



A Robust Analog VLSI Reichardt Motion Sensor

REID R. HARRISON AND CHRISTOF KOCH

Electrical Engineering Department, University of Utah, Salt Lake City, UT 84112, USA
E-mail: harrison@ee.utah.edu

Received June 10, 1999; Revised February 14, 2000

Abstract. Silicon imagers with integrated motion-detection circuitry have been developed and tested for the past 15 years. Many previous circuits estimate motion by identifying and tracking spatial or temporal features. These approaches are prone to failure at low SNR conditions, where feature detection becomes unreliable. An alternate approach to motion detection is an intensity-based spatiotemporal correlation algorithm, such as the one proposed by Hassenstein and Reichardt in 1956 to explain aspects of insect vision. We implemented a Reichardt motion sensor with integrated photodetectors in a standard CMOS process. Our circuit operates at sub-microwatt power levels, the lowest reported for any motion sensor. We measure the effects of device mismatch on these parallel, analog circuits to show they are suitable for constructing 2-D VLSI arrays. Traditional correlation-based sensors suffer from strong contrast dependence. We introduce a circuit architecture that lessens this dependence. We also demonstrate robust performance of our sensor to complex stimuli in the presence of spatial and temporal noise.

Key Words: Reichardt motion detector, analog VLSI, insect vision, motion sensor, robust sensing, biological model

I. Introduction

Sensors are devices that interface the internal world of abstract computation with the external, physical world. At mesoscopic scales, the real world is not discrete so sensors are inherently analog. Sensors often perform local computations in order to extract relevant information from noisy signals. These computations are often performed in analog to save power and to reduce the computational burden on the main processor. The physical devices used to fabricate these “intelligent” analog sensors are subject to built-in inaccuracies, as no two devices will be identical. To be useful outside the laboratory, integrated silicon sensors must thus exhibit two types of robustness: *internal robustness*, that is robustness in the face of substrate and component variability both within and between integrated circuits; and *external robustness*, where sensors respond correctly to the type of complex and noisy scenes extant in the real world.

Much can be learned about real-time sensory processing by investigating biological systems. While

some animals possess larger brains and more sophisticated capabilities than others, *all* animals show complex behaviors and very robust performance in the real world of the type neither understood by scientists nor yet replicated by engineers [1]. The seminal work of Mead and colleagues [2] demonstrates that principles of biological information processing can be implemented in analog VLSI circuits. We continue this line of research to develop a robust visual motion sensor based on the insect visual system.

In this paper, we present a novel implementation of a correlation-based motion sensor found in flying insects. Our design incorporates recent ideas from biological models to reduce the contrast and spatial frequency dependence usually associated with sensors of this type. We systematically test our sensor for robustness by measuring its performance with noisy, naturalistic stimuli.

In Section II, we review motion detection algorithms and corresponding hardware implementations from the past two decades. In Section III, we analyze the biologically-inspired Reichardt motion detection

algorithm. Section IV describes the circuit realization of our sensor, and Section V presents the results of experiments that evaluate the sensor's performance.

II. Relation to Previous Work

During the past 15 years, many analog or hybrid VLSI motion sensors have been developed and tested. Most of these designs incorporate photodetection and motion computation on the same chip. These focal-plane processors typically cannot achieve the high pixel density of dedicated CMOS imagers or CCDs, but rather trade off density for functionality. By extracting motion information at the level of light detection instead of using an external microprocessor, large savings in size, power, and system complexity are achieved.

Nearly every motion detection algorithm devised has been implemented in VLSI in some form. Motion detection algorithms can be divided into two broad classes: feature-tracking or token-based algorithms, and intensity-based algorithms. We will review previous hardware motion sensors of both types.

A. Feature-Tracking or Token-Based Algorithms

Algorithms of this type use feature detectors to identify salient points in the raw image. Binary tokens indicating the absence or presence of a feature are then passed on to a velocity-estimation stage. Two types of feature detectors have been used in silicon motion sensors: spatial feature detectors and temporal feature detectors.

Spatial feature detectors typically look for intensity edges—places where the image changes rapidly over a short distance. These edges are identified in successive discrete-time frames. By matching a feature from time t to the same feature at time $t + \Delta t$, velocity can be measured. Of course, if the feature density is high, the correspondence problem (deciding which edge at time t corresponds to which edge at time $t + \Delta t$) may be difficult, and if the feature density is low then the optic flow field will be sparse. This technique is commonly used in software motion detection schemes due to its compatibility with discrete-time computation, but it has also been used in hardware motion detectors as a front-end before additional continuous-time processing [3,4].

Temporal feature detectors typically look for rapid changes in the image brightness at each pixel—temporal edges. Due to the local nature of this computation, it has been quite popular in analog hardware approaches and has been implemented efficiently in continuous-time circuits [5–8]. These chips measure time-of-travel: the time it takes for an edge to pass from one pixel to an adjacent pixel. They offer the advantage of measuring true image speed over many orders of magnitude, and can operate at contrasts as low as 0.15 as long as the image contains sharp edges [7].

Feature-tracking algorithms—especially those employing relatively simple feature detectors, as VLSI implementations must—may yield spurious responses to weak signals. In hardware motion detectors, features are typically encoded as binary entities which are either present or absent. Weak signals produce features near the threshold of detectability. Physically instantiated feature detectors have thresholds which are not perfectly matched across an array, so a weak signal may trigger one detector but not its neighbor. This makes the correspondence problem difficult and may yield inaccurate results [9].

B. Intensity-Based Algorithms

Intensity-based algorithms use some linearly-filtered version of the image brightness, such as the photocurrent, for motion processing without prior binarization or classification. These algorithms can be subdivided into two types: gradient and correlation algorithms. Gradient algorithms use spatial and temporal derivatives of an image to measure velocity under the assumption that the overall intensity of the image remains constant. As long as this assumption holds, the gradient algorithm measures speed directly and independent of contrast. The use of derivatives makes this algorithm vulnerable to high-frequency noise. However, gradient-based hardware motion sensors have been developed [10,11].

The other class of intensity-based motion detectors measure spatiotemporal correlations caused by moving objects or self-motion. These algorithms include the spatiotemporal motion energy model of Adelson and Bergen [12] and the Reichardt motion detector, first proposed by Hassenstein and Reichardt in 1956 as a model of motion detection in insects

[13]. Over the past four decades, the Reichardt model has been used to explain motion detection in many animals from houseflies to humans [14]. A related algorithm was proposed by Barlow and Levick to explain direction-selective cells in the rabbit's retina [15]. This algorithm has been implemented in hardware using temporal feature detectors [16] and in a purely intensity-based architecture [17]. Template matching schemes can also be considered correlation-based algorithms, but these methods are not generally amenable to analog VLSI implementation.

Delbrück built a continuous-time, continuous-value correlation-based hardware motion detector [18] that used delay lines and quadratic nonlinearities to compute a measure of spatiotemporal motion energy. There was no feature detection; rather, the raw output from the photoreceptors were used to compute motion. Due to the nature of the delay lines, an object moving across the chip's field of view caused a directional response to gradually build over space and time. The circuit exhibited a velocity-tuned response. Like many correlation-based algorithms, the response of this chip was highly contrast-dependent.

Early attempts to implement the intensity-based Reichardt architecture in silicon used translinear, current-mode circuits [19,20]. As we will show in Section III, the response of these traditional Reichardt motion sensors is affected strongly by contrast. Attempting to build contrast-independent Reichardt motion sensors, some have designed circuits that perform an initial binarization of the image based on temporal edges and then delay and correlate these digital signals [21,22]. These circuits would not be expected to perform well in noisy, low-contrast environments without additional image preprocessing. Reichardt-inspired sensors have also been built in discrete hardware and used on mobile robots, although the particular implementation more closely resembled a feature-tracking, time-of-travel scheme [23,24].

