Available online at www.sciencedirect.com
SCIENCE DIRECT®
Vision Research 44 (2004) 2083–2089
Vision Research
www.elsevier.com/locate/visres

A silicon early visual system as a model animal

Tobi Delbrück *, Shih-Chii Liu

Institute of Neuroinformatics, University of Zürich and ETH Zürich, Winterthurerstrasse 190, CH-8057 Zürich, Switzerland

Received 3 July 2003; received in revised form 7 March 2004

Abstract

Examples that show the transfer of our basic knowledge of brain function into practical electronic models are rare. Here we present a user-friendly silicon model of the early visual system that contributes to animal welfare. The silicon chip emulates the neurons in the visual system by using analog Very Large Scale Integration (aVLSI) circuits. It substitutes for a live animal in experiment design and lecture demonstrations. The neurons on this chip display properties that are central to biological vision: receptive fields, spike coding, adaptation, band-pass filtering, and complementary signaling. Unlike previous laboratory devices whose complexity was limited by the use of discrete components on printed circuit boards, this battery-powered chip is a self-contained patch of the visual system. The realistic responses of the chip's cells and the self-contained adjustment-free correct operation of the chip suggest the possibility of implementation of similar circuits for visual prosthetics.

Keywords: Receptive field; Retina cells; Visual system; Neuroprosthesis; Visual physiology tool

1. Introduction

We built this electronic model of the early visual system because Matteo Carandini, a vision physiologist, asked whether we had an existing "neuromorphic" chip that he could use as a model animal for a new stimulation and recording setup he was developing. It was the perfect opportunity to build a practical, easy-to-use device that embodied the principles of neuromorphic chip design that we, along with others, had been developing over the past 15 years.

In this work, we integrated prior developments in neuromorphic analog Very Large Scale Integration (aVLSI) (Liu et al., 2001; Mahowald & Douglas, 1991; Mead, 1989, 1990; Mead & Mahowald, 1988) to make a user-friendly electronic model of the visual system. The high density of aVLSI enables the underlying computation to come much closer to biology than was possible with practical laboratory models built from discrete components on printed circuit boards (e.g. Harmon, 1961; Schweitzer-Tong, 1983). Although the basic functionality of the present chip has been demonstrated in previous aVLSI systems, they have never been specifically integrated into a user-friendly tool for vision physiologists, and none of the previous examples were built for adjustment-free operation. In contrast, this chip requires no parameter adjustments and is preprogrammed for optimal operation.

We do not want to overstate the properties of this device. It is not intended to realize a conceptually new model of the visual system. It is intended as a concrete realization of a practical and manufacturable neuromorphic chip, and has proven to be a useful tool that will continue to have a significant impact on experimental studies of the visual system and the teaching of its principles.

2. The user's perspective

The chip, was christened the "Physiologist's Friend" by the physiologist Kevan Martin. In the classroom, the teacher arranges the device near an overhead projector so that it views the projection screen, which serves as the tangent screen on which stimuli are presented. After plugging the device into a standard powered speaker to amplify its output to classroom volume, the teacher is ready to demonstrate some classic experiments of visual physiology. A slider switch allows the operator to select

* Corresponding author. Tel.: +41-1-635-3038; fax: +41-1-635-3053.
E-mail address: tobi@ini.phys.ethz.ch (T. Delbrück).

0042-6989/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2004.03.021

between ON- and OFF-center retinal ganglion cells and two types of cortical simple cells. Every time the cell spikes, the audience hears a loud "pop." The teacher uses a pen or sheet of paper on the bed of the overhead projector to make bar or edge stimuli on the screen, then proceeds to map the receptive field of the cell. Mapping with the chip's cells is considerably easier than with biological cells because the type of cell being recorded is known; however, the hunt for the actual location of the receptive field can still be entertaining. After explaining the general concept of a receptive field, the teacher can show how the cell responds only to local contrast and not to global illumination change, how a ganglion cell responds transiently to a global change in illumination, and how the photoreceptors adapt over time. Orientation selectivity as well as excitatory and inhibitory subregions can be demonstrated with the cortical cells. While listening to the membrane potential of one of the cortical cells, the audience can hear the excitatory and inhibitory postsynaptic potentials caused by the ganglion cell input. These demonstrations are quite compelling, and thorough discussion of the responses of a cell can easily occupy more than 30 min of interactive lecture experimentation.

