PERGAMON

Neural Networks 14 (2001) 629–643

Neural Networks

www.elsevier.com/locate/neunet

2001 Special issue

Orientation-selective aVLSI spiking neurons

Shih-Chii Liu*, Jörg Kramer, Giacomo Indiveri, Tobias Delbrück, Thomas Burg, Rodney Douglas

Institute of Neuroinformatics, ETH/UNIZ, Winterthurerstrasse 190, CH-8057 Zurich, Switzerland

Received 25 October 2000; revised 21 March 2001; accepted 21 March 2001

Abstract

We describe a programmable multi-chip VLSI neuronal system that can be used for exploring spike-based information processing models. The system consists of a silicon retina, a PIC microcontroller, and a transceiver chip whose integrate-and-fire neurons are connected in a soft winner-take-all architecture. The circuit on this multi-neuron chip approximates a cortical microcircuit. The neurons can be configured for different computational properties by the virtual connections of a selected set of pixels on the silicon retina. The virtual wiring between the different chips is effected by an event-driven communication protocol that uses asynchronous digital pulses, similar to spikes in a neuronal system. We used the multi-chip spike-based system to synthesize orientation-tuned neurons using both a feedforward model and a feedback model. The performance of our analog hardware spiking model matched the experimental observations and digital simulations of continuous-valued neurons. The multi-chip VLSI system has advantages over computer neuronal models in that it is real-time, and the computational time does not scale with the size of the neuronal network.

1. Introduction

The sheer number of cortical neurons and the vast connectivity within the cortex are difficult to duplicate in either hardware or software. Simulations of a network consisting of thousands of neurons with a connectivity that is representative of cortical neurons can take minutes to hours on a fast Pentium, particularly if spiking behavior is simulated. The simulation time of the network increases as the size of the network increases. We have taken initial steps in mitigating the simulation time of neuronal networks by developing a multi-chip VLSI system that can support spike-based cortical processing models. The connectivity between neurons on different chips and between neurons on the same chip are reconfigurable. The receptive fields are effected by appropriate mapping of the spikes from source neurons to target neurons. A significant advantage of these hardware simulation systems is their real-time property; the simulation time of these systems does not increase with the size of the network.

In this work, we show how we synthesized orientation-tuned spiking neurons using the multi-chip system in Fig. 1. The virtual connection from a selected set of neurons on the retina to the target neurons on the multi-neuron transceiver chip is achieved with a PIC microcontroller and an asynchronous event-driven communication protocol. The circuit on this multi-neuron chip approximates a cortical microcircuit (Douglas & Martin, 1991).

We explored the different models that have been proposed for the generation of orientation tuning in neurons of the V1 cortical area. These cortical neurons receive inputs from the LGN neurons that are themselves not orientation-selective. The proposed models can be divided into two classes: feedforward models and feedback models. In the feedforward model, the orientation selectivity of a cortical neuron is conferred only by the spatial alignment of the LGN neurons that are presynaptic to the cortical neuron (Hubel & Wiesel, 1962). In the feedback model, a weak orientation bias provided by the LGN input is sharpened by the intracortical excitatory and/or inhibitory feedback (Ben-Yishai, Bar-Or & Sompolinsky, 1995; Douglas, Koch, Mahowald, Martin & Suarez, 1995; Ferster & Koch, 1987; Sillito, 1992; Somers, Nelson & Sur, 1995). In this work, we looked at both a feedforward model and a feedback model with recurrent inhibition.

1.1. Orientation-selective retinas

Although there have been quite a number of projects in building silicon retinas (Boahen, 1996, 1999; Boahen & Andreou, 1992; Delbrück & Mead, 1994; Liu & Boahen, 1996; Mahowald & Mead, 1991), instances of orientation-selective chips are less common. Most of these

* Corresponding author.

0893-6080/01/$ - see front matter
PII: S0893-6080(01)00054-5

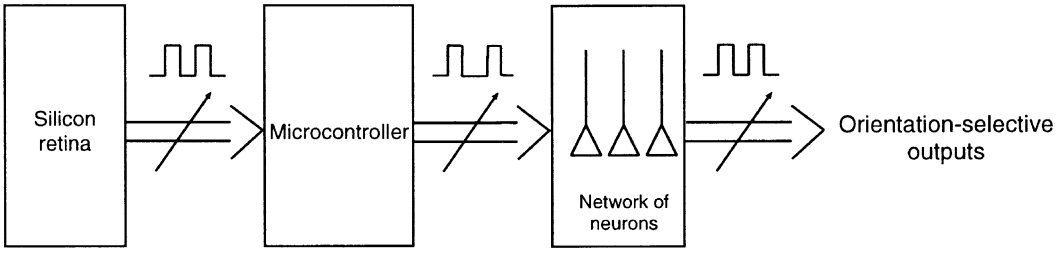

Fig. 1. Block diagram of a neuromorphic multi-chip system in which virtual connections from a set of neurons on a silicon retina onto another set of neurons on a transceiver chip are effected by a microcontroller. The neurons on the transceiver chip are interconnected in a winner-take-all architecture.

implementations are monolithic. Since it is difficult to have both a high fill factor and more sophisticated processing for a monolithic implementation, either the templates for the different orientations are placed to one side of the chip (Etienne-Cummings & Cai, 1997), or the circuitry is included in the pixel, thus decreasing the fill factor (Cauwenberghs & Waskiewicz, 1999; Raffo, Sabatini, Bo & Bisio, 1998; Shi, 2000).

There has been some work in developing neuromorphic multi-chip systems for orientation selectivity (Boahen, Andreou, Hinck, Kramer & Whatley, 1997; Venier, Mortara, Arreguit & Vittoz, 1997; Whatley, Kramer & Douglas, 1997). Projective field processors consisting of a silicon retina, microcontrollers, and multiple receivers have been explored in different projects at the NSF Telluride Workshop (Boahen et al., 1997). The pixels in these receivers consist of a circuit that integrates the incoming spikes on a capacitor and they are not spatially coupled together. The voltage at each pixel is then scanned out for display on a multi-sync monitor. Projective field maps were created by the mapping of a source neuron (sender) to several target pixels on a receiver. Each receiver was sensitive to only one orientation.

The transceiver in the multi-chip system by Venier et al. (1997) consists of pixels that are coupled together through a resistive network. Incoming spikes are integrated at each pixel and orientation selectivity in the network is achieved by setting the resistances so that smoothing is anistropic. All pixels in the transceiver have the same orientation preference. The outputs of the pixels were observed using a non-arbitered event-driven protocol.

A preliminary experiment in creating orientation-selective neurons was performed using a silicon retina and the silicon cortex board (SCX) (Whatley et al., 1997). The DSP on the SCX generates the receptive fields by using a lookup table for the incoming retina addresses. The spike outputs of the orientation-selective neurons can be observed through the address on the output data bus. In our system, the receptive fields of the orientation-selective neurons are created in a manner similar to that shown by Whatley et al. (1997). However, we extend this work and quantify the tuning curves of these neurons to different neuron parameters and explore the selectivity of the neurons in a feedforward model and a feedback model with a winner-take-all circuit. We compare the orientation tuning results with known results from the two models.

2. Connectivity in a neuromorphic multi-chip system

To achieve programmability in the interchip and intrachip wiring of neurons, we take advantage of the fast silicon substrate and replace point-to-point connections with a common digital bus and additional circuitry to handle the communication interface between multiple chips. The multiplexing of addresses from different neurons on a common bus is used in various neuromorphic multi-chip systems (Boahen, 2000; Deiss, Douglas & Whatley, 1999; Higgins & Koch, 1999; Indiveri, Whatley & Kramer, 1999; Lazzaro, Wawrzynek, Mahowald, Sivilotti & Gillespie, 1993; Mahowald, 1994).

This time-division multiplexing of addresses is possible because of the large difference in the time constant (or integration time) of the neurons (in milliseconds) and the bus (in microseconds). The communication protocol used for interchip communication, the address event representation (AER) protocol, was first introduced by Sivilotti (1991) and Mahowald (1994). The spike output of each pixel on the sending array (the sender) is encoded with a unique address. A common bus is used for transmitting the address to a receiver, where an address decoder drives the corresponding pixel in the receiver array.

There are different methods for transmitting the addresses of active pixels from the chip onto the common bus. The most common method is to use an on-chip arbiter that decides the order in which active pixels communicate their address. The addresses are placed in a queue and are transmitted off chip as fast as bus occupancy permits (Boahen, 2000). An alternative non-arbitered method (Mortara, Venier & Vittoz, 1995) allows any active pixel to place its address immediately on the common bus. The pixel addresses are coded in such a way that collisions result in invalid addresses that are ignored by the receiver. A comparison of the performances of these different arbitration schemes has been performed by Boahen (2000). Another arbitration scheme is equivalent to the CSMA (carrier sense/multiple access) protocol used in computer networking (Abusland, Lande & Hvin, 1996). In this

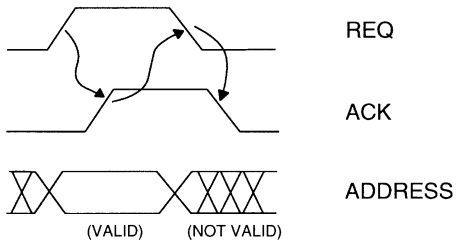

Fig. 2. Timing diagram for communication channel signals between a sender and a receiver. This protocol is known as a four-phase handshake. When the REQ signal from the sender is active, it means that the address on a common digital bus is valid. When the receiver has finished latching in this address, it raises its ACK signal to the sender. The sender then removes its REQ signal and in turn, the receiver removes its ACK signal, thus completing the four-phase handshake. Once ACK goes low, the sender is free to send another request to the receiver.

method, whenever a pixel makes a request, the arbiter checks to see if the bus is busy. If the bus is being used, the request from the pixel is either discarded or it is placed in a queue that is only one event wide. In this way, the time of the event is not arbitrarily delayed by the service time of the arbiter. This scheme is a variant of the arbitered method, in which the queue is only one event deep.

