An Analog VLSI Implementation of a Visual Interneuron: Enhanced Sensory Processing through Biophysical Modeling

Reid R. Harrison and Christof Koch
Computation and Neural Systems Program, California Institute of Technology
MS 139-74, Pasadena, CA 91125
[harrison, koch]@klab.caltech.edu

Flies are capable of rapid, coordinated flight through unstructured environments. This flight is guided by visual motion information that is extracted from photoreceptors in a robust manner. One feature of the fly's visual processing that adds to this robustness is the saturation of wide-field motion-sensitive neuron responses with increasing pattern size. This makes the cell's responses less dependent on the sparseness of the optical flow field while retaining motion information. By implementing a compartmental neuronal model in silicon, we add this "gain control" to an existing analog VLSI model of fly vision. This results in enhanced performance in a compact, low-power CMOS motion sensor. Our silicon system also demonstrates that modern, biophysically-detailed models of neural sensory processing systems can be instantiated in VLSI hardware.

1 Introduction

When building a visual motion sensor for the real world, one must cope with the sparseness of optical flow fields. Many parts of the visual field are featureless, and produce no motion information. When estimating self-rotation, for example, one would like to extract information based on wide-field motion that is robust against these gaps.

Flies have developed a remarkably elegant method for dealing with optical flow sparseness. The optic lobe in the brain of the fly contains several wide-field motion-sensitive neurons that integrate motion information from many elementary motion detectors (EMDs) in large receptive fields to produce estimations of self-rotation [4]. These neurons have been studied for decades, and much is known about their response properties. One property exhibited by some of these cells is called gain control, and seems to make the sensory response robust against gaps in the optical flow field.

Gain control describes the saturating response of these motion-sensitive cells with increasing stimulus size (see Figure 1). As the extent of the stimulating pattern across the visual receptive field increases linearly, the response of the cell saturates, but it saturates at different levels for different stimulus velocities. This cannot be explained by a simple saturating output channel. The wide-field motion-sensitive neuron is integrating motion information spatially, but this integration is nonlinear.

This size-dependent saturation assures that at reasonably high levels of stimulation, the cell is not sensitive to gaps in the optic flow field. (Featureless parts of the visual scene decrease the effective stimulus size.) The cell now encodes the stimulus velocity, which in this case may represent a measure of self-rotation, largely independent of the visual sparseness of the environment. We investigated the biological machinery underlying this sensory processing, and translated it to silicon in the form of a compact, low-power neuromorphic analog VLSI system.

2 Algorithm and Biological Architecture

The neural architecture and biophysical mechanisms underlying gain control in the fly are now understood [2], [6]. Simple linear models of spatial integration result in an output that is linearly dependent on stimulus size (see Figure 2a). Size-dependent saturation comes about if we use a more accurate model of the wide-field motion-sensitive neuron (see Figure 2b). In this model, the EMD outputs are not directly conveyed to the wide-field neuron. Instead, the EMDs modulate synapses, which are modeled as conductances between the intracellular potential and a fixed ion reversal potential. Depending on the type of ion involved, the reversal potential can be

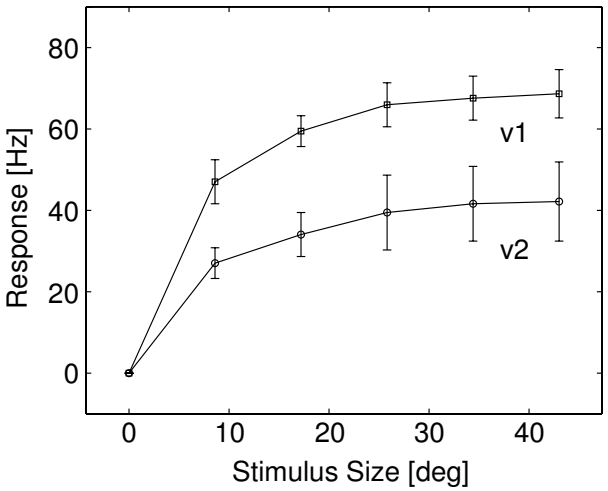

Figure 1. Gain Control in a Wide-Field Motion-Sensitive Neuron in the Fly. The cell response saturates with increasing pattern size, and saturates at different levels depending on the stimulus velocity ($v1 = 72°/\text{sec}$; $v2 = 360°/\text{sec}$). Stimulus was a sinusoidal grating with spatial wavelength of 24° and 29% contrast. Data represent mean ± SEM of extracellular recordings from the lobula plate spiking neuron H1 of four different female blowflies (*Calliphora erythrocephala*). Data reprinted from [6].

