# 2 Fly-Inspired VLSI Vision Sensors

Reid R. Harrison

Engineers have long looked to nature for inspiration. The diversity of life produced by five billion years of evolution provides countless existence proofs of organic machines with abilities that far surpass those of our own relatively crude automata. We have learned how to harness large amounts of energy and thus far exceed the capabilities of biological systems in some ways (e.g., supersonic flight, space travel, and global communications). However, biological information processing systems (i.e., brains) far outperform today's most advanced computers at tasks involving real-time pattern recognition and perception in complex, uncontrolled environments. If we take energy efficiency into account, the performance gap widens. The human brain dissipates 12 W of power, independent of mental activity. A modern microprocessor dissipates around 50 W, and is equivalent to a vanishingly small fraction of our brain's functionality.

Indeed, the human brain is a daunting goal for biologists and engineers alike. Our brain takes several years to fully develop, and contains between $10^{10}$ and $10^{11}$ neurons (nerve cells), each communicating with $10^3$ other cells, on average. Brains of other animals (particularly invertebrates) are much smaller but still perform remarkably complex computations. Insect brains, for example, typically contain between $10^5$ and $10^6$ neurons, yet insects perform sophisticated information-processing tasks rapidly and efficiently.

We have attempted to extract computational principles from the visual system of the fly and apply these principles to an engineered system—an integrated, low-power visual motion sensor. As our engineering tool, we use very large-scale integration (VLSI) of silicon circuits—the most advanced information-processing substrate available today. In particular, we explore continuous-time (unclocked), continuous-value (analog) circuit architectures. This approach was pioneered by Mead and colleagues beginning the in 1980s (Mead, 1989).

The fly is an attractive target for biologically inspired approaches to engineering. Its brain and sensory systems have been studied for decades, so much is known about their operation. Of course, we are still decades (or centuries) away from understanding the entire system,

but a wealth of behavioral and electrophysiological data has led to the development of several models of information processing.

## THE VISUAL SYSTEM OF THE FLY

Vision is a vitally important sense for flying insects. In the housefly's brain, more than half of the 350,000 neurons are believed to have some role in visual processing. The fly's optic lobes contain motion-sensitive neurons that respond to moving stimuli over large portions of the visual field. Many of these neurons have been linked to specific visually guided behaviors that help the animal navigate through a complex environment in a robust manner (Egelhaaf and Borst, 1993).

Insects process visual motion information in a local, hierarchical manner. This information processing begins at the sensor—the retina. Despite the multilens construction of the compound eye, the pattern projected onto the underlying retina is a single image of the visual scene. Photoreceptors in the retina adapt to the ambient light level, and signal temporal deviations from this level. These signals are passed on to the next layer of cells, the lamina. Lamina cells generally show transient or high-pass responses, emphasizing temporal change (Weckström, Juusola, and Laughlin, 1992). The next stage of processing is the medulla, a layer of cells that are extremely difficult to study directly due to their small size. Indirect evidence suggests that local measures of motion (i.e., between adjacent photoreceptors) are computed here. These local, direction-selective motion estimates are integrated by large tangential cells in the lobular plate (Hausen and Egelhaaf, 1989). The housefly has 50–60 tangential cells in each hemisphere of its brain. These are the best-studied cells in the fly visual system, and much is known about their properties.

Lobular plate cells generally respond to motion over large parts of the visual field. Some of these cells seem to be matched filters for the optic flow patterns produced by rotation or translation along particular axes (Krapp and Hengstenberg, 1996). Some of these cells most likely control compensatory motor reflexes that prevent the fly from rotating during flight. Others are sensitive only to small objects moving across the visual field (Egelhaaf, 1985). It is believed that these "figure detection" cells allow the fly to locate nearby objects through motion parallax (Kimmerle, Warzecha, and Egelhaaf, 1997). All of these sensory abilities require that motion first be detected locally between every pair of photoreceptors.