2.1. Arbitered address event representation protocol

We present a summary of the arbitered AER protocol for a point-to-point connection. Every neuron within a system has a unique digital address. These neurons may be distributed across multiple chips. If one of the neurons on a chip is active, that neuron sends a request (PIXEL REQ) to an on-chip arbiter. The on-chip arbiter decides which of the active neurons should place its address on the common bus (see Figs. 2 and 3) and acknowledges that neuron by raising its PIXEL ACK signal while placing that neuron's address on the bus. At the same time, the arbiter issues a chip request (CHIP REQ) to the corresponding receiver. The communication protocol between the sender and the receiver follows a similar protocol to that between the neurons and the on-chip arbiter (see Fig. 2). Meanwhile, the selected neuron removes its PIXEL REQ signal, but the PIXEL ACK signal will not be removed by the arbiter until the sender has received a CHIP ACK signal from the receiver. On receiving this signal, the sender removes its CHIP REQ signal and the address on the bus is now invalid. At the same time, the arbiter takes away the PIXEL ACK signal and selects another active neuron. The sender does not generate a new request to the receiver until the receiver has removed its CHIP ACK signal. This four-phase handshaking protocol applies in this multi-chip system between the retina and the PIC and between the PIC and the transceiver chip.

2.2. Reconfigurable connections using the address event representation protocol

An intermediate chip/system is used to program the virtual connections between different senders and receivers. This intermediate module can be an EPROM (Higgins & Koch, 1999; Higgins, Deutschmann & Koch, 1999), EEPROMs and PIC microcontrollers (Boahen et al., 1997; Cohen, Edwards, Stanford, Cauwenberghs, Andreou & Abshire, 1998), a silicon cortex (SCX) board (Deiss et al., 1999; Whatley et al., 1997), or the receiver itself (Serrano-Gotarredona, Andreou & Linares-Barranco, 1999). These modules are used to create projective fields from the sender neurons. The device that reconfigures the connections should not distort the input spike distribution. Here, we use a PIC to form the connections from the retina to the

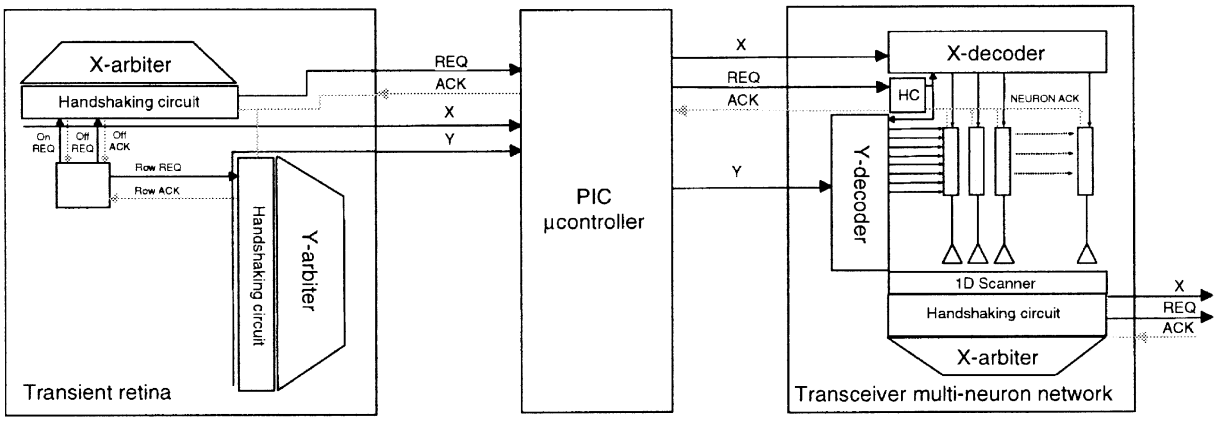

Fig. 3. Detailed block diagram of the three chips and the communication channel signals between them. The active neuron on the retina sends a Row REQ signal to the *Y*-arbiter. (Each pixel in the retina has two neurons which code the ON and OFF responses of the pixel.) If this row is selected, then all the ON and OFF REQ signals from the active neurons on this row will be arbitrated by the *X*-arbiter. The *X* and *Y* addresses of the selected neuron are placed on the bus and the retina handshakes with the PIC. The handshaking circuits on the retina and the transceiver chips (marked by the block HC on the multi-neuron chip) ensure that the four-phase handshaking protocol is not violated. Details of the handshaking circuits and protocol are described in Boahen (2000). The PIC communicates with the multi-neuron chip if the address from the retina falls within one of the stored templates. The address from the PIC is decoded by the multi-neuron transceiver. The *X* address codes the target neuron on this multi-neuron array and the *Y* address codes the synapse that will be stimulated on the selected neuron. The address of the active neuron on this array can also be communicated off-chip to another receiver/transceiver.

multi-neuron array. We will show that the PIC does not distort the interspike stimulus interval (ISI) distribution of spikes from the retina.

3. System architecture

The multi-chip system (Fig. 3) in this work consists of a 16×16 silicon ON/OFF retina, a PIC microcontroller, and a transceiver chip with a ring of 16 integrate-and-fire neurons and a global inhibitory neuron. All three modules communicate using the address event protocol. The communication channel signals consist of the address bits, the REQ signal, and the ACK signal. The PIC and the multi-neuron chip are both transceivers: they can both receive events and send events. The retina with an on-chip arbiter can only send events. Each pixel on the retina responds to both the ON edges and OFF edges of a stimulus. It is not spatially coupled to its neighbors. The irradiance temporal derivatives go to two neurons within the pixel which code separately the positive derivative (the ON response) and the negative derivative (the OFF response). The operation of the retina is described in Section 4. The spike outputs from the retina approximate the ON/OFF outputs of the lateral geniculate nucleus (LGN) neurons. An active neuron starts a four-phase handshaking with the on-chip arbiter as shown in Fig. 3. If the neuron is selected by the arbiter, then the *X* and *Y* addresses which code the column and row locations of this neuron are placed on the output address bus of the chip. The retina then handshakes with the PIC microcontroller. The description of the PIC microcontroller is presented in Section 6.

The multi-neuron chip has an on-chip address decoder for the incoming events and an on-chip arbiter to send events. The *X* address to the multi-neuron chip codes the identity of the neuron and the *Y* address codes the input synapse used to stimulate the neuron. Each neuron can be stimulated externally through any one of eight synapses; six of these are excitatory, and the remaining two are inhibitory. The excitatory neurons of this array are mutually connected via hardwired excitatory synapses. These excitatory neurons also excite a global inhibitory neuron which in turn inhibits all the excitatory neurons. The membrane potentials of the multi-neuron chip can be monitored by an on-chip scanner and the spiking outputs of the neurons can be monitored by the chip's AER output. The address on the output bus codes the active neuron. The architecture and circuit details of this chip are described in Section 7.

The receptive fields of the neurons are created by configuring the connections from a subset of the source pixels on the retina onto the appropriate target neurons on the multi-neuron transceiver chip through the PIC microcontroller. The subsets of retina pixels are determined by user-supplied templates. In this work, the excitatory neurons on the multi-neuron chip model the orientation tuning properties of simple cells in the visual cortex and the global inhibitory neuron models an inhibitory interneuron in the visual cortex. The measured tuning curves of the orientation-selective neurons in both the feedforwad and feedback cases using this multi-chip system are presented in Section 8.

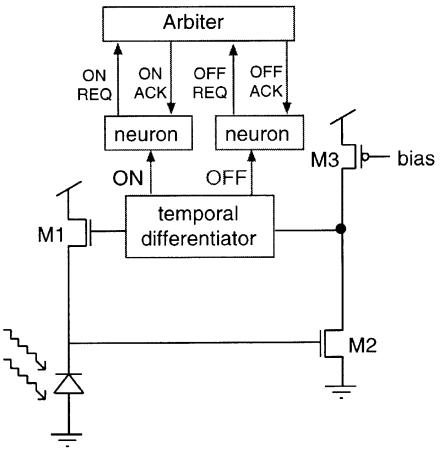

Fig. 4. Pixel of the transient imager. The circuit contains a photodiode with a transistor M1 in a source–follower configuration with a high-gain inverting amplifier (M2, M3) in a negative feedback loop. A rectifying temporal differentiator in the feedback loop extracts transient ON and OFF signals. These signals go to individual neurons that generate the REQ signals to the arbiter. In this schematic, we only show the REQ and ACK signals to the *X*-arbiter. The duration of the ACK signal from the *X*-arbiter is extended within the pixel by a global refractory bias. This duration sets the refractory period of the neuron.

4. The transient retina

The retina is a 16×16 array of pixels that detect transients in irradiances. The retina generates the events that drive the system. Each pixel responds in continuous time to a local change of a brightness distribution projected through a lens onto the surface of the retina. The output of the pixels is coded in the form of asynchronous binary pulses, which are the request signals to the AER communication interface.