In the physiology lab, the Physiologist's Friend substitutes for a live animal. It is arranged to observe the tangent screen and is plugged into a recording setup. The device has been used to train students on experimental protocols, to test data acquisition hardware and software, and to provide known reference receptive fields for automatic mapping routines and the development of new spike-triggered averaging protocols.

3. Architecture

The chip's architecture (Fig. 1) models a small patch of the early visual system, starting with the photoreceptors and ending with orientation-selective cortical simple cells. It uses continuous-time circuits with no clock signals. The only digital events are the action potentials of the neurons.

In the silicon retina part of the chip, we make simplifying assumptions about the retinal circuits that are sufficient for producing a realistic output of the silicon retina to visual stimuli. We include only photoreceptors, a horizontal cell, and bipolar cells in the outer plexiform layer, and sustained ganglion cells in the inner plexiform layer. Amacrine cells are not modeled. In a biological retina, the horizontal cells spatially aggregate the photoreceptor signals and modulate the direct input from these photoreceptors to their postsynaptic bipolar cells. This modulation allows the bipolar cell to code primarily local contrast. The response of the bipolar cell represents the result of a center–surround interaction. The ganglion cells are the output cells of the retina, sending their output to the thalamus in the form of spikes.

On the chip, the visual scene is focused onto a circular array of seven photodiodes that represent the phototransduction stage of the photoreceptors. Fig. 2a shows the chain of analog circuits in each pixel that implements retinal processing. The photodiodes supply photocurrents to a set of seven adaptive photoreceptor circuits (Delbrück & Mead, 1994). The voltage outputs of the photoreceptor circuits are a temporally band-pass filtered, logarithmically encoded representation of the

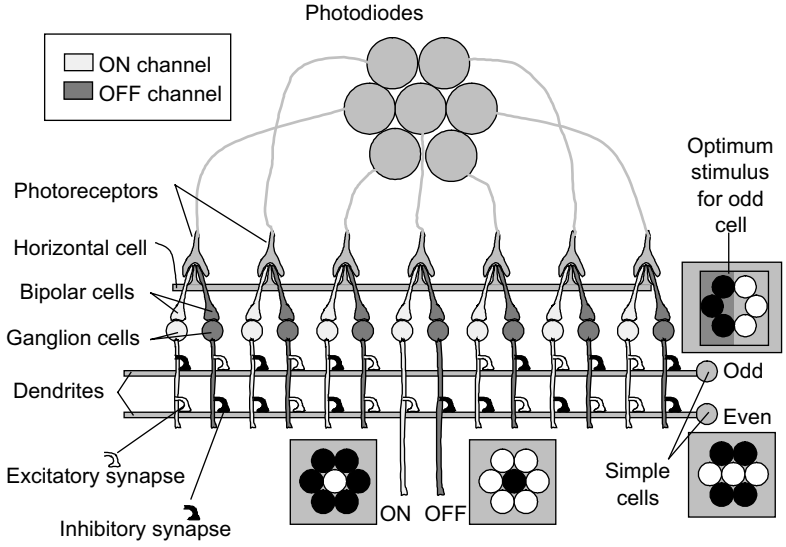

Fig. 1. Schematic architecture of the chip, together with the receptive fields of the ON and OFF ganglion cells and odd- and even-type simple cells. The grayscale in each receptive field shows the optimum stimulus at each pixel. For example, the odd-type simple cell is tuned to vertical orientations and the even-type cell is tuned to horizontal orientations.