above or below the resting potential of the cell, creating excitatory or inhibitory synapses. We connect preferred-direction EMDs to excitatory synapses and null-direction EMDs to inhibitory synapses. Each EMD modulates its corresponding ion channel conductance, which acts to pull the cell away from its resting potential, where it is held by the fixed leakage conductance g_{rest}.

If we measure all voltages relative the the cell resting potential V_{rest}, then the cell potential can be expressed as:

$$V_{cell} = \frac{g_{exc}V_{exc} + g_{inh}V_{inh}}{g_{exc} + g_{inh} + g_{rest}} \qquad (1)$$

Here, V_{exc} and V_{inh} represent the ion reversal potentials, and g_{exc} and g_{inh} represent the ion channel conductances, which are controlled by the preferred-direction and null-direction EMD outputs. Thus V_{cell} saturates at V_{exc} for $g_{exc} \gg g_{inh}, g_{rest}$, and saturates at V_{inh} for $g_{inh} \gg g_{exc}, g_{rest}$.

To generate gain control, this method of spatial integration relies on the particular architecture of the EMD. The Reichardt-type correlation-based EMD exhibits only weak direction selectivity and produces a significant but weaker response for stimuli in its null direction. Strong directionally selectivity is achieved only after the EMD opponent pair signals are sub-

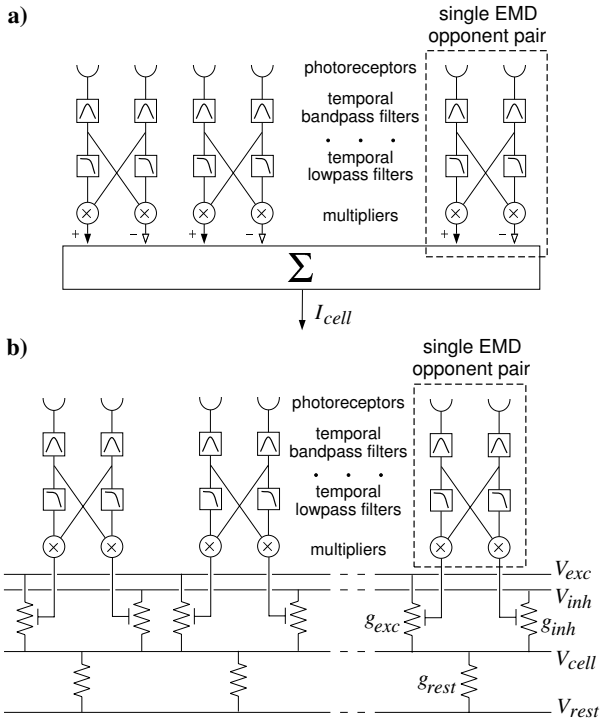

Figure 2. Models of Wide-Field Motion-Sensitive Neurons in the Fly. a) Linear spatial integration. An array of Reichardt-type delay-and-correlate elementary motion detector (EMD) opponent pairs are subtracted locally to achieve strong direction selectivity. The spatially distributed outputs are summed. b) Nonlinear spatial integration. EMD opponent pairs modulate excitatory and inhibitory synaptic conductances of the wide-field neuron. EMD activity can pull the neuron away from its resting potential V_{rest}.