C. Testing Methodology

Most of the sensors above were tested with noiseless, highly artificial stimuli—sharp intensity edges and gratings. Some were tested with low-contrast stimuli, but noise was never intentionally introduced into the

image. There is shot noise in photoreceptor signals, but this is usually small compared to the signal, and is rarely quantified. (A theoretical noise analysis is shown in [6], but noise is introduced into experiments only by lowering the overall light intensity.)

Also, few of these circuits were explicitly tested for interpixel matching characteristics. Different circuit architectures and layouts can be more or less sensitive to stochastic device variation inherent in VLSI processes. Only two studies quantified circuit variation [3,4], while a few others graphically depicted the output of motion detector arrays [8,11,17,18,21].

III. Theoretical Analysis

We now describe the Reichardt motion detection algorithm which underlies our hardware implementation. The Reichardt motion detector uses a correlation-based algorithm, whose output is equivalent to the output of the spatiotemporal motion energy model proposed by Adelson and Bergen [12] and others [25,26].

The basic idea of the Reichardt motion detector is to correlate the signal from one photoreceptor with the delayed signal from an adjacent photoreceptor. This delay-and-correlate algorithm produces a velocity-tuned response that is weakly directionally selective. By subtracting the responses of two opponent half-detectors from each other, only the directional motion signal remains (see Fig. 1(a)) [27].

It is instructive to consider the case where the stimulus is a sinusoidal grating moving at velocity v . Image intensity $i(x, t)$ can be expressed as

$$i(x, t) = I + \Delta I \sin[2\pi f_s(x + vt)] \quad (1)$$

where I is the mean intensity, and f_s is the spatial frequency. The contrast of the grating is $\Delta I/I$. At any single photoreceptor, this moving grating produces a temporal sinusoidal signal with a frequency $f_t = vf_s$. This allows us to rewrite (1) as

$$i(x, t) = I + \Delta I \sin(\omega_t t + \omega_s x) \quad (2)$$

where $\omega_t = 2\pi f_t$ and $\omega_s = 2\pi f_s$. If two photoreceptors have an angular separation of ϕ , then the signals measured by the photoreceptors can be expressed as

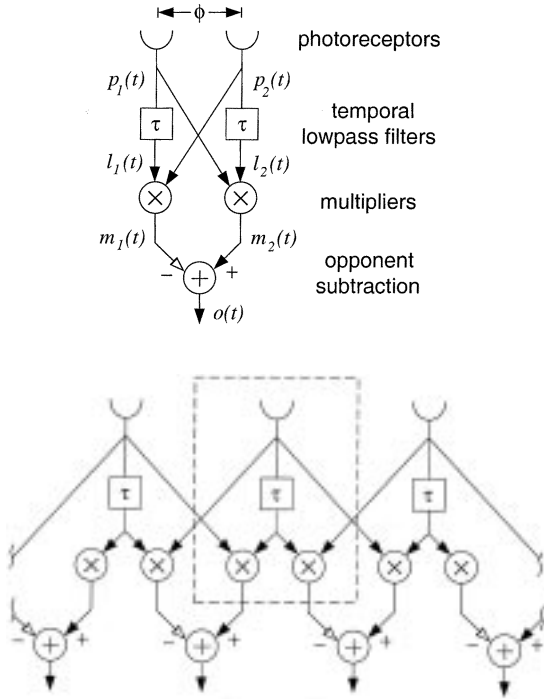


Fig. 1. Reichardt motion detector architecture. (a) The signal from one photoreceptor is correlated with the delayed signal from an adjacent photoreceptor. Direction selectivity is increased by subtracting the responses of two half-detectors in opponency. (b) A 1-D array of Reichardt motion detectors, illustrating the repeated computational element. Subtraction currently performed off-chip for flexibility, but would be easy to implement given the current-mode outputs of the multipliers. A 2-D arrangement is possible with two additional multipliers in each cell.

$$p_1(t) = |H(\omega_t)|\Delta I \sin\left(\omega_t t - \omega_s \frac{\phi}{2}\right) \quad (3)$$

$$p_2(t) = |H(\omega_t)|\Delta I \sin\left(\omega_t t + \omega_s \frac{\phi}{2}\right) \quad (4)$$

We introduce $H(\omega_t)$ as the temporal frequency response of the photoreceptors. For simplicity we ignore the phase contribution of $H(\omega_t)$ as it will be identical in $p_1(t)$ and $p_2(t)$, and thus have no effect on perceived motion. We also assume that the photoreceptors have a highpass behavior which eliminates the DC component of illumination I . We model the photoreceptor response as

$$\frac{H(\omega_t)}{(j\omega_t\tau_{HPF} + 1)(j\omega_t\tau_{photo} + 1)} \quad (5)$$

where τ_{HPF} is the time constant of the DC-blocking highpass filter, τ_{photo} is the time constant defining the photoreceptor bandwidth, and K is a constant of proportionality.

Since we are building continuous-time circuits, a delay is most easily approximated using the phase lag inherent in a first-order lowpass filter. Lowpass filtering each photoreceptor signal yields

$$l_1(t) = \frac{|H(\omega_t)|\Delta I}{\sqrt{\tau^2\omega_t^2 + 1}} \sin\left(\omega_t t - \omega_s \frac{\phi}{2} - \tan^{-1} \tau\omega_t\right) \quad (6)$$

$$l_2(t) = \frac{|H(\omega_t)|\Delta I}{\sqrt{\tau^2\omega_t^2 + 1}} \sin\left(\omega_t t + \omega_s \frac{\phi}{2} - \tan^{-1} \tau\omega_t\right) \quad (7)$$

Correlation is accomplished by multiplying the phase lagged signals with adjacent, non-delayed signals. The results are two ‘‘half-detector’’ responses:

$$m_1(t) = G[\cos(\omega_s\phi + P) - \cos(2\omega_t t - P)] \quad (8)$$

$$m_2(t) = G[\cos(\omega_s\phi - P) - \cos(2\omega_t t - P)] \quad (9)$$

where

$$G = \frac{(|H(\omega_t)|\Delta I)^2}{2\sqrt{\tau^2\omega_t^2 + 1}} \quad (10)$$

$$P = \tan^{-1} \tau\omega_t \quad (11)$$

Once these signals are subtracted in opponency, the final output becomes

$$o(t) = (\Delta I)^2 |H(\omega_t)|^2 \frac{\tau\omega_t}{\tau^2\omega_t^2 + 1} \sin\phi\omega_s \quad (12)$$

This describes the sensitivity of a Reichardt motion detector to a sinusoidal grating with a particular contrast, temporal frequency, and spatial frequency. Notice that the response is a separable function of these three parameters. We can rewrite this equation to make the dependency on the grating velocity v explicit:

$$o(t) = (\Delta I)^2 |H(\omega_s v)|^2 \frac{\tau\omega_s v}{\tau^2\omega_s^2 v^2 + 1} \sin\phi\omega_s \quad (13)$$

Although this response is direction selective (i.e., the sign of $o(t)$ is equal to the sign of v), it does not encode velocity independent of spatial frequency and contrast.

There are no time-dependent terms in this equation. This indicates a DC response to moving patterns. However, if the mean intensity of the image

is not completely removed by prefiltering, or if the opponent subtraction is not perfectly balanced, oscillations at harmonics of the stimulus temporal frequency will be superimposed on the DC response. (For a complete analysis of the Reichardt motion detector in these nonideal cases, see Egelhaaf [28].) These oscillations may be reduced by using an array of Reichardt motion detectors (see Fig. 1(b)) and summing their responses. This has the effect of integrating over different phases of the stimulus and canceling pattern-dependent oscillations [29,30].

IV. Circuit Realization

We chose to implement a fully analog, continuous-time, intensity-based version of the Reichardt motion detector with reduced contrast dependence. To the best of our knowledge, this is the closest approximation to this biological motion sensor that has been built. We are testing the hypothesis that by emulating an evolved neurobiological system, we can develop a robust analog VLSI sensor capable of being used in noisy environments.