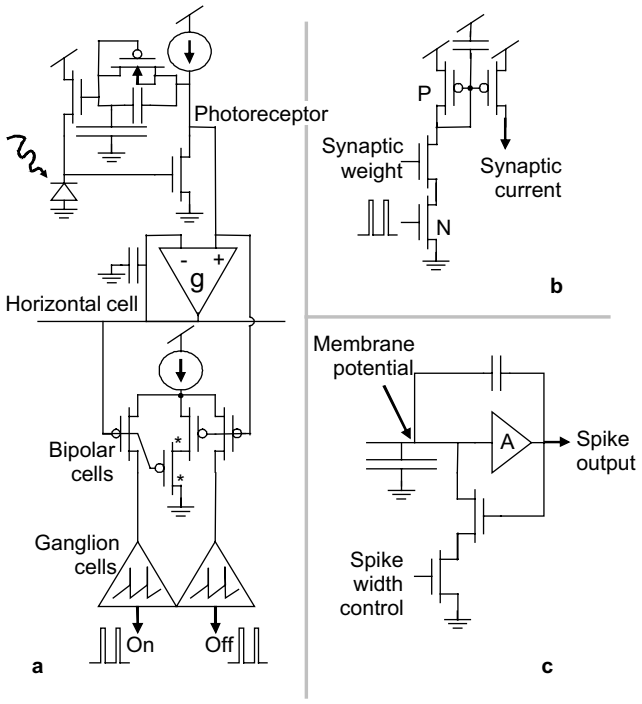

Fig. 2. Simplified schematics of circuits. (a) Retina pixel circuit; (b) synapse circuit; (c) axon-hillock circuit. An N-type transistor (N) sources electrons from the lower voltage and is analogous to a K^+ channel. A P-type transistor (P) sources holes from the higher voltage and is analogous to a Na^+ channel. The size of the capacitor symbol reflects its relative value.

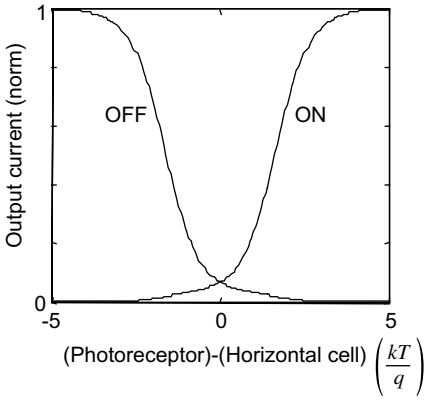

Fig. 3. Response of rectifying bipolar cell synapse circuit (antibump circuit).

local brightness, similar to biological cone responses. Like biological photoreceptors, the electronic photoreceptors have low gain for steady-state illumination, and higher gain for transient signals.

The photoreceptor outputs drive a single horizontal cell and multiple bipolar cell circuits. The horizontal cell circuit computes an average of the photoreceptor outputs using DeWeerth's follower-aggregator circuit (Mead, 1989). Each photoreceptor drives the single horizontal cell through a transconductance amplifier such that the horizontal cell output voltage takes on the average value of the photoreceptor output voltages. By using a very small bias current of 1 pA in the transconductance amplifiers, we set the transconductance so that the photoreceptor input to the horizontal cell is delayed by about 100 ms. This delay means that the horizontal cell's response lags behind the outputs of the photoreceptors, thus leading to a large portion of the transient part of the subsequent ON and OFF bipolar cells' responses. The remaining transient part of their responses is due to the slower photoreceptor adaptation.

Each ON and OFF bipolar cell pair is formed out of an antibump circuit (Delbrück, 1991) that transforms the voltage difference between its two inputs—the photoreceptor and the horizontal cell outputs—into rectified ON and OFF currents. The ON and OFF currents saturate when the voltage difference is sufficiently large. The central correlating transistors (marked $*$ in Fig. 2a) cause both ON and OFF currents to be about 1/20 of their saturation value when the difference between the inputs is zero, which occurs when contrast is absent. The response of this circuit is shown in Fig. 3. As with biological bipolar cells, rectification ensures that most of the dynamic range of each bipolar cell encodes only its own sign of contrast; it also ensures that the subsequent ON and OFF ganglion cell activity is low unless local contrast is present.