### Photoreception

Each eye of the blowfly *Calliphora erythrocephala* consists of approximately 6000 individual lenses. Beneath each lens is a cluster of eight

light-sensitive cells. Each lens and its associated photoreceptors form a unit called an ommatidium. Six of the eight photoreceptors are used to implement neural superposition, a technique to increase the effective lens diameter by pooling the responses of neighboring ommatidia. The other two photoreceptors do not seem to be involved in the detection of motion. Mutants with these photoreceptors impaired cannot discriminate colors, but show no motion-related deficits (Heisenberg and Buchner, 1977). From an information-processing perspective, each ommatidium records one "pixel" of the external world's image. Inter-ommatidial angular spacing is 1.1–1.3° (Land, 1997). This angular resolution is approximately 150 times worse than the 0.008° resolution in the foveated region of the human retina (Wandell, 1995). (This is roughly equivalent to having 20/3000 vision.) Although the ommatidia are arranged in a hexagonal array, it is useful to think of the equivalent array size in terms of the standard rectangular array used by computer monitors and digital cameras. Taking the square root of the number of ommatidia (6000), we see that *Calliphora*'s eye is roughly equivalent to a 77 by 77-pixel array covering one visual hemifield. The eye of the fruit fly *Drosophila* contains only 700 ommatidia, resulting in an equivalent array size of 26 by 26 (Land, 1997). Today's cheap digital cameras provide 640 by 480-pixel images, and emerging photo-quality digital cameras provide 1800 by 1200 pixels or more—around two orders of magnitude more photoreceptors than a fly's eye. Typical cameras concentrate these pixels into a 40° field of view, while each fly eye sees nearly a complete hemisphere.

It is remarkable that flies are capable of such impressive navigation when one considers their low-resolution eyes. This limited spatial acuity is a consequence of the compound eye design. In order to increase spatial acuity, more ommatidia are required. However, the resolving capability of each ommatidium is limited by diffraction, so each lens must also be made larger. If we wanted to build a compound eye with the acuity of the human fovea (0.008°), it would have a radius of 11.7 meters! The visual acuity of the largest insect eye in nature (that of the aeshnid dragonfly) reaches 0.24° in its most acute zone, still 30 times coarser than the human fovea (Land, 1997).

Although inferior to human eyes spatially, fly vision far exceeds ours temporally. Human vision is sensitive to temporal modulations up to 20 or 30 Hz, while fly photoreceptors respond to temporal frequencies as high as 300 Hz (Autrum, 1958).

## SIGNAL PROCESSING IN THE PERIPHERAL OPTIC LOBE

The laminar region, also called the first optic ganglion, contains cells that exhibit transient responses to step intensity changes. The large monopolar cells (LMCs) in this ganglion ignore the DC light level but

amplify temporal changes (Weckström, Juusola, and Laughlin, 1992). This high-pass response has been shown to optimize information transfer through this region (Laughlin, 1994). Laminar cells do not exhibit motion-specific responses. There is a strong retinotopic organization from the retina through the lamina to the next layer, the medulla. Every ommatidium has an associated neural "cartridge" beneath it in these underlying ganglia, suggesting many identical processing units operating in parallel (Strausfeld, 1976).

Cells in this second optic ganglion are extremely small and difficult to record from, and little is know about their structure or function. DeVoe recorded from medullar cells in *Calliphora* and reported a wide variety of response characteristics: transient temporal responses, sustained responses, directional motion responses, and nondirectional motion responses (DeVoe and Ockleford, 1976; DeVoe, 1980).

### The Tangential Cells of the Lobular Plate

The third optic ganglion is also known as the lobula-lobular plate complex. At this point in the optic lobe, the retinotopic organization ends with massive spatial convergence. Information from several thousand photoreceptors converges onto 50–60 tangential cells. These cells have broad dendritic trees that receive synaptic input from large regions of the medulla, resulting in large visual receptive fields (Hausen, 1982a, b; Hengstenberg, 1982; Hausen, 1984; Krapp and Hengstenberg, 1996).

A subset of these neurons were found to respond primarily to horizontal motion, and these cells were given names beginning with "H." H1 is a spiking neuron that responds to back-to-front optic flow. HSS, HSE, and HSN are graded potential (nonspiking) neurons covering the southern, equatorial, and northern regions of the visual field, respectively. Collectively called the HS cells, these neurons are depolarized by full-field visual motion from the front to the back of the eye, and hyperpolarized by back-to-front motion. They have been shown to encode horizontal motion as effectively as the spiking H1 cell (Haag and Borst, 1997). Each HS cell integrates signals from an ipsilateral retinotopic array of elementary motion detectors (EMDs), units in the medulla that estimate local motion in small areas of the visual field. The HS cells synapse onto descending, spiking neurons, which relay information to the motor centers of the thoracic ganglion. Another class of neurons, the VS cells, respond to vertical motion. Recently, it has been shown that these HS and VS cells are not simply responsive to motion along one axis, but rather act as matched filters for complex patterns of optic flow that would be produced by body rotations (Krapp and Hengstenberg, 1996).