The transient detector is composed of an adaptive photoreceptor (Delbrück & Mead, 1994) that has a rectifying temporal differentiator (Kramer, Sarpeshkar & Koch, 1997) in its feedback loop (shown in Fig. 4). Positive temporal irradiance transients (dark-to-bright or ON transitions) and negative irradiance transients (bright-to-dark or OFF transitions) appear at two different outputs of each pixel. The ON and OFF outputs are separately amplified with tunable gains and they drive two different neurons within the pixel. Each neuron generates a request signal to the on-chip arbiter if its input exceeds a chosen threshold. The ON and OFF outputs have different AER addresses. The ACK pulse from the *X*-arbiter triggers a reset pulse to the requesting neuron. The duration of this reset signal is controlled by a global refractory bias and determines the refractory period for the succeeding request from the same channel. Depending on the refractory period of the neuron

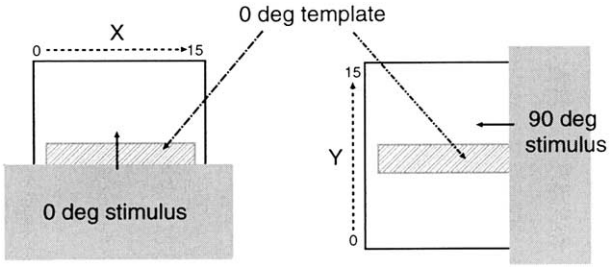

Fig. 5. Outline of the configuration of the orientation-selective experiments. A stimulus (edge) moves over the retina. The retina comprising 16×16 pixels is sensitive to bright-to-dark and dark-to-bright irradiance transitions which are coded on the ON channel and the OFF channel, respectively, of each pixel. The orientations of the stimuli are shown here. The lighter hatched area indicates the template of pixels used for detecting 0° oriented stimuli (or the horizontal template). Spikes from the set of neurons within this area are routed to a neuron on the transceiver chip that is sensitive to 0° orientation.

and the magnitude and duration of the irradiance transient, the pixel responds with either a single spike (the non-bursting mode) or a burst of spikes (the bursting mode).

The pixels are arranged on a rectangular grid. The position of a pixel along a row is encoded with a 4-bit column address (X address) and its position along a column with a 4-bit row address (Y address) as shown in Fig. 5. By setting the threshold and the respective gain factors appropriately, the circuit can be made to respond only to ON transients or only to OFF transients or to both types of transients. The circuit details are presented in Kramer (2001).

5. Response of retina to oriented stimuli

A rotating drum with a single black and white strip was placed about 6 cm from the retina. The lens on the retina had a focal length of 8 mm with a field of view of approximately 9.5°. The spike addresses and spike times generated by the retina at an image speed of 7.9 mm/s (or 89 pixels/s) in response to the rotating drum were recorded using an Agilent 16702A logic analyzer. The pixel pitch was 88.8 μm and the length of the array was 1332 μm. The orientations of the stimuli are defined in Fig. 5. An example of the spike addresses derived from both the ON and OFF channels of the retina in response to the rotating stimulus oriented at 90° is depicted in Fig. 6. In this experiment, the retina was set in the non-bursting mode, so only one spike is produced by each pixel that detects the stimulus. The retina address on the ordinate is defined as $16Y + X$. The average time of travel between the ON and OFF edges of the stimulus is about 1 s. The spike addresses during the time of travel of the OFF edge of a 0° oriented stimulus through the entire array (Fig. 7(a)) indicate that almost all the pixels along a row transmit their addresses sequentially as the edge passes by. This sequential ordering can be seen because the stimulus is oriented slightly different from 0°. If the stimulus was perfectly at 0°, then any of the pixels of one row could

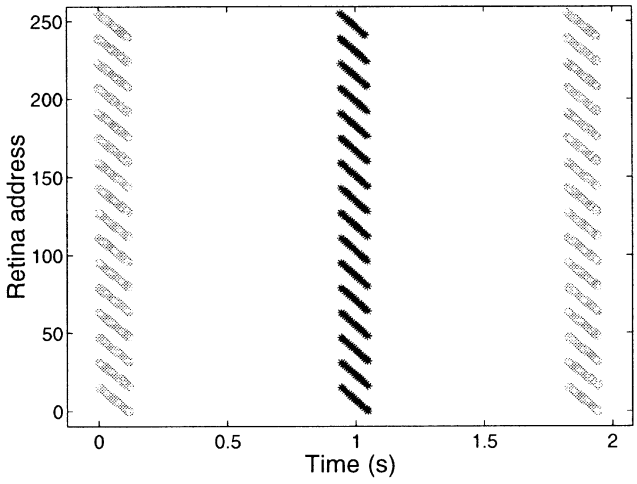

Fig. 6. ON-transient and OFF-transient spikes recorded from the 16×16 retina in response to an edge (oriented at 90°) moving over the array. The retina was set in the non-bursting mode; so each pixel should spike only once every time it sees the stimulus edge pass by. Each circle (OFF-transient) or asterisk (ON-transient) in the figure marks an address issued by the retina and recorded by the logic analyzer. The ON and OFF spikes from the pixels on the retina are coded by unique X and Y addresses. The address on the ordinate is defined as $16Y + X$. The Y address and X address codes the row and column location, respectively, of the active pixel. The polarity of the response (ON/OFF) is coded by an extra bit of the X address.

transmit their address and there would be no ordering of the pixel addresses. The same observation can be made for the OFF-transient spikes recorded in response to a 90° oriented stimulus (Fig. 7(b)). Each column of the retina fires sequentially in response to the stimulus.

The interspike interval (ISI) distributions of the retina spikes collected in the experiment above are shown in Fig. 8(a) for the 0° orientation and Fig. 8(b) for the 90° orientation. In this non-bursting mode, the arbiter on the retina can easily cope with the rate of incoming requests from the active pixels because the pixels spike only once per stimulus presentation, and only a small percentage of the pixels spike at the same time. In the worst case, all 16 points on one row or column will simultaneously make a request to the arbiter. Under these conditions, the ISI distribution reflects a jitter in the latency of the activated pixel that arises from the mismatch in the circuitry among the pixels in the array; and the tilt in the orientation of the stimulus. The arbiter itself does not introduce any distortion in the ISI distribution of the retina spikes.

When the retina is in the bursting mode, each pixel produces several spikes in response to a stimulus. The intraburst firing rate of the pixel is tunable via the refractory period bias. The ISI distributions of the retina spikes in response to two orientations of the stimulus are shown in Fig. 9. We can test the ability of the arbiter to service the increase in the incoming requests from the pixels by increasing the intraburst firing rate of the pixel until the minimum time interval between subsequent CHIP REQ signals from the arbiter stops decreasing indicating saturation of the

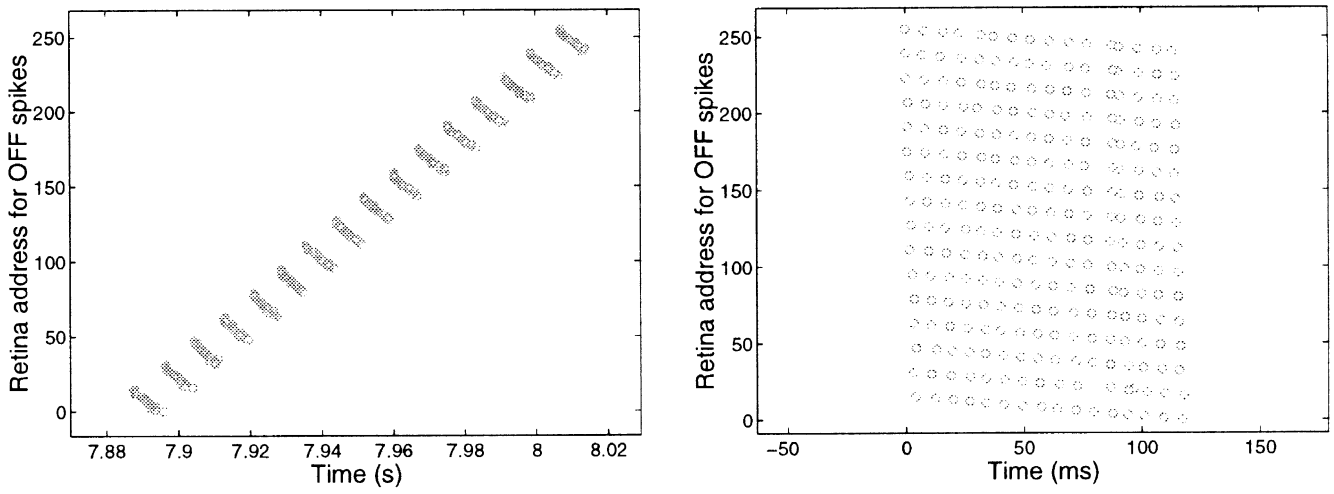

Fig. 7. The addresses of the spikes from the retina in the non-bursting mode are recorded for a (a) 0° and (b) 90° oriented stimulus, respectively. The figure shows the time progression of the stimulated pixels (OFF spikes are marked with circles) as the 0° oriented stimulus (see Fig. 5 for the orientation definition) passes over each row in (a). The address on the ordinate is defined as $16Y + X$. A similar observation is true of (b) for the ordering of the OFF-transient spikes when each column on the retina is stimulated by the 90° oriented stimulus.

channel capacity. The data in Fig. 9 show this limiting scenario where a large percentage of spikes have an ISI of around 2–4 μs. In this work, the arbiter has been tuned so it can only handle incoming requests with a minimum latency of 2 μs. The arbiter can be made to process events with a frequency of up to 10 MHz, but we have slowed down the arbiter because this implementation generated a few spurious addresses at faster speeds. At this latency, the AER interface can support an array of 10^5 neurons in which only 10% of the neurons are active at any one time, and the active neurons have an average spiking rate of 100 Hz.

5.1. Orientation-selective templates

Before using the PIC to create the receptive fields of the neurons on the multi-neuron transceiver chip, we further analyzed the recorded spikes from the experiment shown in Fig. 8 by binning the OFF spikes from selected sets of pixels on the retina. These sets were defined by two templates of size 5×9 and 9×5 which were centered in the middle of the retina. The templates are analogous to the receptive fields of the neurons. The PIC uses these templates to map the input spikes from the retina to two target neurons on the transceiver chip.