The 14 ON and OFF bipolar currents drive their individual ON and OFF ganglion cells. We used Mead's axon-hillock circuit (Mead, 1989) (Fig. 2c), which models the generation of action potential in the soma of a neuron, to implement the ganglion and cortical cells. Although a more realistic spike initiation circuit has been developed (Mahowald & Douglas, 1991), the response of this simple somatic model is sufficiently realistic to satisfy physiologists who have used this chip.

The ganglion cells excite or inhibit the two cortical cells through simple synaptic circuits (Boahen, 1997); the function of the lateral geniculate nucleus is ignored. The silicon ganglion cells form monosynaptic excitatory and inhibitory connections onto cortical cells. (The known disynaptic inhibitory connections onto cortical cells are modeled as direct inhibitory connections.) Where the axons (wires) from the ganglion cells cross the dendrites (wires) of the simple cells, we made either excitatory or inhibitory synaptic connections. The synapses (Fig. 2b) inject a current pulse (a packet of charge) onto the dendrite when they receive a spike. Excitatory synapses inject charge, and inhibitory synapses remove charge. All the excitatory synapses share a common weight, as do the inhibitory synapses. The excitatory weight is twice that of the inhibitory weight, so that the simple cells have a spontaneous firing rate.

The connections from the ganglion cells onto the cortical cells are arranged to create push-pull models

(Hubel & Wiesel, 1962) of the odd- and even-type simple receptive fields shown in Fig. 1 (odd and even refer to the symmetry of the receptive field). By a push-pull model, we mean that the cell is excited by one polarity of contrast at a spatial location and inhibited by the opposite polarity at the same location. For example, the odd-type simple cell is excited by ON ganglion cells and inhibited by OFF ganglion cells from the right side of the pixel array, and excited by OFF ganglion cells and inhibited by ON ganglion cells from the left side of the array. The odd-type simple cell is maximally excited by the black and white edge shown overlaying its receptive field in Fig. 1. We found that this push-pull scheme results in a more robust orientation response than a purely excitatory scheme with standing inhibition. This makes sense, because from an engineering perspective, a differential signal is widely used to nullify the effects of common-mode variation. Our motivation for using this push-pull mechanism was its practicality, although some recent experimental evidence from intracellular cortical recordings supports the notion of a push-pull mechanism (Anderson, Carandini, & Ferster, 2000; Borg-Graham, Monier, & Fregnac, 1998; Hirsch, Alonso, Reid, & Martinez, 1998), at least in some simple cells.

If we simply wanted to compute these cortical responses, it would not be sensible from an architectural point of view to create spikes and then immediately transform them back into analog synaptic potentials on the same chip. After all, spikes are meant for long-range transmission of neural information. The reason for this choice was to more accurately reproduce the responses of cortical cells as the result of quantal input.

Purchasers of chips such as operational amplifiers or analog-to-digital converters expect them to be usable with a minimal number of external components and adjustments. A novel feature of this neuromorphic chip is the biasing circuit (Delbrück & van Schaik, 2004) that generates the internal parameters, i.e., the bias currents and reference voltages that determine the time constants and synaptic weights. In the field of neuromorphic design, these parameters have traditionally been set using external passive components, like potentiometers. These components require careful adjustment for correct operation, making it difficult to build systems with identical behavior. They also increase the size of the final system. On this chip, the bias circuit generates 12 internal bias currents and reference voltages that are nearly independent of transistor thresholds and supply voltage variations. Thus, the chips require no calibration, and they all behave nearly identically.

Briefly, the biasing is accomplished as follows: Widlar's bootstrapped current mirror loop generates a known master reference current (Vittoz & Fellrath, 1977). Other bias currents and reference voltages are derived from this master current by Bult and Geelen's current splitter (Bult & Geelen, 1992). A single off-chip resistor sets the master reference current; therefore, the excitability of all the neurons can be simultaneously (but not individually) scaled over several decades. The ratio of the largest current (nominally 10 μA, the photoreceptor bias) to the smallest (nominally 1 pA, the synapse onto the horizontal cell) is 10^7. This huge ratio is many orders of magnitude higher than that for most analog chips, which typically require only a few similar currents for biasing amplifiers or static logic. It was possible to design a bias generator for the present chip that was fully functional in the first version because the operation of the chip is relatively simple, and values for the parameters could be estimated by hand calculation. We verified these estimates by simulating the operation of the circuits on the entire chip, which required several hours of CPU time per second of real time. A design kit (Delbrück & van Schaik, 2004) that automates the production of these bias generators is available.