A. Circuit Architecture

Before computing any information about motion, we must first measure light intensity at each pixel. We use an adaptive photoreceptor circuit developed by

Delbrück and Mead [31]. This is a four-transistor circuit that uses a substrate photodiode and source follower (M_1) to convert incident light into a logarithmically encoded voltage (see Fig. 2(a)). A high gain amplifier (M_2 and M_3) and feedback network (C_1 and C_2) amplify the voltage signal by a factor of 18. The adaptive element (M_4) acts as a nonlinear feedback element that conducts only if the voltage across it exceeds several hundred millivolts. This allows the photoreceptor to adapt to large changes in illumination. Thus we maintain a large dynamic range over a wide operating range. At low bias current levels, the bandwidth of the photoreceptor is limited by the parasitic output capacitance C_p . For a detailed discussion of this circuit, see Delbrück and Mead [31].

The adaptive photoreceptor signal is sent to a $g_m C$ highpass filter (see Fig. 2(a)). We use a source follower to provide a low-impedance driver, but in future designs we will leave this out and compensate for the increased output capacitance by increasing the photoreceptor bias current I_{pr} . We use a highpass filter for two reasons. First, the ac coupling eliminates any systematic offsets caused by device variation in the adaptive photoreceptor. Second, by fixing the DC component of the signal to V_a , we can eliminate any common-mode effects later in the circuit.

The delay is accomplished with a first-order $g_m C$ lowpass filter (see Fig. 2(b)). The bias transistor in the circuit was made several times minimum size to improve time constant matching across the chip. By

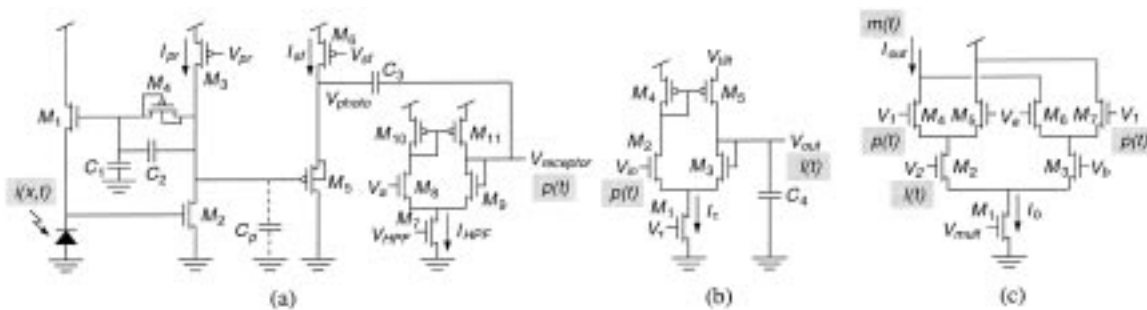


Fig. 2. Circuit schematic. Shaded labels indicate corresponding signals from Fig. 1(a). (a) Adaptive photoreceptor (M_1 – M_4 , C_1 – C_2) with source follower (M_5 – M_6) and temporal highpass $g_m C$ filter (M_7 – M_{11} , C_3) to remove the DC component of V_{photo} . (b) Temporal lowpass $g_m C$ filter. This circuit's phase lag acts as a delay. (c) Gilbert multiplier. This circuit multiplies delayed and nondelayed photoreceptor signals.

operating this circuit at low current levels, we can achieve time constants useful for motion detection (10–100 ms) with reasonably sized capacitors (on the order of 1 pF).

Correlation is approximated by a Gilbert multiplier (see Fig. 2(c)). The input V_2 comes from the lowpass filter, and V_1 comes from the highpass filtered photoreceptor from an adjacent pixel (see Fig. 2(c)). The voltage V_a is the reference voltage used by the highpass filter, and V_b is another DC bias voltage set a few tens of millivolts below V_a . We operate these field-effect transistors (FETs) in subthreshold, where their drain current I_d , ignoring channel-length modulation effects, is given by

$$I_d = I_0 e^{\kappa V_g / V_T} (e^{-V_s / V_T} - e^{-V_d / V_T}) \quad (14)$$

where I_0 is a process-dependent constant, V_g , V_s and V_d are the gate, source, and drain voltages referenced to the bulk potential, κ is the gate efficiency factor (typically around 0.7), and V_T is the thermal voltage kT/q [2]. Sub-threshold FETs exhibit exponential behavior, much like the BJTs with which the Gilbert multiplier was originally built. We take a single-ended current-mode output from the circuit, which gives us

$$I_{out} = \frac{I_b}{2} + \frac{I_b}{2} \tanh \left[\frac{\kappa(V_1 - V_a)}{2V_T} \right] \tanh \left[\frac{\kappa(V_2 - V_b)}{2V_T} \right] \quad (15)$$

where I_b is the bias current. For small-signal inputs, this can be approximated as

$$I_{out} = \frac{I_b}{2} + \frac{\kappa^2 I_b}{8V_T^2} (V_1 - V_a)(V_2 - V_b) \quad (16)$$

An older version of this circuit used a pFET mirror to eliminate the DC component $I_b/2$ and double the signal amplitude. Device mismatch in the mirror introduced additional offsets, so it was not used in this design. We could in principle use both single-ended outputs, but that would require more wires routed through each pixel, which would consume more area.

For the multiplier to work properly, the common-mode voltage of the lower inputs (V_2 and V_b) must be lower than the common-mode voltage of the upper inputs (V_1 and V_a). Simulation results show that acceptable behavior is obtained with a difference of only 50 mV. In order to lower the DC level of the

lowpass filter output, we lowered the source voltage of the output FET in the current mirror of the $g_m C$ filter (see Fig. 2(b)). By placing the V_{ilt} bias a few tens of millivolts below V_{dd} , we lower the DC output level by $(V_{dd} - V_{ilt})/\kappa$. This difference in source voltages increases the time constant of the lowpass filter, but we can compensate by raising I_τ . The difference in source voltages also creates an asymmetry in the up-going and down-going slew rates of the filter, but in practice this does not seem to have a significant effect on the overall circuit performance.

It can be shown from (15) that the circuit output saturates for differential inputs greater than about $4V_T \approx 100$ mV. Rather than restrict our signals to this small linear region, we exploit the nonlinear behavior of the circuit to improve our motion detection algorithm. It has been shown that by adding saturating nonlinearities before the correlation stage, the contrast dependence of a Reichardt detector can be reduced [32]. To understand this effect, consider (15) in the extreme case where both differential inputs are much greater than V_T :

$$\begin{aligned} I_{out} &\approx \frac{I_b}{2} + \frac{\kappa^2 I_b}{8V_T^2} \text{sgn}(V_1 - V_a) \text{sgn}(V_2 - V_b) \\ &\approx \frac{I_b}{2} + \frac{\kappa^2 I_b}{8V_T^2} \text{sgn}[(V_1 - V_a)(V_2 - V_b)] \end{aligned} \quad (17)$$

where the sign function $\text{sgn}(x)$ is defined as

$$\text{sgn}(x) = \begin{cases} 1, & x \geq 0 \\ -1, & x < 0 \end{cases} \quad (18)$$

Incorporating this ‘‘saturated multiplier’’ model into our analysis from Section III, we can rewrite (8) and (9) as

$$m_1(t) = \text{sgn}[\cos(\omega_s \phi + P) - \cos(2\omega_i t - P)] \quad (19)$$

$$m_2(t) = \text{sgn}[\cos(\omega_s \phi - P) - \cos(2\omega_i t - P)] \quad (20)$$

where P is given by (11). We normalize for the constant prefactor in (17) and neglect the constant DC component since the opponency subtraction will cancel this current.