The spike histograms for the 0° oriented stimulus and for the 90° oriented stimulus using the two templates are shown in Fig. 10(a) and (b), respectively. The solid curve corresponds to the spikes collected using the horizontal template (9×5) and the dotted curve corresponds to the spikes collected using the vertical template (5×9). The histograms show the possibility of using a threshold to distinguish the orientation of the stimulus. Similar results were obtained from the retina in the bursting mode. This analysis shows that a feedforward model as

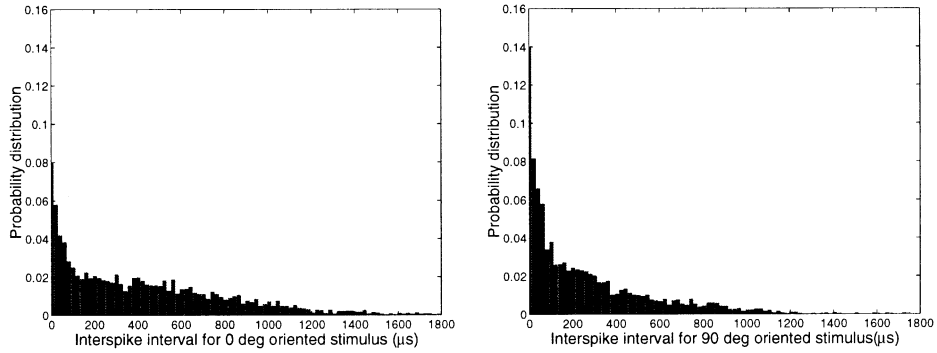

(a) (b)

Fig. 8. Probability distribution of the ISIs of the spike outputs from the retina in the non-bursting mode in response to a moving edge. The plot has been cropped for ISIs above 1800 μs so that we can see the detail in the smaller ISI distributions. (a) ISI distribution of the retina spikes in response to a moving edge which is oriented at 0°. The pixels on each row are stimulated close together in time as the stimulus moves over the row. (b) ISI distribution of the retina spikes in response to a moving edge oriented at 90°.

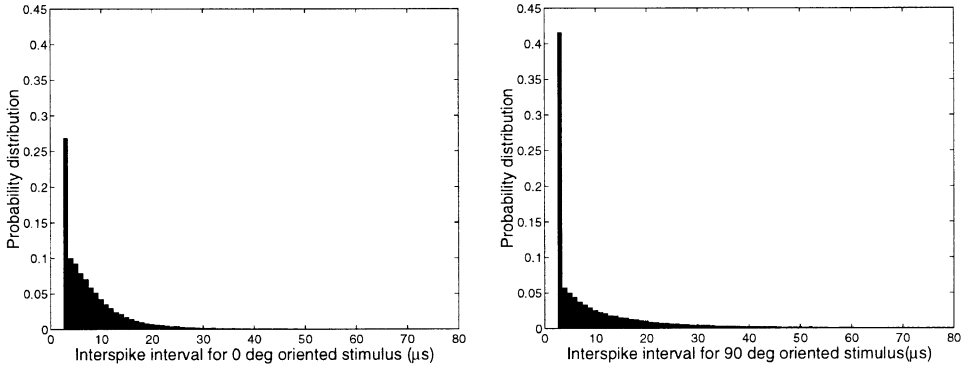

(a) (b)

Fig. 9. Probability distribution of the ISIs of the spike outputs from the retina in the bursting mode in response to a moving edge. The intraburst firing rate of the pixel is larger than the process time of the arbiter. Because of technical reasons which are described in the main text, the arbiter has been tuned so that it handles incoming requests with a minimum latency of 2 μs. The plot has been cropped for ISIs above 80 μs so that we can see the detail in the smaller ISI distributions. (a) ISI distribution of retina spikes in response to a moving edge oriented at 0°. (b) ISI distribution of retina spikes in response to a moving edge oriented at 90°. Since more requests are made by the retina pixels to its on-chip arbiter, the ISI of the spikes should reflect the speed of the arbiter in processing the requests. Most ISIs were around 2–4 μs.

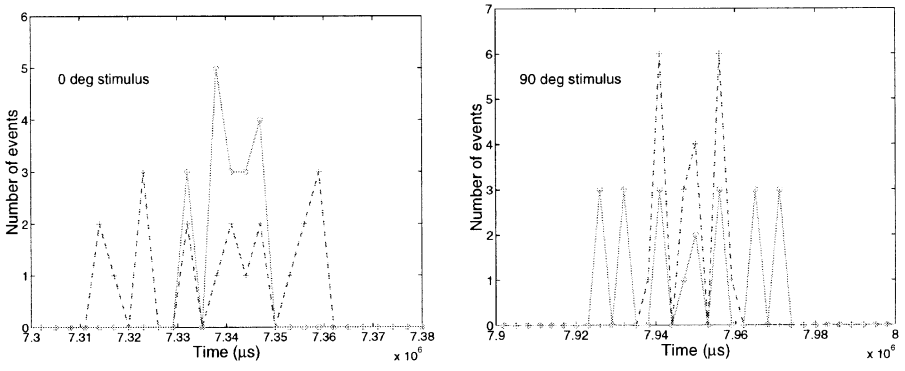

(a) (b)

Fig. 10. Interval histogram (bin size 10 ms) of retina OFF spikes collected from the experiment shown in Fig. 8. (a) Stimulus was oriented at 0°. Two templates were used for the histogram. The histogram of the spikes using the horizontal template of 9×5 pixels (sensitive to 0° orientation) is shown by the solid curve and the histogram of spikes from the vertical template of 5×9 pixels is shown by the dashed curve. The maximum number of spikes per bin time using the horizontal template is greater than that collected with the vertical template. The histogram shows the possibility of a process to distinguish the two stimulus orientations using the two different templates. (b) Stimulus was oriented at 90°. The maximum number of spikes per bin time for the vertical template is now higher than that for the horizontal template. A threshold can be used to separate the outputs of the two templates.

proposed by Hubel and Wiesel (1962) can lead to orientation-selective neurons.

6. AER mapping by PIC microcontroller

We used a PIC 16C74 microcontroller[1] to interface between the silicon retina and the multi-neuron chip (Fig. 11). The 16C74 is an 8-bit processor running at an instruction rate of about 4–5 MIPS. It has four 8-bit digital ports. The retina events are captured on two ports (one port plus one bit) while another port is used to transmit events to the receiver chip. The microcontroller filters each event to decide if it lies in one or more of the receptive fields (RFs) of the neurons on the receiver. If it does, an event is transmitted to the appropriate receiver neuron.

The PIC is connected to an RS232 serial port on a laptop computer, from which the templates (or RF sizes) are set. The program on the PIC consists of about 500 lines of assembly language,[2] which is assembled to about 300 14-bit words of code. The program runs in a tight loop looking for a sender event. Commands sent from the laptop interrupt

[1] www.microchip.com

[2] Program is available from http://www.ini.unizh.ch/~tobi/aerpic

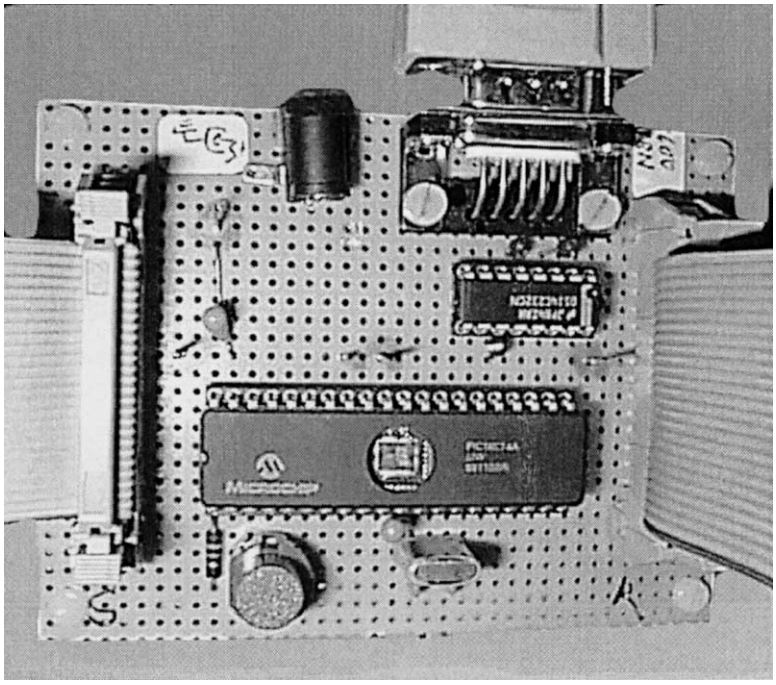

Fig. 11. The PIC interface board, comprising a PIC 16C74 microcontroller, a 20 MHz crystal, a DS14C232 RS232 interface (single +5 V supply), and ribbon connectors for sender and receiver.

this tight loop, and are parsed. The PIC has been programmed to provide other functions. It can generate a trigger spike when it detects a particular spike address, it can stimulate a receiver neuron at a specified frequency, and it can project spikes from the sender to multiple receiver addresses by adding a table of offset values to each sender's address.

As soon as the PIC receives an event, it filters the event against possible receptive fields and takes appropriate action. Once the PIC has completed its action (including the handshaking process to the transceiver, if the spike is within one of the templates), it acknowledges the sender. The exact time taken for transmission of a single spike from the sender to the receiver depends on the location of the spike address in the list of RFs. Typically, it is about 15 μs (Fig. 12). The cycle time of 15 μs can be reduced by using a faster processor in place of the PIC. The retina and transceiver chips can handle handshaking cycle times on the order of 100 ns.