The physical layout of the chip is shown in Fig. 4a. Each retina pixel requires 55 transistors and 5 capacitors. The entire chip has fewer than 800 transistors and 100 capacitors. It was fabricated in an ancient but economical 1.6-μm technology and uses an active area of about 3 mm^2. (By comparison, state-of-the-art circuits are fabricated in 90-nm technology, in which the current device would occupy an area of about 0.01 mm^2.)

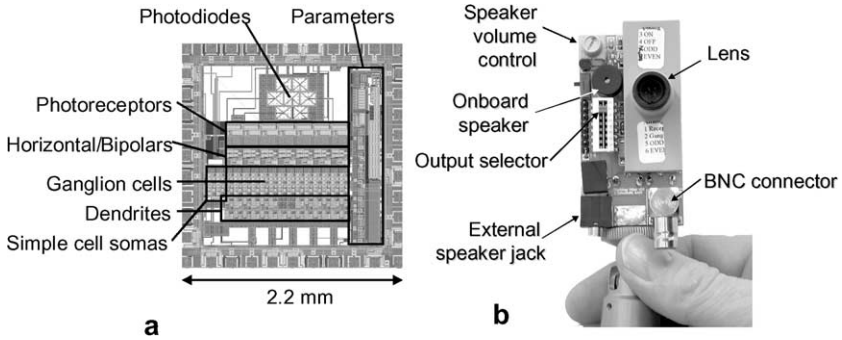

Fig. 4. Physical layout of chip and complete system. (a) Chip layout; the chip is fabricated in a standard 1.6 μm double-metal, double-poly CMOS process through MOSIS (www.mosis.org). (b) Complete system; a 9 V battery (on the back) supplies power for about 100 h of operation.

A printed circuit board (Fig. 4b) carries the chip along with its optics and battery, and provides a built-in speaker, volume control, and connections to standard physiology rigs or an external speaker. A mini tripod makes it easy to aim the chip at the tangent screen on which stimuli are presented, while a slider switch enables the user to select the desired output. The spiking outputs of the central complementary pair of ganglion cells are available for demonstration, as are the output of the central photoreceptor and the membrane potential of a ganglion cell from a different pixel. The digital spikes and membrane potentials of the two cortical simple cells can also be demonstrated.

The chip's power consumption is less than 1 mW at 5 V, of which at least 90% goes to powering the photoreceptor circuits. The complete 9 V-battery-powered system draws between 4 and 20 mA, depending upon the onboard speaker volume.

4. Cell response characteristics

How do the silicon retinal cells respond to simple visual stimuli? Fig. 5 shows the responses of some retinal cells—the central photoreceptor, the horizontal cell, and the central ON- and OFF-type ganglion cells—to two types of transient stimuli displayed on a laptop computer monitor. One of the stimuli is a full-field flash, and the other is a bright flashing spot on the central pixel. Fig. 5a shows that the full-field flash produces transient changes in the spike outputs of the ON and OFF ganglion cells in response to the rising and falling edges of the flash. These transient changes occur because the output of the horizontal cell lags behind the responses of the photoreceptors. When one type of ganglion cell fires more, the complementary type fires less. However, the sum of the firing rates is not constant. The ON and OFF ganglion cell firing rates in response to a uniform stimulus are approximately 20 times lower than the peak firing rate of either ganglion cell because of the rectification in the bipolar cell circuit (Fig. 3).