We can continue the analysis of an EMD array with nonlinear multipliers by using the time averages of $m_1(t)$ and $m_2(t)$, which are given by

$$\langle m_1(t) \rangle = 1 - \frac{2}{\pi} (\omega_s \phi + \tan^{-1} \tau \omega_t) \quad (21)$$

$$\langle m_2(t) \rangle = 1 - \frac{2}{\pi} (\omega_s \phi - \tan^{-1} \tau \omega_t) \quad (22)$$

When these signals are subtracted in opponency, the time average of the output is given by

$$\begin{aligned} \langle o(t) \rangle &= \frac{4}{\pi} \tan^{-1} \tau \omega_t \\ &= \frac{4}{\pi} \tan^{-1} \tau v \omega_s \end{aligned}$$

Spatial integration across a small array of motion detectors will have the effect of integrating over different phases of the stimulus. This will remove the time-dependent components of the motion detector output.

We see in (23) that in the limit of full multiplier saturation the output is sensitive to pattern velocity and spatial frequency but is independent of pattern contrast. We use the inherent saturation in the Gilbert multiplier to achieve this behavior without adding additional hardware. We shall demonstrate this reduced contrast dependence in Section V.

Fig. 3 shows the layout for one 1-D motion sensor, corresponding to the circuit element outlined in Fig. 1(b). All experimental results shown below were measured from arrays of this circuit, which was fabricated in a $1.2\ \mu\text{m}$ double-poly, double-metal n -well CMOS process, yielding a pixel size of $61\ \mu\text{m} \times 199\ \mu\text{m}$ with 32 transistors and a four capacitors totaling 3.0 pF. In order to build a 2-D motion sensor, we need add only two more multiplier circuits and additional interpixel and output wiring. Only two wires in each direction are required for nearest-neighbor communication, making 2-D layout practical. An additional interpixel wire may be required if opponent subtraction is performed locally. A 2-D version of this circuit is currently being fabricated in the same process, and the pixel size is $138\ \mu\text{m} \times 138\ \mu\text{m}$ —a 57% increase in area.

B. Supply Voltage Requirements

By computing motion in parallel, we do not need time constants less than a millisecond at any pixel. The fastest known visual systems (those of houseflies) have bandwidths of less than 200 Hz, and humans can barely perceive the flicker of a 60 Hz monitor. This low bandwidth requirement allowed us

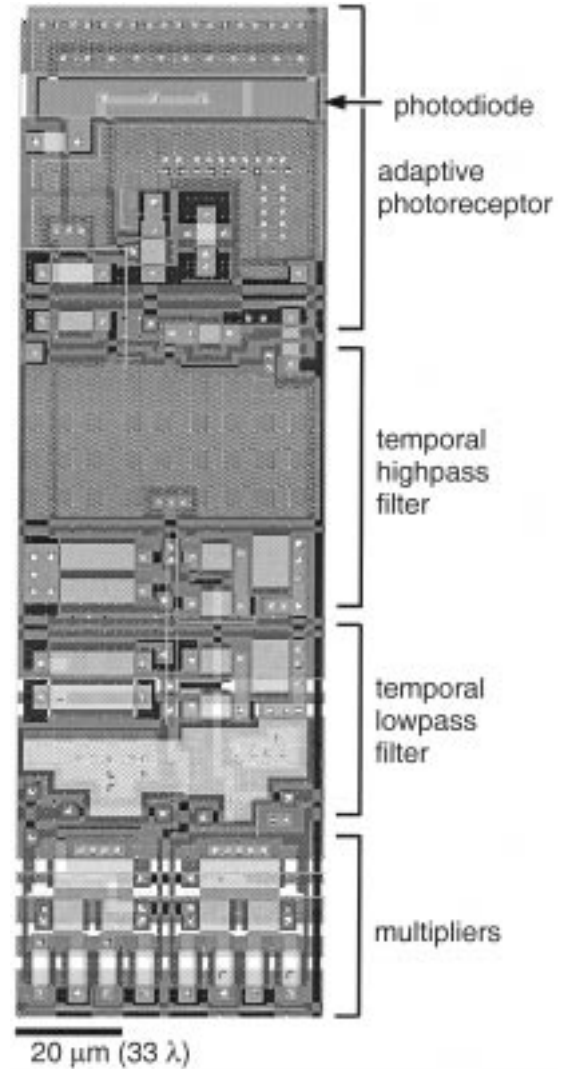


Fig. 3. Motion detector layout. Cell measures $61\ \mu\text{m} \times 199\ \mu\text{m}$ in a standard $1.2\ \mu\text{m}$ process with 32 transistors and 4 capacitors totaling 3.0 pF. In order to build a 2-D motion sensor, we need only add two more multiplier circuits, and some additional interpixel wiring.

to operate the entire circuit in subthreshold (drain currents typically less than $1\ \mu\text{A}$). Subthreshold operation allowed us to operate at $V_{dd} = 2.5\ \text{V}$ despite the Gilbert multiplier circuit, where three transistors are in series between the power supply and ground.

In order to operate an above-threshold transistor in the saturation region (where the drain voltage has no

effect on the drain current, neglecting channel-length modulation), the drain-to-source voltage V_{ds} must exceed $V_{gs} - V_t$, where V_t is the threshold voltage. In subthreshold operation, the condition for saturation is $V_{ds} \geq 4V_T$, where V_T is the thermal voltage, and $4V_T$ is about 100 mV at room temperature. This low voltage overhead requirement allows us to operate with a power supply voltage of 2.5 V even though the circuit was fabricated in a 5 V process. The source follower buffer for the photoreceptor is the subcircuit limiting our voltage ceiling. Once this is removed from later designs, operation with a supply voltage of less than 2 V should be possible.

C. Interpixel Variation

Device mismatch is inherent in any physical circuit, and is an important consideration when designing analog circuits that will be repeated many times across a chip. We want every motion sensor on a die to exhibit similar performance. The large number of sensors on a single chip precludes off-chip trimming of each circuit to achieve matching. Some floating-gate circuits are capable of storing correction factors locally, but these would add to the size and complexity of each circuit [33].

We use two types of devices in our EMD: transistors and capacitors. Parallel plate capacitors in CMOS, even very small ones, match very well [34]. In order to study transistor matching, we fabricated arrays of several hundred *n*MOS and *p*MOS transistors, and measured their matching characteristics.

Transistor mismatch is most simply modeled as a voltage source in series with the gate. This gate voltage variation is particularly important in subthreshold operation due to the increased relative transconductance in this region. We modeled this source as a Gaussian distribution of mean zero and standard deviation σ . We estimated σ for *n*FETs and *p*FETs of various sizes by measuring I–V characteristics across the transistor arrays.

As expected, σ decreased with the square root of transistor area (data not shown). This behavior has been observed in other studies of subthreshold FETs [35]. We also found that the value of σ for *p*FETs was 2.7 times as large as the value for *n*FETs of equal gate area, on average.

Using this knowledge, we can apportion our

limited layout area in a way that maximizes interpixel matching. We analyzed subcircuits in our motion sensor to determine what effect variation in an individual device would have on the entire subcircuit. For example, analysis of the five-transistor lowpass filter (see Fig. 2(b)) reveals that mismatch in each of four transistors—the two *n*FETs in the differential pair (M_2 and M_3) and the two *p*FETs in the current mirror (M_4 and M_5)—contribute equally to the intercircuit variance of the output voltage:

$$\sigma_{out}^2 = \sigma_{M_2}^2 + \sigma_{M_3}^2 + \sigma_{M_4}^2 + \sigma_{M_5}^2 \quad (24)$$

Given a limited area for circuit layout, it follows that to minimize mismatch, we must apportion the chip area as follows:

$$\frac{A_P}{A_N} = \sqrt{\frac{N_P \sigma_p}{N_N \sigma_n}} \quad (25)$$

where A_P and A_N are the layout areas devoted to *p*FETs and *n*FETs, N_P and N_N are the number of *p*FETs and *n*FETs that contribute equally to the total circuit variance, and $\sigma_p = 2.7\sigma_n$ in our technology. We used this consideration—allocating more area for *p*FETs due to their worse matching properties—when designing the layout. We also tried to reduce the total number of *p*FETs in the circuit (e.g. removing the *p*FET current mirror from the Gilbert multipliers).