This cycle time is sufficient to transmit about six transceiver neuron addresses within a retina ISI of 100 μs. The statistics of the ISI distribution would then not be distorted by the presence of the PIC. We show the ISI distribution of the OFF-transient retina spikes in the non-bursting mode (Fig. 13(a)) in response to a stimulus oriented at 0° and the ISI distribution of spikes from the PIC (Fig. 13(b)) to the transceiver neuron that is sensitive to 0° orientation. The statistical distribution in both cases looks similar, thus the PIC does not distort the distribution of the retina spikes. However, if the retina is in the bursting mode with an average ISI below 10 μs, or the retina pixel is routed to more than six transceiver neurons, then the statistics of the distribution to the transceiver would be altered by the PIC's presence.

7. Multi-neuron network

The transceiver chip consists of a ring of 16 integrate-and-fire neurons that excite a global inhibitory neuron. The inhibitory neuron inhibits all the excitatory neurons (see Fig. 14). The excitatory neurons are locally coupled to their nearest neighbors. Each excitatory neuron can be driven externally through six excitatory synapses and two inhibitory synapses using the AER protocol. The circuit of a neuron and excitatory synapse are shown in Fig. 15. The synapse circuit (M1–M4) in the left box of the figure was originally described in Boahen (1996). The presynaptic spike drives the transistor M4, which acts like a switch. The bias voltages V_w and V_e set the strength and the time constant of the synapse. The strengths of the inhibitory synapses on this chip are set by two global biases. The strengths of the six excitatory synapses are set by two other global biases and by different transistor sizes for M3 in each synapse.

The circuit in the right box of Fig. 15 implements a linear threshold integrate-and-fire neuron with an adjustable voltage threshold, spike pulse width and refractory period. If a DC current I_{inj} is sourced into the neuron's membrane capacitance C_m, its membrane voltage V_{mem} increases accordingly. The transconductance amplifier M5–M9 compares V_{mem} to a threshold voltage V_{thr}. As soon as V_{mem}

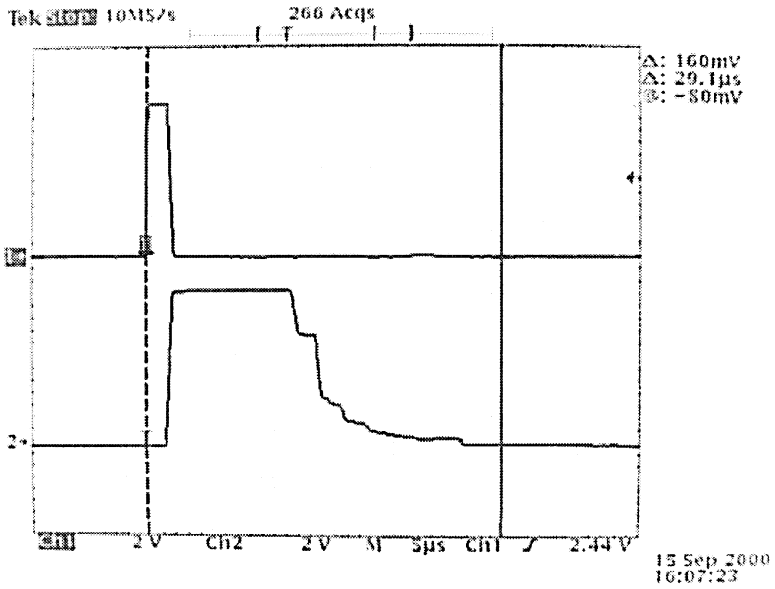

Fig. 12. Averaged traces of the REQ signal from the retina (top) and the PIC's ACK signal (bottom). The PIC's acknowledge signal goes high about 2 μs after the retina's request. This latency is due to the fact that the PIC has an instruction rate of 5 MIPS and it takes an average of 10 instructions for the PIC to acknowledge the retina. The PIC's acknowledge signal goes low about 15–25 μs after the PIC has either rejected the spike, or sent a spike to the appropriate receiver neuron.

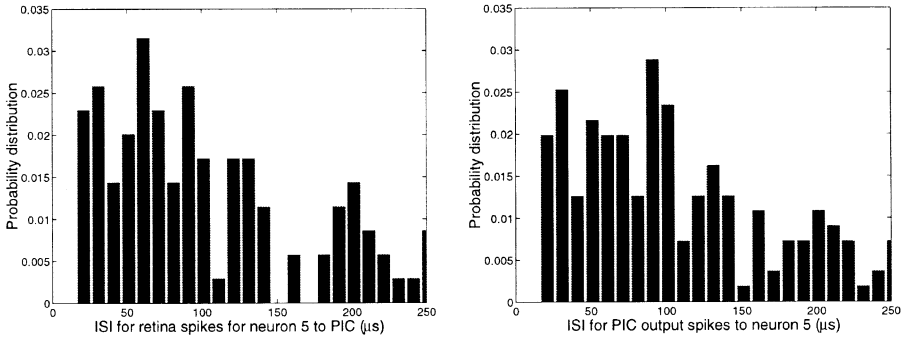

(a) (b)

Fig. 13. ISI distributions at the input and output of the PIC. (a) ISI distribution of the retina OFF-transient spikes in the non-bursting mode. The spikes were collected in response to a preferred stimulus oriented at 0°. (b) ISI distribution of the mapped spikes from the output of the PIC to the target neuron on the transceiver chip whose preferred orientation is 0°. The statistical distributions for both curves look similar.

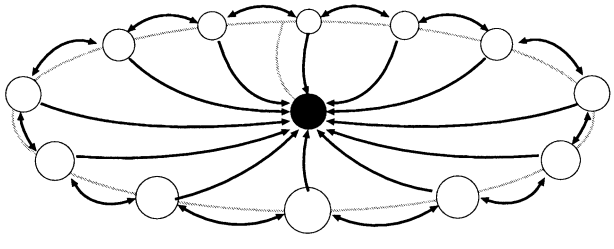

Fig. 14. Architecture of the ring of integrate-and-fire neurons chip. The unfilled circles represent the excitatory neurons. The filled circle represents the global inhibitory neuron. The gray line symbolizes the global inhibitory connection to all the excitatory neurons, while the dark arrows denote excitatory connections.

exceeds V_{thr}, the amplifier's output voltage switches to the positive rail V_{dd}. The two inverters (M10–M12 and M13–M15) implement a non-inverting high-gain amplifier; hence V_{out} also switches to V_{dd}. The capacitor C_{fb}, together with C_m implement a capacitive divider that provides positive feedback to the node, V_{mem}. This feedback speeds up the circuit's response and provides hysteresis to ensure that small fluctuations of V_{mem} around V_{thr} (e.g. due to noise) do not make V_{out} switch erratically.

When V_{out} switches to V_{dd}, the transistor M21 is switched on. This switch discharges the charge on the membrane capacitor C_m to flow to ground through transistors M20

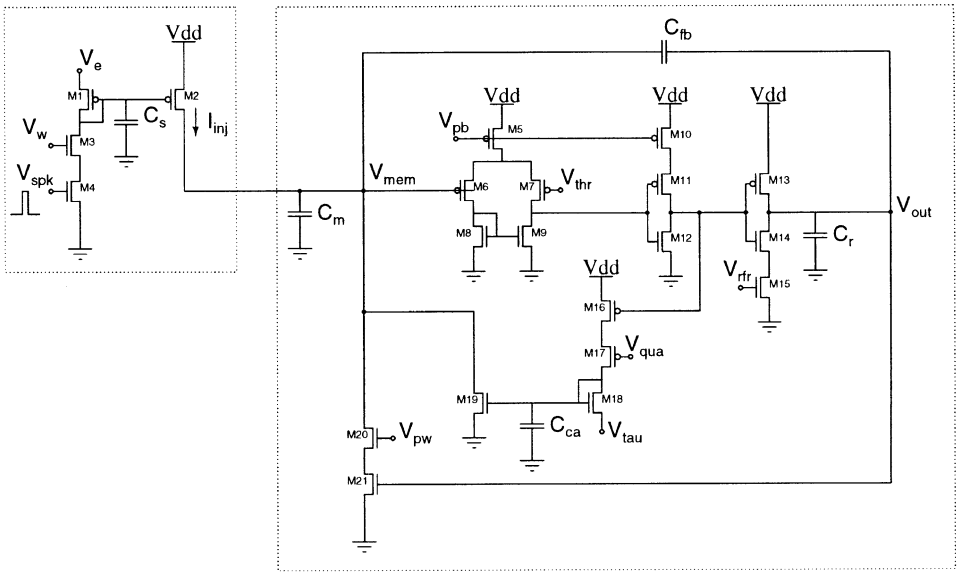

Fig. 15. Circuit diagram of an excitatory synapse (left box) connected to a linear threshold integrate-and-fire neuron (right box).

and M21 at a rate controlled by V_{pw}. This bias voltage can be used to control the spike's pulse width. As C_m is discharged, V_{mem} decreases. Once V_{mem} decreases below V_{thr}, the transconductance amplifier switches to ground. The first inverter then switches to V_{dd}, but the output of the second inverter V_{out} does not immediately go to zero; it decreases linearly at a rate set by V_{rfr}. In this way, transistor M21 is kept on, even after V_{mem} has decreased below the neuron's threshold voltage. As long as the gate voltage of M21 is sufficiently high, the neuron is in its refractory period. Spikes cannot be generated because the current injected in the membrane capacitance C_m is discharged directly to ground (by setting V_{pw} so that the current through M20 is larger than I_{inj}). Once the transistor M21 is turned off, a new spike can be generated in a time that is inversely proportional to the magnitude of I_{inj}. Fig. 16(a) shows a spike trace measured from one of the 16 neurons, as it was being stimulated by a 100 Hz pulse-train (see V_{mem} of Fig. 15).