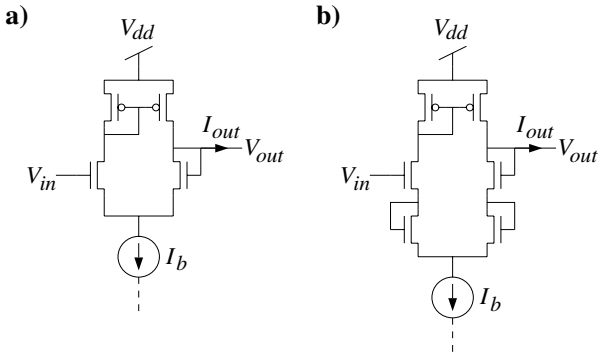

Figure 3. Variable Conductance Circuits. a) Four-transistor follower. The output node is driven towards the potential at V_{in}. The output conductance of this circuit is proportional to the bias current I_b. b) Six-transistor follower. Same behavior as the four-transistor circuit, but source-degeneration diode-connected transistors double the circuit's linear range.

tracted [1]. Importantly, it has been demonstrated that the degree of directional selectivity is itself a function of stimulus velocity [2]. Thus, a particular stimulus velocity will generate a characteristic ratio of preferred-direction to null-direction EMD output. This in turn will lead to a characteristic voltage that the cell is driven towards. We can rewrite Equation 1 as:

$$V_{cell} = V_{stimulus}\left(\frac{g_{stimulus}}{g_{stimulus} + g_{rest}}\right) \tag{2}$$

Where:

$$V_{stimulus} = \frac{g_{exc}V_{exc} + g_{inh}V_{inh}}{g_{exc} + g_{inh}} \tag{3}$$

$$g_{stimulus} = g_{exc} + g_{inh} \tag{4}$$

Since conductances sum in parallel, a larger pattern size will cause more EMD opponent pairs to be active, producing a larger $g_{stimulus}$ and driving V_{cell} away from V_{rest} and towards the $V_{stimulus}$ that is characteristic of the stimulus velocity.

3 Silicon Implementation

We have previously developed an analog VLSI circuit that models the Reichardt-type EMD shown in Figure 2 [3]. These delay-and-correlate motion detectors exhibit spatiotemporal and directional responses similar to those studied in the fly. The EMD circuits produce current-mode outputs, which allowed us easily to sum the outputs of an EMD array simply by tying the wires together. But as we have seen, this linear integration across a receptive field does not produce a saturating size response. To produce this response, we developed a silicon implementation of the biophysical model shown in Figure 2b.

The silicon implementation of excitatory and inhibitory synaptic conductances consumes little silicon area or power. To achieve a variable conductance, we use a differential transconductance amplifier in a follower configuration (i.e., the output directly connected to the negative input) so that the output node is driven toward the input voltage (see Figure 3a) [5]. As long as the follower is operated in the subthreshold regime ($I_b < 100$ nA) and in the linear region of operation (the input and the output voltages differ by less than 75 mV) the circuit has an output conductance g_{out}, which is given by:

$$g_{out} = \frac{\partial I_{out}}{\partial V_{out}} = \frac{q\kappa}{2kT}I_b \tag{5}$$

Where κ is the back-gate coefficient, a CMOS process parameter that typically has a value close to 0.7. If we assume constant temperature, the circuit produces an output conductance proportional to the bias current I_b, approximately $g = (14\ \mathrm{V}^{-1})I_b$ at room temperature. By adding two additional diode-connected transistors in a source-degeneration configuration [7], we extend the linear range of this circuit from ± 75 mV to ± 150 mV (see Figure 3b).

We use this six-transistor circuit to model a population of ion channels whose ion reversal potential is specified by the input voltage. Thus each compartment of our cell has three ion channel populations modeled by three follower circuits. One channel has a constant conductance g_{rest} and sets the resting potential of the cell, V_{rest}. The other two model the excitatory and inhibitory inputs from the local opponent pair of EMDs. These conductances, g_{exc} and g_{inh}, are controlled by the preferred- and null-direction EMDs and have reversal potentials above (V_{exc}) and below (V_{inh}) the resting potential, respectively. The current-mode output of each EMD directly biases a variable conductance circuit.

4 Experiments

We fabricated a one-dimensional array of 28 compartments in a standard 1.2-μm CMOS VLSI process. Each compartment contained a preferred- and null-direction EMD and the corresponding variable conductance circuits described above. The chip included integrated photoreceptors, so imaging and computation were performed on a single die. We mounted a standard 4 mm CCTV camera lens over the chip, which gave the photoreceptors an angular spacing of

0.73°. We presented high-contrast drifting sinusoidal stimuli to the chip with a computer monitor (Sony Multiscan 17seII). We were able to achieve frame rates of 70 Hz, and screen resolution far exceeding the photoreceptor spacing.