V. Experimental Results

All of our experiments were carried out on a 1×22 array of motion sensors fabricated on a $2.2 \text{ mm} \times 2.2 \text{ mm}$ die in a standard $1.2 \mu\text{m}$ CMOS process. A 2.6 mm focal length lens was mounted directly over the chip, giving a 35° field of view across the entire array. The angle ϕ between adjacent photoreceptors was 1.5° , comparable to the eyes of many flying insects [36]. The chip was biased to an appropriate operating range, and the bias settings were unchanged during all experiments, except where explicitly stated.

For experiments involving spatial integration over many sensors, the individual output currents were summed on two wires, one for the rightward-facing half-receptors (i.e. the m_1 signal in Fig. 1(a)), and one for the leftward-facing half-receptors (i.e. the m_2 signal in Fig. 1(a)). The currents were measured with off-chip sense amplifiers. The two opponent signals

were subtracted to yield a direction selective response.

We presented computer-generated visual stimuli on a standard monitor (Sony Multiscan 17se II) with a refresh rate of 70 Hz. Our software was able to update the screen at approximately the same rate. The bandwidth of the adaptive photoreceptors was set sufficiently low to attenuate screen refresh artifacts by 20 dB. This also prevented the photoreceptors from responding to the 120 Hz signal in ac incandescent lighting.

We generated visual stimuli with spatial resolution far exceeding the motion sensor array resolution. We used a 64-value gray scale to generate sinusoidal gratings and other complex stimuli of varying contrasts.

A. Direction Selectivity

Fig. 4 shows the output of a single Reichardt motion sensor and the summed output of the 22-element sensor array in response to a sinusoidal grating drifting along the sensor axis. The sensor array is highly direction selective, giving responses of opposite sign to motion in opposite directions. The

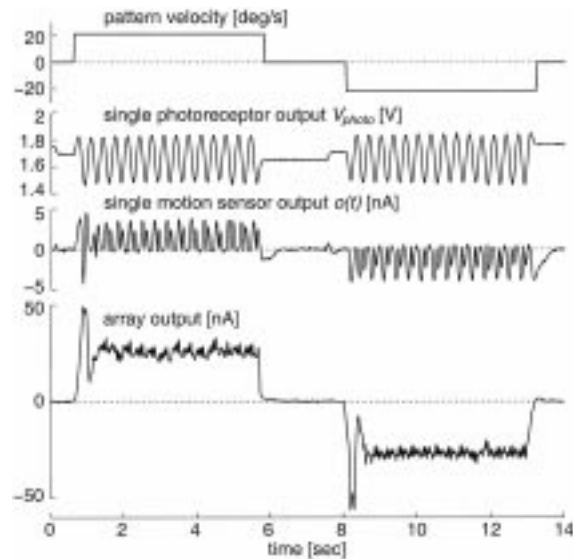


Fig. 4. Direction selectivity. A sinusoidal grating ($f_t = 3.0$ Hz, $f_s = 0.14$ cpd, contrast = 1) moved along the motion detector axis in alternating directions. Spatial integration over an array of 22 Reichardt detectors eliminates much of the pattern dependence seen in the single sensor trace.

individual sensor shows a high degree of pattern dependence superimposed on a DC direction selective response. Much of this pattern dependence is caused by device mismatch in the Gilbert multiplier. If the differential pairs are not perfectly matched, the output contains components of the raw input signals.

Pattern dependence is greatly reduced by spatially integrating over a small group of motion sensors that see different phases of the stimulus. Pattern dependence has also been observed in motion-sensitive cells in flying insects, where it is also reduced by spatial integration [29,30]. In principle, pattern dependence could also be removed through temporal integration (averaging over time), but this would limit the response time of the sensor. We chose spatial integration, which sacrifices resolution, but maintains temporal bandwidth. The transients observed at the onset of motion are also observed in biological motion-sensitive cells and have been shown to be a consequence of summing many EMDs which see different phases of a periodic stimulus [32].

Although the sensors are configured in a 1-D arrangement, they respond to motion in other directions. Fig. 5 shows the mean response of the sensors to a sinusoidal grating moving in other directions. Dashed lines indicate negative responses. The sensor exhibits a cosine tuning curve, similar to many motion-sensitive cells observed in animals [37]. This is a consequence of the spatial-frequency dependence in (12). For spatial wavelengths much greater than the photoreceptor spacing ϕ , the $\sin \phi \omega_s$ term can be approximated as $\phi \omega_s$. The temporal frequency of the stimulus moving at an angle θ stays the same as it rotates, but the spatial frequency as seen by two adjacent photoreceptors is proportional to $\cos \theta$.

This type of directional tuning is desirable because it responds strongly to slightly off-axis motion, but tends to ignore perpendicular motion. As discussed in Section IV, by adding two additional multipliers to each pixel, we can create an additional motion sensor that is oriented perpendicular to the original sensor. Together, these sensors would return signals proportional to the sine and cosine of image motion, so the direction of motion could be determined. Of course, this sensor only measures the normal component of local motion, and does not solve the aperture problem [38].

Our sensor preferentially ignores stimuli in the




Fig. 5. Directional tuning of the Reichardt detector array. A sinusoidal grating ($f_t = 3.0$ Hz, $f_s = 0.05$ cpd, contrast = 1) moved at various directions θ relative to the motion detector axis. Magnitude of normalized mean response is shown. Line shows $\cos \theta$; dashed section indicates negative response.

orthogonal direction, so if one is interested in motion only along one axis where off-axis motion is rare (such as estimating the translation or rotation of a ground vehicle), a simple 1-D sensor array might suffice, since vertical motions would be largely ignored.

B. Spatiotemporal Frequency Tuning

Next, we varied the temporal and spatial frequencies of the sinusoidal gratings. Fig. 6 shows the mean response of the sensor array, as well as the standard deviation of the signal over ten temporal cycles of the stimulus. The error bars give an indication of the magnitude of residual pattern dependence. Noise levels were far below the deterministic pattern-dependent fluctuations observed. Theoretical fits to are plotted as dashed lines. These fits use (12), including parameters from (5): the first-order temporal lowpass behavior of the adaptive photoreceptor ($\tau_{photo} = 30$ ms) and the first-order temporal highpass filter that follows ($\tau_{HPF} = 200$ ms).

The same parameters were used for all fits in Fig. 6: $\tau = 80$ ms, $\phi = 1.5^\circ$, and a fixed constant of proportionality. The circuit behaves as a Reichardt