Fig. 16(b) shows the output frequency of the neuron in response to the magnitude of the input DC current, for three different values of the refractory period bias setting V_{rfr}. This plot is commonly referred to as an F–I curve. Transistors M16–M19 implement a *spike frequency adaptation* mechanism (Boahen, 1996). With every spike, M16 is switched on and a fixed amount of charge (set by V_{qua}) is dumped onto the capacitor C_{ca}. The integrated charge on the capacitor is balanced by the charge that leaks off through the transistor M18. The final charge on the capacitor sets the current that flows through M19. This current is subtracted from the input current, and the neuron's spike frequency decreases accordingly. The voltages V_{qua} and V_{tau} can together be used to set the gain and time constant of the integrator. Fig. 17(a) shows the time taken for the neuron to reach its adapted spike frequency for four different values of V_{qua}. The neuron was stimulated with a DC current such that it spiked at a rate of 100 Hz (in steady state, without spike

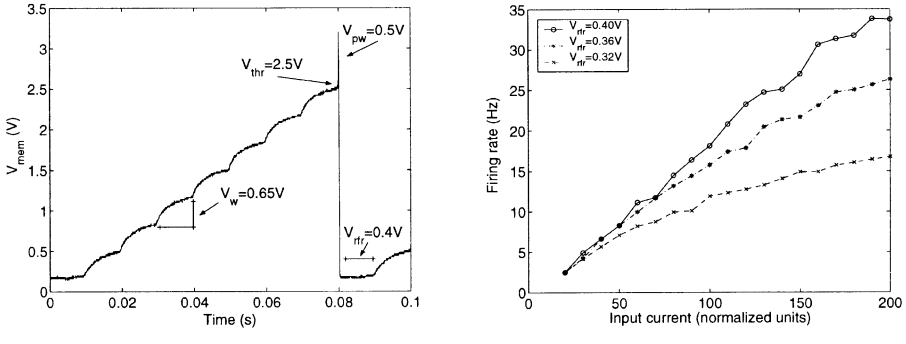

(a) (b)

Fig. 16. (a) The membrane potential of one of the 16 excitatory neurons as the neuron was being stimulated by a 100 Hz input spike train. The circuit's bias parameters and their effect on the neuron's spiking properties are labeled in the figure. (b) Plot of the frequency of the output spikes versus the magnitude of a DC input current, for three values of V_{rfr}.

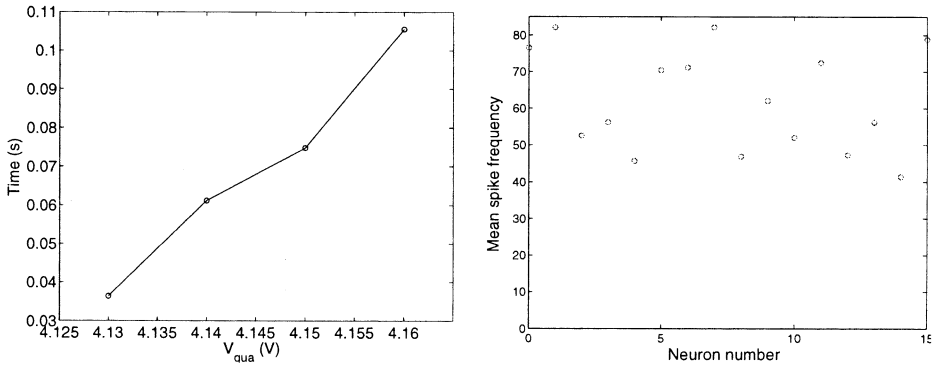

Fig. 17. (a) Spike frequency adaptation times: In this experiment, the neuron was stimulated by a constant current so that it spiked at a constant rate of 100 Hz (in the absence of any spike frequency adaptation). In the presence of spike frequency adaptation, the output frequency of the neuron adapted from 100 to 67 Hz. The time that the neuron takes to reach its steady state firing frequency changes as a function of V_{qua}. (b) Distribution of the mean output spiking rates of the 16 neurons on the chip for the same presynaptic stimulations. Due to mismatches, the steady state firing rate varies among the neurons.

frequency adaptation). In the example shown, spike frequency adaptation brought the steady state firing rate down to 67 Hz. Depending on the value of V_{qua}, this effect can be immediate, requiring only the generation of one single spike, or more gradual.

The neurons on this chip have an inherent mismatch which arises from both the synapse and the soma of the neuron. This mismatch leads to a variance in the firing rates of the neurons that is observed when the neurons are stimulated by the same presynaptic frequency (see Fig. 17(b)).

8. Responses of orientation-selective neurons

The receptive fields of the orientation-selective neurons were synthesized by mapping the outputs of a selected set of neurons on the retina as shown in Fig. 18. In these experiments, only two neurons on the transceiver chip are mapped for orientation selectivity. These two neurons have orthogonal preferred orientations. The local excitatory coupling between the neurons was disabled. There is no self excitation to each neuron so we explored only a feedforward model and a feedback model using inhibition. We varied the size and aspect ratio of the receptive fields of the neurons by changing the template size used in the mapping of the retina spikes to the transceiver chip. The template size and aspect ratio determine the orientation responses of the neurons. However, the orientation response of these neurons also depends on the time constant of the neuron. On this multi-neuron chip, we do not have an explicit transistor that allows us to control the time constant of the neuron. Instead, we generated a leak current through M19 in Fig. 15 by controlling the source voltage of M18, V_{tau}. By increasing V_{tau}, we decrease the time constant of the neuron. Since the neuron charges up to threshold through the summation of the incoming EPSPs, if the ISIs of the incoming spikes are too large, the neurons will not be able to charge up to threshold. If the ISI of the presynaptic spikes is shorter than the time constant of the neuron, the neuron will be able to

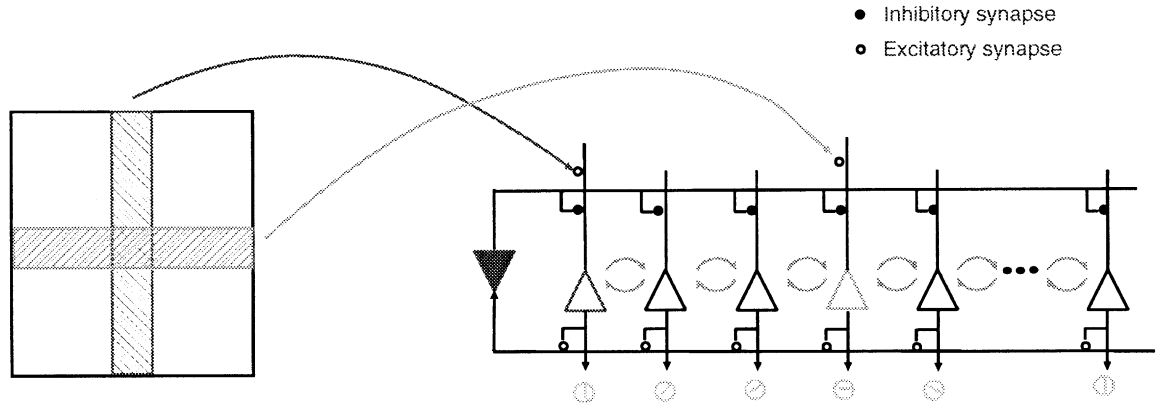

Fig. 18. Spikes from a selected set of neurons within the two rectangular regions on the retina are mapped by the PIC onto the corresponding orientation-selective neurons on the transceiver chip. The right-side-up triangles mark the somas of the excitatory neurons and the upside-down triangle marks the soma of the global inhibitory neuron. Only two neurons, which are mapped for orthogonal orientations, were used in this experiment.

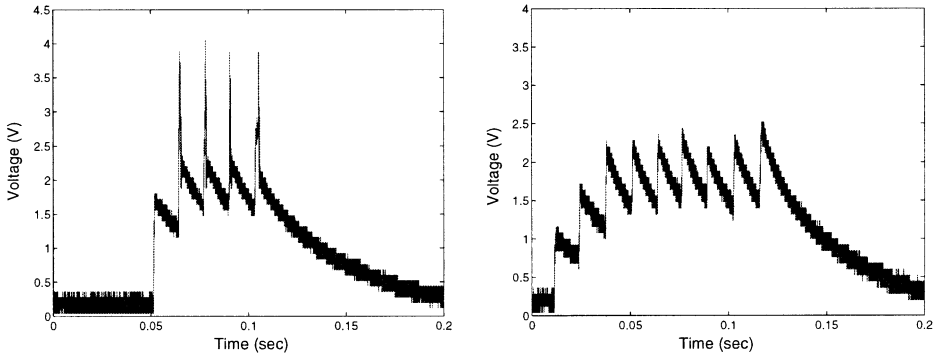

(a) (b)

Fig. 19. The response of the membrane potential of a neuron to a stimulus of its preferred orientation (a) and to a stimulus of its non-preferred (orthogonal) orientation (b). Since the ISIs of the presynaptic spikes are smaller in the case of the preferred orientation, the neuron is able to charge up to threshold more often. The membrane potential leaks away when there are no presynaptic spikes. The time constant of the membrane is about 10 ms.

charge up to threshold as shown in Fig. 19(a). This figure shows the response of the membrane potential to its preferred stimulus orientation, while Fig. 19(b) shows the membrane potential of the neuron in response to the non-preferred stimulus orientation. The incoming EPSPs cannot sum to threshold, because the ISIs of the input spikes are too large. The synaptic weight also affects the responses of the orientation-selective neurons as the weight determines the number of EPSPs needed to drive the neuron above threshold.