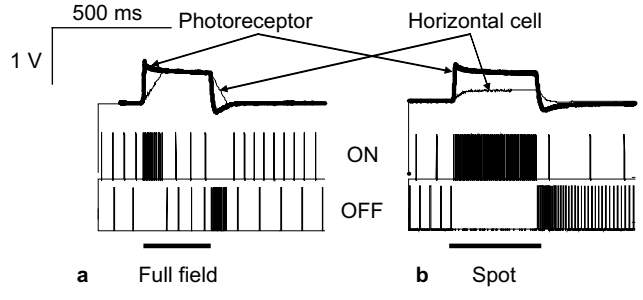

Fig. 5. Responses of chip's retinal cells. (a) Responses to a transient global increase in illumination (indicated by the bar); (b) responses to a transiently flashed bright spot on the central pixel.

Fig. 5b shows that the flashing bright spot on the central photoreceptor causes a large increase in its output voltage but only a small change in the output of the horizontal cell because the neighboring photoreceptor outputs remain unchanged. The resulting output of the central ON bipolar cell leads to a sustained increase in the spike output of the ON ganglion cell. The OFF ganglion cell responds to the removal of the bright spot because of a combination of the dark afterimage of the bright spot in the photoreceptor and the delayed response of the horizontal cell. These retinal cell responses are reasonable facsimiles of biological retinal cell responses (Dowling, 1987). The silicon ganglion cell responses most closely resemble the characteristics of biological sustained ON- and OFF-type ganglion cells.

The silicon photoreceptors adapt to a sustained change in global illumination over a time of about 30 s after their initial transient response to the step change in illumination. This photoreceptor adaptation time constant is longer than that observed in biological photoreceptors, but it is good for classroom demonstrations. The adaptation is nonlinear: the rate of photoreceptor adaptation is exponential in the voltage difference between the photoreceptor output voltage and the eventual adapted value. The start of this prolonged adaptation can be seen in the photoreceptor responses in Fig. 5.

We determined the spatial receptive field of the odd-type simple cell (Fig. 6a) by applying the reverse spike correlation method to its response to white noise drifting grating stimuli (Ringach, Sapiro, & Shapley, 1997). We also recorded the orientation tuning responses of the cell by using a drifting sinusoidal grating of optimal spatial and temporal frequency. Fig. 6b–d shows the spike raster plots, poststimulus spike histograms, and orientation tuning curve. Given the construction of the receptive field, it is no surprise that the response tuning is roughly sinusoidal. These responses are reasonable facsimiles of the responses of a biological nondirection-selective simple cell.

4.1. Differences from biological cells

What features of the biological cell responses are not captured by this silicon emulation? One difference is that the model cells adapt more slowly. To make the chip more usable for classroom demonstrations, we intentionally implemented no adaptation, except for the photoreceptor adaptation. Another difference is that the silicon neurons fire more regularly. Fig. 7a shows the intracellular membrane potential of the chip's ganglion and simple cells in steady state. The ganglion cell spike output is very regular, while the simple cell output, although noisier, does not show the interspike interval variations typical of biological cortical neurons. Fig. 7b shows the interspike interval distributions of both cells.