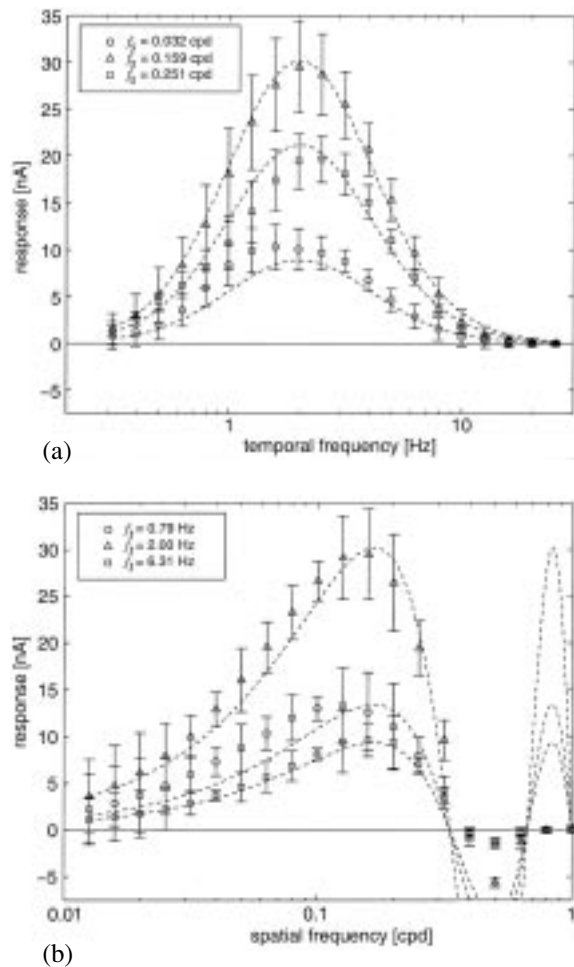


Fig. 6. Spatiotemporal frequency sensitivity of the Reichardt detector array. (a) Temporal frequency sensitivity. (b) Spatial frequency sensitivity. In this figure and in those following, error bars show one standard deviation of the time response computed over 10 stimulus periods, and represent residual deterministic pattern dependence such as that seen in Fig. 4. Noise levels were small by comparison. Dashed lines show fits to (12). Spatial aliasing is reduced by the finite photoreceptor size as well as by slightly defocused optics.

motion detector over a wide range of spatial and temporal frequencies. Spatial aliasing should produce response reversals at $N/2\phi$, where $N = 1, 2, 3, \dots$. Indeed, the first reversal can be seen near $f_s = 1/2\phi \approx 0.33$ cpd. This reversal is also observed in flies, and has been used to measure their interommatidial angle ϕ [39]. The effect of aliasing at higher spatial frequencies is reduced by the finite photoreceptor size as well as by slightly defocused

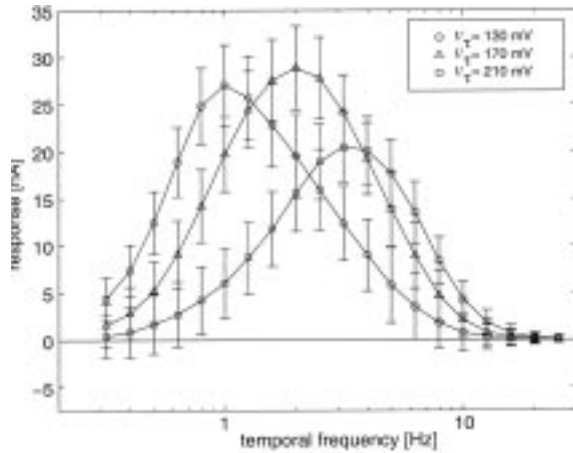


Fig. 7. Tuning the temporal frequency response. By varying the lowpass filter bias voltage V_L , we shift the sensitivity of the motion detectors ($f_s = 0.12$ cpd; contrast = 1).

optics. Both effects attenuate at high spatial frequencies.

The spatial frequency dependence of the motion sensor is determined by the angular spacing of the photoreceptors. This can be changed only by altering the layout or by changing the focal length of the lens. This would have the effect of shifting the response curve to higher or lower frequencies. Spatial prefiltering of the image could also be introduced to alter the spatial frequency response.

The location of the peak in temporal frequency response is determined by the time constant of the lowpass filter. By changing the bias V_L , we can shift the area of maximum sensitivity (see Fig. 7). This might allow us, in future versions of the EMD, to introduce temporal adaptation to sustained motion stimuli that shifts the response curve toward dominant stimulus frequencies, a phenomenon observed in biological motion sensors [40]. We could also build multiple EMDs tuned to a wide range of temporal frequencies.

The above experiments were repeated with square wave gratings, and the results were similar (data not shown).

C. Contrast Dependence

One of the biggest disadvantages of the Reichardt motion detector is its strong (quadratic) dependence on contrast. This not only confounds contrast with

spatiotemporal frequencies in the response, it also greatly amplifies the effect of high-contrast features, while attenuating low-contrast features. Studies of natural scenes have shown that low contrasts are much more common [41], so we do not want them to be underrepresented by a motion sensor. Motion sensors that respond only to high-contrast edges may work well in the laboratory, but will have sub-optimal performance in the real world.

As suggested in Section IV, by exploiting the nonlinear, saturating nature of the Gilbert multiplier we can greatly reduce this contrast dependence. Fig. 8 shows the motion sensor array's response to drifting sinusoidal gratings of varying contrast. While the contrast dependence is quadratic for low contrasts (dashed line), it saturates at moderate contrast levels as the Gilbert multipliers begin to saturate. Other stimulus information is preserved, as is shown by the saturation to different levels for stimuli with different temporal frequencies. Thus for moderate to high contrasts, our sensor's dependence on contrast is greatly reduced. By increasing the photoreceptor gain, we should be able to further decrease contrast dependence in future designs. The response to low contrasts is weak, but direction information is still encoded. Similar results were obtained for square wave gratings (data not shown).




Fig. 8. Contrast sensitivity. By operating the Gilbert multipliers outside of their linear range, we achieve reduced contrast dependence for moderate to high contrasts for two different velocities ($f_s = 0.12$ cpd). Dashed line shows quadratic dependence, as predicted by the simple Reichardt model in (12). Despite saturating behavior, temporal frequency information is still encoded.

D. Interpixel Variation

As discussed in Section IV, good interpixel matching is essential before large arrays of local motion sensors will be feasible. We measured the matching characteristics of our motion sensors across the 22-sensor 1-D array that spanned 1.3 mm. We used a sinusoidal grating of fixed spatial frequency, and measured the mean response of each sensor in the array as we varied its velocity from $-200^\circ/\text{s}$ to $+200^\circ/\text{s}$. Fig. 9 shows the mean and standard deviations of the 22 responses measured across the chip. The ratio of standard deviation to maximum mean response varied between 0.10 and 0.25, depending on the temporal frequency. The individual sensors perform similarly, indicating good matching of gains, DC levels, and time constants.

E. Robustness to External Noise

Stimulus contrast is difficult to characterize in complex natural scenes that contain many spatial frequencies. To demonstrate and quantify robustness, we devised stimuli more complex than drifting sine or square waves. Analysis of natural scenes has revealed that spatial frequency power spectra of natural scenes exhibit a $1/f^N$ behavior, where $N \approx 2$ [42]. Temporal frequency power spectra also follow this behavior, resulting in amplitude spectra with $1/f$

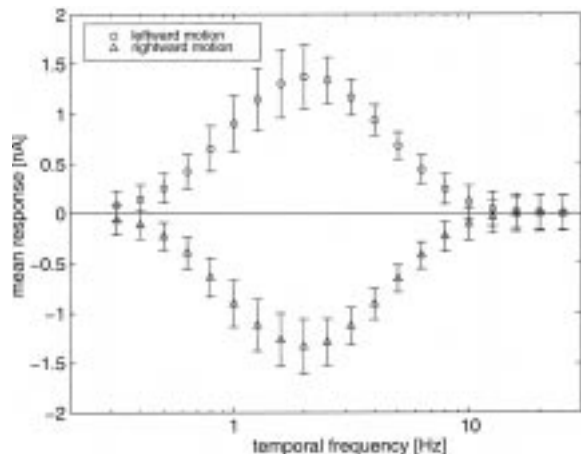


Fig. 9. Interpixel variation. Here, error bars show standard deviation of mean single-element response across the 22-element array for leftward and rightward motion at different temporal frequencies or, equivalently, speeds ($f_s = 0.12$ cpd; contrast = 1).

shapes. To better approximate natural scenes, we generated 1-D patterns with a $1/f$ spatial frequency amplitude spectrum. By randomizing the phase spectrum, any number of distinct patterns with “natural” second-order statistics can be synthesized. Our sensor array responds well to these $1/f$ patterns.