We first investigated the feedforward model by using a template size of 5×7 ($3^\circ \times 4.2^\circ$) for one neuron and 7×5 ($4.2^\circ \times 3^\circ$) for the second neuron. (We have repeated the following experiments using smaller template sizes, 3×5 and 1×3 and the experimental results were pretty much the same. We show results from the 5×7 template whose aspect ratio is closest to 1.) The aspect ratio of this template was 1.4. The time constant of the neuron and synaptic gain and strength were adjusted so that both neurons responded optimally to the stimulus. The inhibitory connection from the global inhibitory neuron to the two excitatory neurons was disabled. As in the retina experiments in Section 5, a rotating drum with a black and white strip was placed in front of the retina. Data were collected from the multi-neuron chip for different orientations of the drum (hence the stimulus) spaced at 30° intervals. The image speed of the stimulus was approximately 7.9 mm/s (or 89 pixels/s). The stimulus was presented to the retina approximately 500–1000 times and the output spikes from the orientation-selective neurons were collected using a logic analyzer. The retina was set in the non-bursting mode; the refractory period of the neuron was around 500 μs. Only the OFF transients from the retina were used as input spikes to the transceiver neurons. The ISI distributions of the mapped input spikes to the two neurons when the retina was presented with a 0° oriented stimulus are shown in Fig. 20. The input distribution of spikes to the neuron that is sensitive to this orientation (Fig. 20(b)) shows a larger percentage of spikes with low ISIs.

Since the orientation-selective neurons respond with only

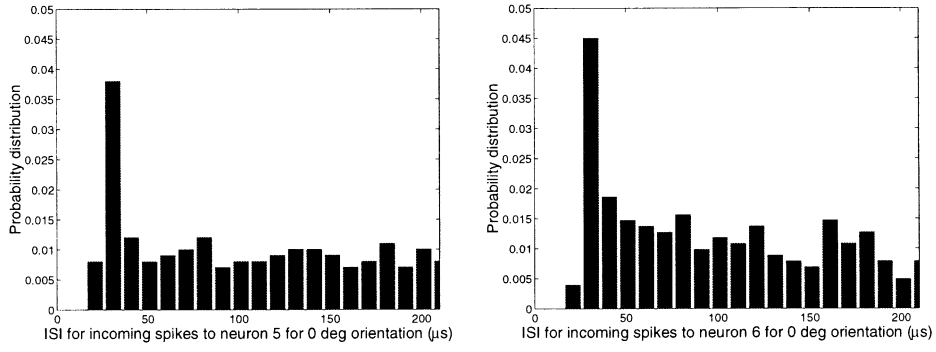

(a) (b)

Fig. 20. ISI distribution of OFF-transient retina spikes (a) that are mapped to neuron 5 (sensitive to vertical oriented stimulus) and of retina spikes (b) that are mapped to neuron 6 (sensitive to horizontal oriented stimulus). The retina spikes were collected during a stimulus presentation at 0° orientation. Neuron 6 has a larger number of spikes at low ISIs compared to neuron 5.

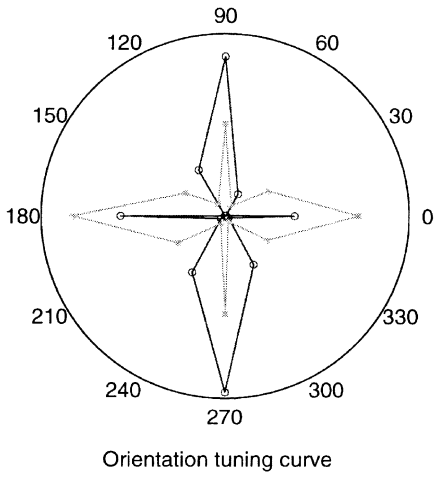

Fig. 21. Orientation tuning curves of two neurons that have orthogonal preferred orientations in the case of a feedforward model in the absence of inhibition. The receptive field sizes of the neurons are 5×7 and 7×5, respectively. The responses of the neurons in this figure and the remaining figures are the number of spikes collected per stimulus presentation. The maximum radius of the polar plot has been normalized to the maximum response of both neurons. The data were collected for stimulus orientations spaced at 30° intervals. The neuron (neuron 5) that responds preferably to a 90° oriented stimulus (solid curve) has a smaller response to a stimulus at 0° orientation. The same observation is true for the other neuron (dashed curve), neuron 6. The orientation-selective (OS) index for neuron 5 is 0.323 and for neuron 6 it is 0.195. Refer to the main text for the definition of the OS index.

1–3 spikes every time the stimulus moves over the retina, we normalized the total spikes collected in these experiments to the number of stimulus presentations. The results are shown as a polar plot in Fig. 21 for the two neurons that are sensitive to orthogonal orientations. The radius of the plot was normalized to the largest neuron response. The figure shows that each neuron responds better to its preferred orientation even though they both responded to the non-preferred orientations. The neuron responded with more spikes to the orthogonal orientation than the in-between orientations, because there are a small number of retina spikes that arrive with a small ISI when the orthogonal-oriented stimulus moves across the template space of the retina (see Fig. 5). We used an orientation-selective (OS) index to quantify the response of the neuron to all orientations. This index is defined as

$$\mathrm{OS} = \frac{\mathrm{R(preferred)} - \mathrm{R(nonpreferred)}}{\mathrm{R(preferred)} + \mathrm{R(nonpreferred)}}$$

where R(preferred) is the response at the preferred orientation and R(nonpreferred) is the response at the orthogonal orientation. As an example, R(preferred) for neuron 5, which is sensitive to vertical orientations, is R(90) + R(270) and R(nonpreferred) is R(0) + R(180).

We next investigated the feedback model in the presence of global inhibition. In this case, the system acts like a soft winner-take-all circuit. Since the inhibitory neuron is driven by all of the excitatory neurons, it is itself not orientation-selective. We tuned the coupling strengths between the excitatory neurons and the inhibitory neuron so that we obtained the optimal response to the same stimulus presentations as in the feedforward case. The new tuning curves are plotted in Fig. 22(a). The orientation tuning curves of both neurons were sharpened by the recurrent inhibition. The neuron with the higher firing rate suppresses the response of the other neuron (cross-orientation inhibition). The response of the inhibitory neuron is shown in Fig. 22(b). Its output is a combination of the responses of both neurons.

The orientation tuning curves for both neurons in the

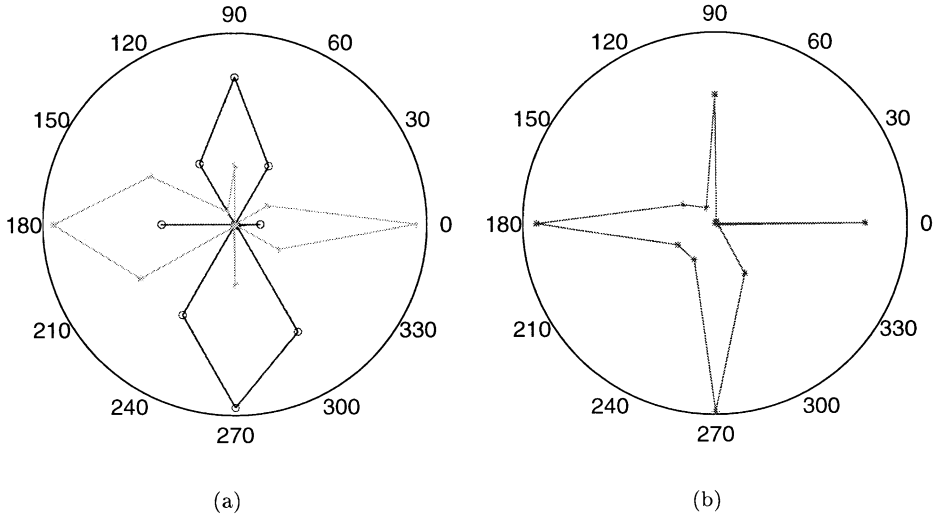

Fig. 22. Orientation tuning curves of the two orientation-selective neurons and the global inhibitory neuron in the presence of inhibition (feedback model). (a) Tuning curves of the two orientation-selective neurons described in Fig. 21. The neurons now respond less to the nonpreferred orientation. The neuron with the higher firing rate suppresses the response of the other neuron (cross-orientation inhibition). The output firing rates are also lower in this case (approximately half of the firing rates in the absence of inhibition). The OS index for neuron 5 is 0.546 and for neuron 6 is 0.4973. (b) Orientation tuning curve of the global inhibitory neuron. The neuron is not orientation-selective.

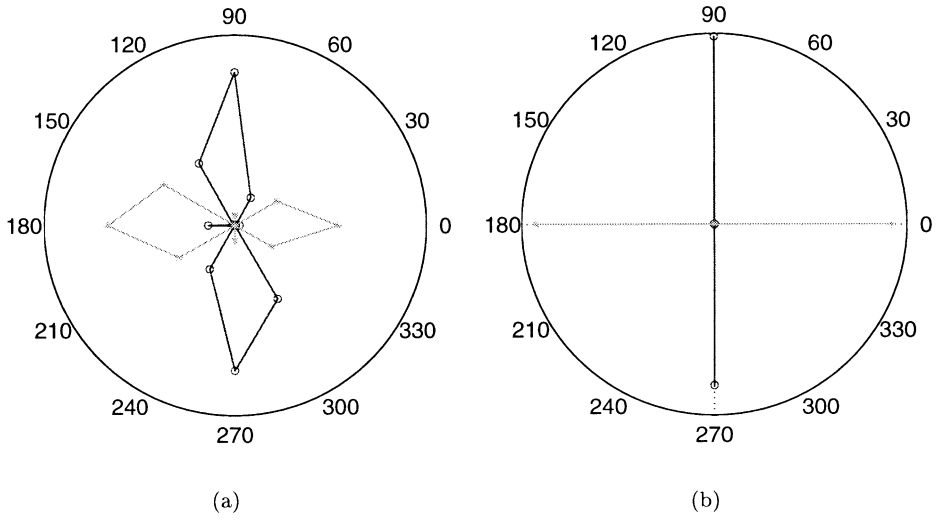

Fig. 23. Responses of the two orientation-selective neurons whose receptive fields are created using templates with aspect ratios of 1.8 and 3 in the feedback configuration. The orientation tuning curves of these neurons whose receptive field sizes are 5×9 and 9×5 are shown in (a). The tuning curves are sharper than the curves in Fig. 22. The OS index for neuron 5 is 0.81 and for neuron 6 is 0.79. In (b), the tuning is even sharper when we used template sizes of 3×9 and 9×3 to create the receptive fields of the neurons. The OS index for both neuron 5 and neuron 6 is 1. As expected, using a template with a larger aspect ratio leads to a better orientation selectivity for the neurons. The plots have been normalized for the maximum response recorded from the neurons.

presence of inhibition were determined using two other template sizes (5×9 and 3×9). These sizes translate to aspect ratios of 1.8 and 3, respectively. The results are shown in Fig. 23(a) for the template aspect ratio of 1.8 and Fig. 23(b) for the template aspect ratio of 3. In the case of the template with an aspect ratio of 3, the neurons showed no response to the non-preferred stimulus, even though the number of EPSPs that the neurons receive per stimulus presentation is approximately half that for the other template sizes. As expected, using a template with a larger aspect ratio leads to a better orientation selectivity for the neurons.