We varied the extent of our stimulus across the chip's field of view. (All sizes are expressed as the fraction of the photoreceptor array that was stimulated.) First we investigated the behavior of the linear spatial summation architecture, achieved simply by tying together the output currents of all the EMDs in the array (see Figure 4a). The output increases linearly with pattern size. Next we investigated nonlinear spatial integration by directing the EMD output currents to the variable conductance circuits. We measured the voltage of our 28-compartment cell model as a function of pattern size (see Figure 4b). The cell response saturates with increasing pattern size. We also see that it saturates at different levels for different velocities.

We can control the nature of this saturation by increasing or decreasing g_{rest}, the conductance that tries to hold the cell at the resting potential. We investigated pattern size saturation for three different values of g_{rest} (see Figure 5). As g_{rest} decreases, it becomes easier to pull the cell away from V_{rest} and only a few EMD opponent pairs need contribute to dominate the response. Rapid saturation decreases pattern size dependence but increases the sensitivity of the chip to small-field motion that might not be produced by self-motion.

This method for achieving gain control is extremely power efficient. Since the excitatory and inhibitory synapses are powered by the current-mode output of the existing EMD circuitry, the only power cost we incur for adding gain control is the power consumed by the leakage conductance g_{rest}, which is simply $V_{dd}I_b$. For the experiments shown in Figure 5, this power varied between 14 pW and 5.4 nW per compartment, depending on the value of g_{rest}. This is negligible when compared to the $\sim$10 μW consumed by the EMD circuitry.

5 Conclusions

We developed and implemented a neuromorphic visual motion sensor with improved robustness against the sparseness commonly found in real-world optical flow fields. This robustness is exhibited in a response which saturates with increasing pattern size, yet retains information on the stimulus velocity. This gain control is achieved by adapting to silicon a more biophysically-realistic model of the underlying neu-

