Sensor, 2001, submitted for publication.
- [25] E. Fragnière, A. van Schaik, and E. Vittoz, "Design of an analogue VLSI model of an active cochlea," *J. Analog Integrated Circuits Signal Processing*, vol. 13, no. 1/2, pp. 19–35, May 1997.
- [26] R. Desimone and J. Duncan, "Neural mechanisms of selective visual attention," *Annu. Rev. Neurosci.*, vol. 18, pp. 193–222, 1995.
- [27] M. C. Mozer and M. Sitton, "Computational modeling of spatial attention," in *Attention*, H. Pashler, Ed. East Sussex, U.K.: Psychology Press, 1998, pp. 341–395.
- [28] G. Indiveri, R. Mürer, and J. Kramer, "Active vision using an analog VLSI model of selective attention," *IEEE Trans. Circuits Syst. II*, vol. 48, pp. 492–500, May 2001.



Giacomo Indiveri received the Laurea degree in electrical engineering from the University of Genoa, Italy, in 1992 and won a postgraduate fellowship within the "National Research Program on Bioelectronic Technologies" from which he graduated (*cum laude*) in 1995.

From 1994 to 1996, he worked as a Postdoctoral Fellow in the Department of Biology at the California Institute of Technology, Pasadena, working on the design of analog VLSI subthreshold neuromorphic architectures for low-level visual tasks and motion detection. He is a Research Assistant at the Institute Neuroinformatics of the Swiss Federal Institute of Technology and the University of Zurich, Switzerland. His current research interests include the design and implementation of neuromorphic systems for modeling selective attention neural mechanisms and for exploring the computational properties of networks of silicon integrate and fire neurons. He is co-teacher of two classes on the analysis and design of analog VLSI Neuromorphic Systems at the Swiss Federal Institute of Zurich and co-organizer of the Workshop Neuromorphic Engineering, held annually in Telluride, CO.