To test the robustness of our sensor, we added either spatial or temporal $1/f$ noise to a moving stimulus. We chose not to use white noise because it contains relatively little power in the low frequencies that the motion sensor array can resolve. Fig. 10(a) shows an 1-D $1/f$ stimulus moving rightward at a constant speed as a function of position and time. The slope of the pattern corresponds to its velocity in this spatio-temporal diagram. We introduced spatial noise by adding a fixed $1/f$ pattern (with different random phase spectrum) to the scene. The signal-to-noise ratio (SNR) of this stimulus is +25 dB, where SNR is defined as

$$\text{SNR}_{\text{dB}} = 10 \log_{10} \frac{\sigma_{\text{signal}}^2}{\sigma_{\text{noise}}^2} \quad (26)$$

In Fig. 10(b), we have reduced the SNR to -8 dB. After adding noise to a stimulus, we rescaled the image to cover the entire dynamic range our monitor was able to display, from white to black. This allowed us to cover a greater range of SNRs, but reduced the signal contrast as more noise was added.

Fig. 10(b) also shows three “snapshots” of the stimulus at different times. Our sensor array can easily discriminate its direction of motion in this case. It is clear from looking at the snapshots that there are no obvious “features” that one could track from one frame to the next. Although human observers have no problem seeing the direction of motion at this SNR level, features seem to wax and wane in a seemingly random manner as the stimulus moves across the retina.

We also added $1/f$ temporal noise. Fig. 11(a) shows the position vs. time graph of the stimulus with $\text{SNR} = +25$ dB. The perception is of motion in one direction, with full-field flicker or illumination modulation added. Fig. 11(b) shows the case where $\text{SNR} = -8$ dB, along with three time traces at fixed locations, such as three photoreceptors might see. There are no obvious temporal features that move from one trace to the next.

We tested the direction selectivity of the sensor array as a function of SNR for spatial noise (see Fig.




Fig. 10. Adding spatial noise. Diagrams show an one-dimensional $1/f$ stimulus moving rightward at a constant velocity. In both cases, $1/f$ spatial noise has been added with an effective (a) SNR of +25 dB (b) or of -8 dB. Also shown are three “snapshots” of each image. Notice that under low SNR conditions, few obvious spatial features are apparent that could be reliably tracked to estimate image motion.

12) and temporal noise (see Fig. 13). Both figures show the mean and standard deviation of time courses from the sensor array. The increase in mean response can be accounted for by the increasing signal contrast as SNR is raised while dynamic range is kept constant. Even at very low SNR levels, the sensor is still able to discriminate between leftward and rightward motion with high accuracy. Our analog, correlation-based motion algorithm is robust against noisy inputs that would foil other feature-based motion-detection schemes. Our sensor behaves quite robustly in the face of temporal noise, thanks to the opponent subtraction that eliminates common-mode signals. Note that the gradient algorithm would

fail here by reporting motion in the direction of the spatial intensity gradient, since the assumption that the image intensity is constant does not hold [43].

F. Power

One entire Reichardt motion sensor consumes 50 nW of power at $V_{dd} = 2.5$ V under normal indoor illumination of 10 cd/m^2 . Photocurrent contributes significantly to overall power consumption. Under increased illumination of 2500 cd/m^2 , the power consumption increased to 110 nW. The circuit consumes approximately the same amount of power




Fig. 11. Adding temporal noise. Diagrams show a 1-D $1/f$ stimulus moving rightward at a constant velocity. In both cases, $1/f$ temporal noise has been added. (a) SNR = +25 dB. (b) SNR = -8 dB. Also shown are three “photoreceptor signals” from each image. Notice that when the SNR is low, no obvious temporal features exist that could be used to estimate image motion.

with or without motion present. This is the lowest power requirement of any motion sensor we are aware of. A 100×100 array of 2-D Reichardt motion sensors would consume less than 1 mW.

VI. Discussion

We have described a low-power silicon implementation of the Reichardt motion detector, originally used to explain motion detection in insects and evaluated the robustness of our EMD to internal, fixed-pattern noise (interpixel device variation) and external spatial

and temporal noise. We reduced contrast dependence by exploiting circuit nonlinearities. While our sensor does not measure image velocity with high precision like feature-based motion detectors, it could be used in noisy environments to return qualitative motion information when high precision is not essential. Possible applications include autonomous navigation on small vehicles. This circuit was recently used to estimate and cancel vehicle rotation in real-world robotic experiments [44].

This EMD is now being implemented in a $0.35 \mu\text{m}$ CMOS process. An individual motion sensor in this technology consumes $3000 \mu\text{m}^2$ of silicon real-estate.




Fig. 12. Robustness with spatial noise. Performance of Reichardt detector array on direction discrimination with a $1/f$ pattern moving at a constant velocity of $14^\circ/s$ while $1/f$ spatial noise was added. Error bars show one standard deviation of the time response, and represent residual deterministic pattern dependence such as that seen in Fig. 4.




Fig. 13. Robustness with temporal noise. Performance of Reichardt detector array on direction discrimination with a $1/f$ pattern moving at a constant velocity of $14^\circ/s$ while $1/f$ temporal noise was added. Error bars show one standard deviation of the time response, and represent residual deterministic pattern dependence such as that seen in Fig. 4.

At this scale a 2-D Reichardt detector array with the same resolution as a housefly eye (approximately 80×80) could be fabricated on a $7 \text{ mm} \times 7 \text{ mm}$ chip, with less than $700 \mu\text{W}$ power dissipation. As mentioned in Section II, the addition of focal-plane processing limits resolution. However, flying insects

offer an existence proof that very rapid, vision-based navigation through unstructured 3-D environments does not require high-resolution imaging. It will be interesting to investigate algorithms that use low-resolution optic flow maps for real-time navigation. We hope this circuit can be used as a building block for increasingly sophisticated hardware sensory systems.

Acknowledgments

The authors would like to thank Alexander Borst, Martin Egelhaaf, Charles Higgins, Laurent Itti, Brad Minch, Rahul Sarpeshkar, and M.V. Srinivasan for valuable discussions. We also thank Oliver Landolt and Alberto Pesavento for useful comments on the manuscript. This work was supported by the Center for Neuromorphic Engineering at Caltech as a part of NSF's Engineering Research Center program, and by ONR.


References

1. C. Koch and G. Laurent, "Complexity and the Nervous System." *Science* 284, pp. 96–98, 1999.
2. C. Mead, *Analog VLSI and Neural Systems*. Reading, MA, Addison-Wesley, 1989.
3. R. Etienne-Cummings, J. Van der Spiegel, and P. Mueller, "A focal plane visual motion measurement sensor." *IEEE Transactions on Circuits and Systems I* 44, pp. 55–66, 1997.
4. G. L. Barrows, "Feature tracking linear optic flow sensor for 2-D optic flow measurement," in *Proceedings of the International Conference on Automation, Robotics, and Computer Vision (ICARCV'98)*, December 1998.
5. J. Kramer, "Compact integrated motion sensor with three-pixel interaction." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 18, pp. 455–460, 1996.
6. R. Sarpeshkar, J. Kramer, G. Indiveri, and C. Koch, "Analog VLSI architectures for motion processing: from fundamental limits to system applications." *Proceedings of the IEEE* 84, pp. 969–987, 1996.
7. J. Kramer, R. Sarpeshkar, and C. Koch, "Pulse-based analog VLSI velocity sensors." *IEEE Transactions on Circuits and Systems II* 44, pp. 86–101, 1997.
8. C. M. Higgins, R. A. Deutschmann, and C. Koch, "Pulse-based 2D motion sensors," in *IEEE Transactions on Circuits and Systems II*, to appear.
9. C. M. Higgins, personal correspondence, May 20, 1999.
10. J. Tanner and C. Mead, "An integrated analog optical motion sensor," in *VLSI Signal Processing II*, New York, IEEE Press, pp. 59–76, 1986.
11. R. A. Deutschmann and C. Koch, "Compact real-time 2-D gradient based analog VLSI motion sensor," in *Proceedings of*