9. Conclusions

We demonstrated a programmable multi-chip VLSI system that can be used for exploring spike-based processing models. This system has advantages over computer neuronal models in that it is real-time and the computational time does not scale with the size of the neuronal network. The spiking neurons can be configured for different computational properties. Interchip and intrachip connectivity between neurons can be programmed using the AER protocol. In this work, we created the receptive fields for the orientation-tuned spiking neurons by mapping the transient spikes from a silicon retina using a microcontroller. We have not mapped onto all the neurons on the transceiver chip because the PIC microcontroller we used is not fast enough to create receptive fields for more neurons without distorting the ISI distribution of the incoming retina spikes. We are in the process of testing faster boards that can handle the mapping of more neurons.

We evaluated the responses of the orientation-tuned spiking neurons for different receptive field sizes and aspect ratios and also in the absence and presence of feedback inhibition. In a feedforward model, the aVLSI spiking neurons show orientation selectivity similar to digital simulation results obtained from continuous-valued neurons. Adding feedback inhibition increased the selectivity of the spiking neurons.

We can extend the multi-chip VLSI system in this work to a more sophisticated system that supports multiple senders and multiple receivers. Such a system can be used, for example, to implement multi-scale cortical models (Grossberg, Mingolla & Williamson, 1995). The success of this system opens up the way for more elaborate spike-based emulations in the future.

Acknowledgements

We acknowledge Timothy Horiuchi for discussions on the transceiver chip design and David Lawrence for the software driver developed in this work. This work was supported in part by the Swiss National Foundation Research SPP grant and the Köbler Foundation.

References

Abusland, A., Lande, T., & Hvin, M. (1996). Architecture for stochastically pulse-encoded analog signals. In *Proc. IEEE Intl. Symp. Circ. & Syst.* (vol. III, pp. 401–404).

Ben-Yishai, R., Bar-Or, R. L., & Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proc. Natl. Acad. Sci. USA*, *92* (9), 3844–3848.

Boahen, K. A. (1996). A retinomorphic vision system. *IEEE Micro*, *16* (5), 30–39.

Boahen, K. A. (1999). Retinomorphic chips that see quadruple images. In *Proceedings of the Seventh International Conference on Microelectronics for Neural, Fuzzy and Bio-inspired Systems; Microneuro '99* (pp. 12–20), Los, Alamito, CA. IEEE Comp. Soc. Press.

Boahen, K. A. (2000). Point-to-point connectivity between neuromorphic chips using address events. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, *47* (5), 416–434.

Boahen, K. A., & Andreou, A. G. (1992). A contrast sensitive silicon retina with reciprocal synapses. *Advances in Neural Information Processing Systems*, *4*, 764–772.

Boahen, K. A., Andreou, A., Hinck, T., Kramer, J., & Whatley, A. (1997). Computation- and memory-based projective field processors. In T. Sejnowski, C. Koch & R. Douglas (Eds.), *Telluride NSF Workshop on Neuromorphic Engineering, Telluride, CO*.

Cauwenberghs, G., & Waskiewicz, J. (1999). Analog VLSI cellular implementation of the Boundary Contour System. In M. Kearns, S. Solla & D. Cohn, *Advances in neural information processing systems* (pp. 756–763). , vol. 11. Cambridge, MA: MIT Press.

Cohen, M., Edwards, T., Stanford, M., Cauwenberghs, G., Andreou, A., & Abshire, P. (1998). One-dimensional AER-based remapping project. In *Telluride NSF Workshop on Neuromorphic Engineering, Telluride, CO*. http://www.ini.unizh.ch/telluride98/report98

Deiss, S. R., Douglas, R. J., & Whatley, A. M. (1999). A pulse-coded communications infrastructure for neuromorphic systems. In W. Maass & C. M. Bishop, *Pulsed neural networks* (pp. 157–178). Boston, MA: MIT Press.

Delbrück, T., & Mead, C. (1994). *Analog VLSI phototransduction by continuous-time, adaptive, logarithmic photoreceptor circuits*. CNS Memo 30, California Institute of Technology, Pasadena, CA.

Douglas, R., & Martin, K. (1991). A functional microcircuit for cat visual cortex. *J. Physiol.*, *440*, 735–769.

Douglas, R., Koch, C., Mahowald, M., Martin, K., & Suarez, H. (1995). Recurrent excitation in neocortical circuits. *Science*, *269* (5226), 981–985.

Etienne-Cummings, R., & Cai, D. (1997). A general purpose image processing chip: Orientation detection. In M. Jordan, M. Kearns & S. Solla, *Advances in neural information processing systems* (pp. 873–879). , vol. 10. Norwell, MA: MIT Press.

Ferster, D., & Koch, C. (1987). Neuronal connections underlying orientation selectivity in cat visual cortex. *Trends in Neurosciences*, *10* (12), 487–491.

Grossberg, S., Mingolla, E., & Williamson, J. (1995). Synthetic aperture radar processing by a multiple scale neural system for boundary and surface representation. *Neural Networks*, *8* (7/8), 1005–1028.

Higgins, C. M., & Koch, C. (1999). Multi-chip neuromorphic motion processing. In *Conference on Advanced Research in VLSI, Atlanta, GA*.

Higgins, C. M., Deutschmann, R. A., & Koch, C. (1999). Pulse-based 2-d motion sensors. *IEEE Transactions on Circuits and Systems II*, *46* (6), 677–687.

Hubel, D., & Wiesel, T. (1962). Receptive fields, binocular interaction and functional architecture. *J. Physio. (Lond.)*, *160*, 106–154.

Indiveri, G., Whatley, A. M., & Kramer, J. (1999). A reconfigurable neuromorphic VLSI multi-chip system applied to visual motion computation. In *Proceedings of the Seventh International Conference on Microelectronics for Neural, Fuzzy and Bio-inspired Systems: Microneuro '99* (pp. 37-44), Los Alamitos, CA. IEEE Computer Society.

Kramer, J. (2001). An integrated optical transient sensor (submitted).

Kramer, J., Sarpeshkar, R., & Koch, C. (1997). Pulse-based analog VLSI velocity sensors. *IEEE Trans. Circuits and Systems II: Analog and Digital Signal Processing*, *44* (2), 86–101.

Lazzaro, J., Wawrzynek, J., Mahowald, M., Sivilotti, M., & Gillespie, D. (1993). Silicon auditory processors as computer peripherals. *IEEE Transactions on Neural Networks*, *4* (3), 523–528.

Liu, S., & Boahen, K. (1996). Adaptive retina with center-surround receptive field. In D. Touretsky, M. Mozer & M. Hasselino, *Advances in neural information processing systems* (pp. 678–684). , vol. 8. Norwell, MA: MIT Press.

Mahowald, M. (1994). *An analog VLSI system for stereoscopic vision*, Boston: Kluwer.

Mahowald, M., & Mead, C. (1991). The silicon retina. *Scientific American*, *264* (5), 76–82.

Mortara, A., Venier, P., & Vittoz, E. A. (1995). A communication scheme for analog VLSI perceptive systems. *IEEE Journal of Solid-State Circuits*, *30* (6), 660–669.

Raffo, L., Sabatini, S., Bo, G., & Bisio, G. (1998). Analog VLSI circuits as physical structures for perception in early visual tasks. *IEEE Transactions on Neural Networks*, *9* (6), 1483–1494.

Serrano-Gotarredona, T., Andreou, A., & Linares-Barranco, B. (1999). AER image filtering architecture for vision-processing systems. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, *46* (9), 1064–1071.

Shi, B. (2000). A low-power orientation-selective vision sensor. *IEEE Transactions on Circuits and Systems*, *47* (5), 435–440.

Sillito, A. (1992). GABA mediated inhibitory processes in the function of the geniculo-striate system. In R. Mize, R. Marc & A. Sillito, *GABA in the retina and central visual system* (pp. 349–384). Elsevier.

Somers, D., Nelson, S., & Sur, M. (1995). An emergent model of orientation selectivity in cat visual cortex simple cells. *Journal of Neuroscience*, *15* (8), 5448–5465.

Venier, P., Mortara, A., Arreguit, X., & Vittoz, E. A. (1997). An integrated cortical layer for orientation enhancement. *IEEE Journal of Solid-State Circuits*, *32* (2), 177–185.

Whatley, A., Kramer, J., & Douglas, R. (1997). ON/OFF retina to silicon cortex. In T. Sejnowski, C. Koch & R. Douglas (Eds.), *Telluride NSF Workshop on Neuromorphic Engineering, Telluride, CO*.