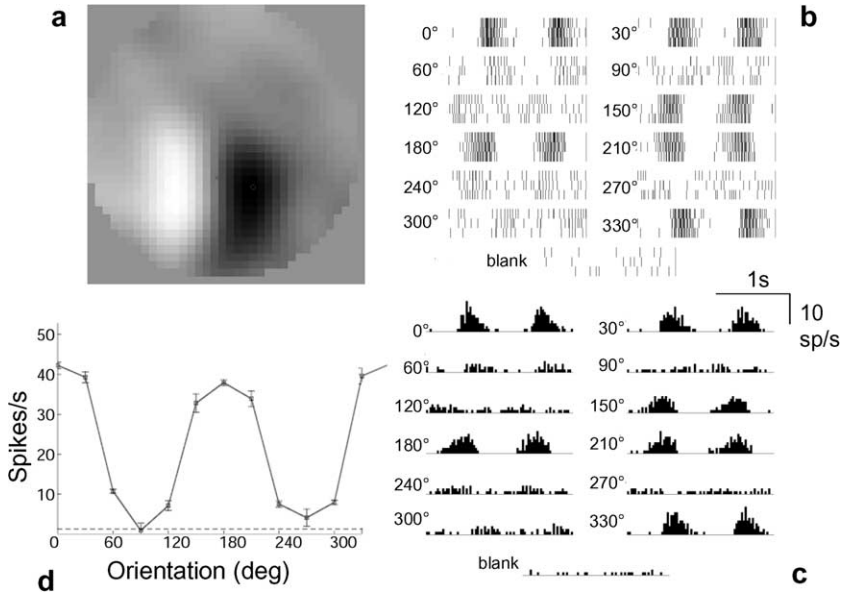

Fig. 6. Receptive-field characteristics of the odd-type simple cell. (a) Measured receptive field (see Section 4); it is mirrored horizontally by the lens optics compared with the receptive field in Fig. 1. The receptive field of the cell happened to be at the lower part of the test area. (b) Spike rasters for different orientations of a drifting sinusoidal grating (2 Hz, 0.4 cycles/deg). (c) Spike histograms for the data in (b). (d) First harmonic response tuning; dotted line shows the baseline firing rate.

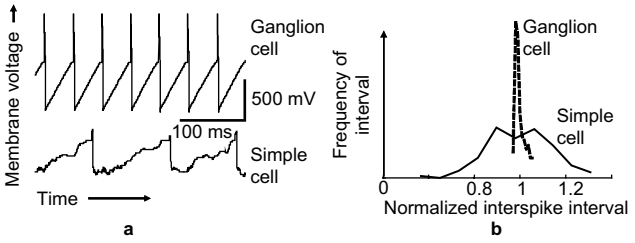

Fig. 7. Responses of a ganglion cell and a simple cell in steady state while the chip is viewing a blank scene. (a) Membrane potentials of the cells; (b) normalized interspike interval distributions.

The simple cell receives quantal input from competing excitatory and inhibitory ganglion cell synapses. Transistor mismatch makes the ganglion cells spike at slightly different rates and the synaptic weights are also randomly mismatched, which causes a complex but regular spiking pattern in the simple cell. The responses of the ganglion and simple cells to dynamic stimuli (Fig. 6b) show more trial-to-trial variability than their steady-state responses, probably because of variations in the adaptation state of the photoreceptor during stimulus presentation.

5. Discussion

This aVLSI chip is one of an evolving population of chips built using a neuromorphic approach. Many chips in the neuromorphic field are built with the express intent of understanding the function of the nervous system by emulating its structure. In some work, the aim is to present a new silicon model of neural function without regard to practical application. In other work—like the one described here—the aim is simply to fabricate a widely accepted structure for use in a practical application. In this case, we sought to make an electronic substitute of the visual system for use by vision physiologists. The discoveries in this kind of work come from the actual building of functional electronic emulations of the nervous system. Despite the simplifications of the underlying biological circuits on this chip, the silicon cells behave and sound like real cells well enough that they can be used in lecture demonstrations and physiology labs. Some neuroscientists question the validity of this work, asking what it teaches about the workings of the brain. The same could have been asked of the early aeronautical engineers: i.e., what would they learn about birds' flight by building a flying machine. By building systems like the Physiologist's Friend, we are learning how to make physical devices that compute more like the brain than do the synchronous logic devices that currently dominate artificial computation.

The construction of this chip achieves three other specific goals. First, the realistic responses of the silicon cells suggest that prostheses can use neuromorphic circuits to emulate neural structures (Mead, 1990). Emulation, in contrast to simulation, can result in more compact realization and greatly reduced power consumption (Sarpeshkar, 1998), at the cost of reduced precision, flexibility, and, possibly, increased development time.

Second, the inclusion of circuits that generate the wide range of required bias currents demonstrates that neuromorphic circuits can be reliably manufactured and used with a minimal number of adjustments or external components.