- the *International Conference on Advances in Focal Plane Arrays and Electronic Cameras*, Zürich, 1998.
12. E. H. Adelson and J. R. Bergen, "Spatiotemporal energy models for the perception of motion." *Journal of the Optical Society of America A* 2, pp. 284–299, 1985.
 13. B. Hassenstein and W. Reichardt, "Systemtheoretische Analyse der Zeit-, Reihenfolgen-, und Vorzeichenauswertung bei der Bewegungsperzeption des Rüsselkäfers *Chlorophanus*." *Z. Naturforsch.* 11b, pp. 513–524, 1956.
 14. A. Borst and M. Egelhaaf, "Principles of visual motion detection." *Trends in Neuroscience* 12, pp. 297–306, 1989.
 15. H. B. Barlow and W. R. Levick, "The mechanism of directionally selective units in the rabbit's retina." *Journal of Physiology* 178, pp. 447–504, 1965.
 16. T. Horiuchi, J. Lazzaro, A. Moore, and C. Koch, "A delay-line based motion detection chip," in *Advances in Neural Information Processing Systems* 3, R. Lippman, et al., Eds., San Mateo, CA, Morgan Kaufman, pp. 406–412, 1991.
 17. R. G. Benson and T. Delbrück, "Direction selective silicon retina that uses null inhibition," in *Advances in Neural Information Processing Systems* 4, R. Lippman, J. Moody, and D. S. Touretzky. Eds., San Mateo, CA, Morgan Kaufman, 1992.
 18. T. Delbrück, "Silicon retina with correlation-based velocity-tuned pixels." *IEEE Transactions on Neural Networks* 4, pp. 529–541, 1993.
 19. A. G. Andreou, K. Strohhahn, and R. E. Jenkins, "Silicon retina for motion computation," in *Proceedings of the 1991 IEEE International Symposium on Circuits and Systems*, pp. 1373–1376, Singapore, June 1991.
 20. R. R. Harrison and C. Koch, "An analog VLSI model of the fly elementary motion detector" in *Advances in Neural Information Processing Systems*. M. I. Jordan, M. J. Kearns, and S. A. Solla. Eds., Cambridge, MA: MIT Press, 10, pp. 880–886, 1998.
 21. A. Moini, A. Bouzerdoum, K. Eshragian, A. Yakovlev, X. T. Nguyen, A. Blanksby, R. Beare, D. Abbott, and R. E. Bogner, "An insect vision-based motion detection chip." *IEEE Journal of Solid State Circuits* 32, pp. 279–284, 1997.
 22. H.-C. Jiang and C.-Y. Wu, "A 2-D velocity- and direction-selective sensor with BJT-based silicon retina and temporal zero-crossing detector." *IEEE Journal of Solid State Circuits* 34, pp. 241–247, 1999.
 23. J.-M. Pichon, C. Blanes, and N. Franceschini, "Visual guidance of a mobile robot equipped with a network of self-motion sensors," in *Proceedings of SPIE* 1195, pp. 44–53, 1989.
 24. N. Franceschini, J. M. Pichon, and C. Blanes, "From insect vision to robot vision." *Philosophical Transactions of the Royal Society London B* 337, pp. 283–294, 1992.
 25. J. P. H. van Santen and G. Sperling, "Elaborated Reichardt detectors." *Journal of the Optical Society of America A* 2, pp. 300–321, 1985.
 26. A. B. Watson and A. J. Ahumada, "Model of human visual-motion sensing." *Journal of the Optical Society of America A* 2, pp. 322–342, 1985.
 27. 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, pp. 9363–9367, 1990.
 28. M. Egelhaaf, A. Borst, and W. Reichardt, "Computational structure of a biological motion-detection system as revealed by local detector analysis in the fly's nervous system." *Journal of the Optical Society of America A* 6, pp. 1070–1087, 1989.
 29. W. Reichardt and M. Egelhaaf, "Properties of individual movement detectors as derived from behavioural experiments on the visual system of the fly." *Biological Cybernetics* 58, pp. 287–294, 1988.
 30. S. Single and A. Borst, "Dendritic integration and its role in computing image velocity." *Science* 281, pp. 1848–1850, 1998.
 31. T. Delbrück and C. A. Mead, "Analog VLSI phototransduction by continuous-time, adaptive, logarithmic photoreceptor circuits." CNS Memo 30, California Institute of Technology, 1996.
 32. M. Egelhaaf and A. Borst, "Transient and steady-state response properties of movement detectors." *Journal of the Optical Society of America A* 6, pp. 116–127, 1989.
 33. R. R. Harrison, P. Hasler, and B. A. Minch, "Floating-gate CMOS analog memory cell array," in *Proceedings of the IEEE International Symposium on Circuits (ISCAS'98)* 2, pp. 204–207, 1998.
 34. B. A. Minch, C. Diorio, P. Hasler, and C. Mead, "The matching of small capacitors for analog VLSI," in *Proceedings of the IEEE International Symposium on Circuits (ISCAS'96)* 1, pp. 239–241, 1996.
 35. A. Pavasović, A. G. Andreou, and C. R. Westgate, "Characterization of subthreshold MOS mismatch in transistors for VLSI systems." *Analog Integrated Circuits and Signal Processing* 6, pp. 75–85, 1994.
 36. M. F. Land, "Visual acuity in insects." *Annual Rev. Entomol.* 42, pp. 147–177, 1997.
 37. J. M. Zanker, "On the directional sensitivity of motion detectors." *Biological Cybernetics* 62, pp. 177–183, 1990.
 38. D. Marr, *Vision*. New York, W.H. Freeman & Co., 1982.
 39. K. Götz, "Die optischen Übertragungseigenschaften der Komplexaugen von *Drosophila*," *Kybernetik* 2, pp. 215–221, 1965.
 40. R. R. de Ruyter van Steveninck, W. H. Zaagman, and H. A. K. Mastebroek, "Adaptation of Transient Responses of a Movement-Sensitive Neuron in the Visual System of the Blowfly *Calliphora erythrocephala*." *Biological Cybernetics* 54, pp. 223–236, 1986.
 41. D. L. Ruderman and W. Bialek, "Statistics of natural images: scaling in the woods." *Physic Review Letters* 73, pp. 814–817, 1994.
 42. D. W. Dong and J. J. Atick, "Statistics of natural time-varying images." *Network* 6, pp. 345–358, 1995.
 43. B. Klaus and P. Horn, *Robot Vision*. Cambridge, MA, MIT Press, 1986.
 44. R. R. Harrison and C. Koch, "A robust analog VLSI motion sensor based on the visual system of the fly." *Auton. Robots* 7, pp. 211–224, 1999.



Reid R. Harrison was born in DeFuniak Springs, Florida, in 1972. He received the B.S. degree in electrical engineering from the University of Florida in 1994. After working for a short time at the Jet Propulsion Laboratory and Los Alamos National Laboratory, he joined the Computation and Neural Systems program at the California Institute of Technology. He was awarded his Ph.D. from Caltech in 2000. His research interests include silicon implementations of biological sensory systems, low-power analog VLSI circuits, and neuroprosthetics.



Christof Koch was born in Kansas City, Missouri, in 1956 and studied physics and philosophy at the University of Tübingen in Germany. He was awarded his Ph.D. in biophysics from the Max Planck Institute for Biological Cybernetics in Tübingen in 1982. He worked until 1986 at MIT's Artificial Intelligence Laboratory before joining the Computation and Neural Systems (CNS) program at the California Institute of Technology where he is now a Full Professor and Executive Officer of the CNS department. His research interests include the biophysical mechanisms underlying information storage and processing in individual neurons, neuro-morphic vision chips, and the neuronal basis of visual awareness and consciousness.