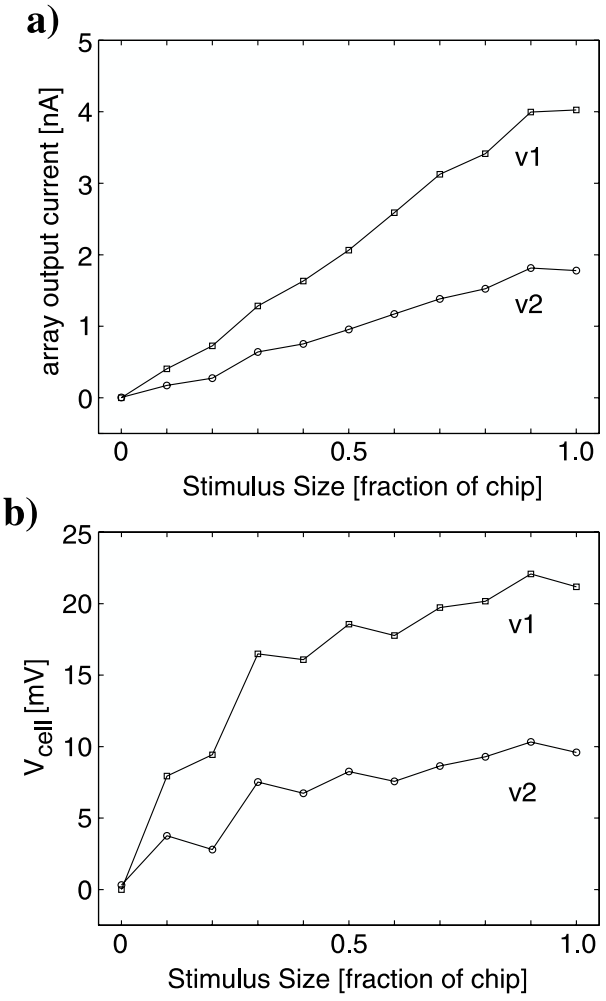

Figure 4. Gain Control in the Silicon System. a) Linear spatial integration. The outputs from the EMD array are summed; gain control is not observed. b) Nonlinear spatial integration with synaptic conductances. As is the biological system, the chip response saturates with increasing pattern size, and saturates at different levels depending on the stimulus velocity ($v1$ = 8.8°/sec; $v2$ = 4.4°/sec). The ion reversal potentials were V_{exc} = +150 mV; V_{inh} = −80 mV relative to V_{rest}. The asymmetry was necessary to counteract observed transistor mismatch. Stimulus was a sinusoidal grating with spatial wavelength of 2.9° and >99% contrast. Data represent mean output during stimulus presentation.

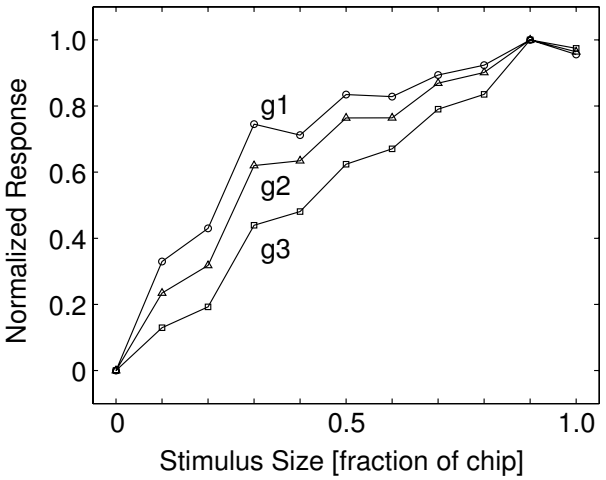

Figure 5. Varying leakage conductance. The degree of saturation can be tuned by changing the leakage conductance. Data are normalized for each value of g_{rest} ($g1 = (26\ \mathrm{G}\Omega)^{-1}; g2 = (1.0\ \mathrm{G}\Omega)^{-1}; g3 = (0.10\ \mathrm{G}\Omega)^{-1}$).

ral computation. We built a compartmental model of a neuron in hardware, and used the properties this system to improve our sensory processing. This represents a significant advance towards building small, low-power sensors that extract useful information from the real world in a robust manner.

The addition of this circuit to our existing EMD circuits (see [3]) advances our analog VLSI model of the fly's visual system. While analog hardware models lack the high precision of software simulations, they are physically instantiated and compute in real time. Thus they can be tested and evaluated in real-world situations. These advantages may prove important as neural models of sensorimotor systems grow more complex.

Acknowledgements

This work was supported by the Center for Neuromorphic Systems Engineering at the California Institute of Technology as a part of NSF's Engineering Research Center program, by ONR, and by DARPA. We thank Alexander Borst for invaluable discussions.

References

[1] A. Borst and M. Egelhaaf. Direction selectivity of blowfly motion-sensitive neurons is computed in a two-stage process. *Proceedings of the National Academy of Sciences USA*, 87:9363–9367, 1990.

[2] A. Borst, M. Egelhaaf, and J. Haag. Mechanisms of dendritic integration underlying gain control in fly motion-sensitive interneurons. *Journal of Computational Neuroscience*, 2:5–18, 1995.

[3] R. R. Harrison and C. Koch. An analog vlsi model of the fly elementary motion detector. In M. I. Jordan, M. J. Kearns, and S. A. Solla, editors, *Advances in Neural Information Processing Systems 10*, pages 880–886. MIT Press, 1998.

[4] H. G. Krapp and R. Hengstenberg. Estimation of self-motion by optic flow processing in single visual interneurons. *Nature*, 384:463–466, 1996.

[5] C. Mead. *Analog VLSI and Neural Systems*. Addison-Wesley, 1989.

[6] S. Single, J. Haag, and A. Borst. Dendritic computation of direction selectivity and gain control in visual interneurons. *Journal of Neuroscience*, 17:6023–6030, 1997.

[7] L. Watts, D. A. Kerns, R. F. Lyon, and C. A. Mead. Improved implementation of the silicon cochlea. *IEEE Journal of Solid-State Circuits*, 27:692–700, 1992.