Third, the use of this device in physiology labs contributes concretely to the three R's of animal welfare: replacement, reduction, and refinement. Our physiologist colleagues in Zürich and at several other labs in North America and Europe use this system in experiment design, student training, and lecture demonstrations. We hope that this chip, as well as its more sophisticated descendents, takes on a role in physiology labs similar to that of voltmeters and screwdrivers.

Resources

Other resources for the Physiologist's Friend are available at www.ini.unizh.ch/~tobi/friend, including a Java computer program modeled after the chip and designed as a learning tool that simulates a small patch of the visual system. Using the mouse and keyboard to control virtual stimuli, users can hear and see the responses of a variety of cells, including mystery cells.

Acknowledgements

M. Carandini indirectly initiated this project and measured some of the chip's responses. E. Chicca, G.M. Ricci, and S. Bovet assembled a preliminary version of the chip together with G. Indiveri, J. Kramer, and the authors. A. van Schaik and O. Landolt contributed to the bias generator design, An anonymous reviewer provided constructive comments. This work was funded by ETH Zürich and the University of Zürich.

References

Anderson, J. S., Carandini, M., & Ferster, D. (2000). Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. *Journal of Neurophysiology, 84*, 909–926.

Boahen, K. A. (1997). The retinomorphic approach: pixel-parallel adaptive amplification, filtering, and quantization. *Journal of Analog Integrated Circuits and Signal Processing, 13*, 53–68.

Borg-Graham, L. J., Monier, C., & Fregnac, Y. (1998). Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature, 393*, 369–373.

Bult, G., & Geelen, G. (1992). An inherently linear and compact MOST-only current division technique. *IEEE Journal of Solid-State Circuits, 27*, 1730–1735.

Delbrück, T. (1991). "Bump" circuits for computing similarity and dissimilarity of analog voltages. In *Proceedings of the International Joint Conference on Neural Networks* (pp. I-475–I-479).

Delbrück, T., & Mead, C. A. (1994). Analog VLSI adaptive logarithmic wide-dynamic-range photoreceptor. In *Proceedings of the IEEE International Symposium on Circuits and Systems, 4* (pp. 339–342).

Delbrück, T., & van Schaik, A. (2004). A. Bias generators with wide dynamic range. To be presented at the *IEEE International Symposium on Circuits and Systems (ISCAS 2004)*.

Dowling, J. E. (1987). *The Retina: An Approachable Part of the Brain*. Cambridge, MA: The Belknap Press of Harvard University Press.

Harmon, L. D. (1961). Studies with artificial neurons, I: properties and functions of an artificial neuron. *Kybernetik, 3*, 89–101.

Hirsch, J. A., Alonso, J. M., Reid, R. C., & Martinez, L. M. (1998). Synaptic integration in striate cortical simple cells. *Journal of Neuroscience, 18*, 9517–9528.

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology, 160*, 106–154.

Liu, S. C., Kramer, J., Indiveri, G., Delbrück, T., Burg, T., & Douglas, R. (2001). Orientation-selective aVLSI spiking neurons. *Neural Networks, 14*, 629–643.

Mahowald, M., & Douglas, R. (1991). A silicon neuron. *Nature, 354*, 515–518.

Mead, C. A. (1989). *Analog VLSI and neural systems*. Reading, MA: Addison-Wesley Publishing.

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

Mead, C. A., & Mahowald, M. A. (1988). A silicon model of early visual processing. *Neural Networks, 1*, 91–97.

Ringach, D. L., Sapiro, G., & Shapley, R. (1997). A subspace reverse-correlation technique for the study of visual neurons. *Vision Research, 37*, 2455–2464.

Sarpeshkar, R. (1998). Analog versus digital: extrapolating from electronics to neurobiology. *Neural Computation, 10*, 1601–1638.

Schweitzer-Tong, D. E. (1983). The photoneuromime: an artificial visual neuron for dynamic testing of computer-controlled experiments. *Behavior Research Methods and Instrumentation, 15*, 9–12.

Vittoz, E., & Fellrath, J. (1977). CMOS analog integrated circuits based on weak inversion operation. *IEEE Journal of Solid-State Circuits, SC-12*, 224–231.