### Visually Guided Behaviors

Flies rely heavily on visual motion information to survive. In the fly, motion information is known to underlie many important behaviors including stabilization during flight, orienting toward small, rapidly moving objects (Egelhaaf and Borst, 1993), and estimating time-to-contact for safe landings (Borst and Bahde, 1988). Some motion-related tasks, like extending the legs for landing, can be executed less than 70 msec after stimulus presentation. Wagner reports a 30 msec reaction time for male flies chasing prospective mates (Wagner, 1986). The computational machinery performing this sensory processing is fast, small, low power, and robust.

Flies use visual motion information to estimate self-rotation and to generate a compensatory torque response to maintain stability during flight. This well-studied behavior is known as the optomotor response. It is interesting from an engineering standpoint because it extracts relevant information from a dynamic, unstructured environment on the basis of passive sensors and uses this information to generate appropriate motor commands during flight. This system is implemented in biological hardware that is many orders of magnitude smaller and more power efficient than charge-coupled device (CCD) imagers coupled to a conventional digital microprocessor.

Much of the computation underlying the optomotor control system is performed by the HS cells (Geiger and Nässel, 1981, 1982; Egelhaaf et al., 1988; Hausen and Wehrhahn, 1990; Egelhaaf and Borst, 1993). This well-studied system estimates rotation from optic flow and uses this information to produce a stabilizing torque with the wings (Götz, 1975; Warzecha and Egelhaaf, 1996).

## THE REICHARDT MOTION DETECTOR MODEL

Much is understood about the mechanisms used by flies to extract visual motion information from their retinal signals. The most-well-known model of fly motion detection is the Reichardt model, first proposed by Hassenstein and Reichardt in 1956 to explain visually mediated behaviors in a walking beetle (Hassenstein and Reichardt, 1956). This model is a type of correlation-based motion detector, which measures spatiotemporal correlations caused by moving objects.

This model has been successful at explaining both detailed electrophysiological responses of motion-sensitive neurons to visual stimuli (Egelhaaf and Borst, 1989; Zanker, 1990) and visually guided behavioral responses (Reichardt and Poggio, 1976; Reichardt and Egelhaaf, 1988; Borst, 1990; Warzecha and Egelhaaf, 1996). Modified versions of the Reichardt model have also been used to explain motion perception

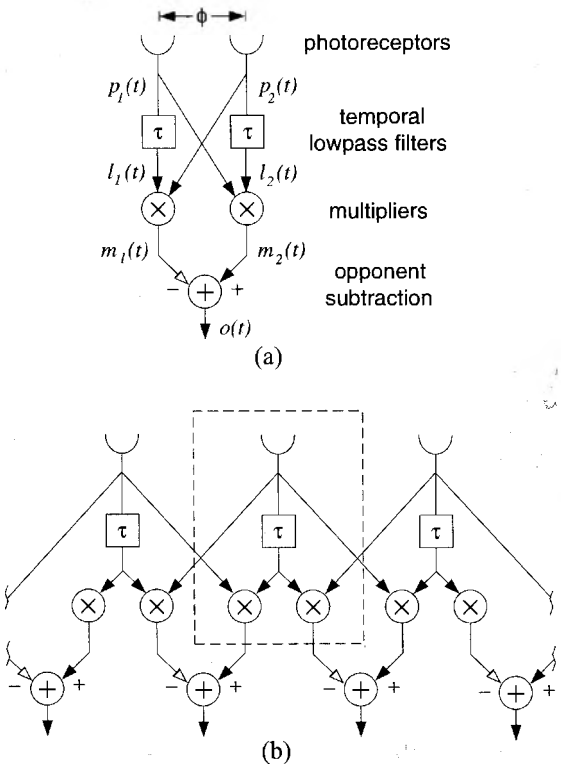


**Figure 2.1** Reichardt motion detector architecture. (a) The signal from one photoreceptor is correlated with the delayed signal from an adjacent photoreceptor. Direction selectivity is increased by subtracting the responses of two half-detectors in opponency. (b) A 1-D array of Reichardt motion detectors, 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.

properties in vertebrates, including humans (Borst and Egelhaaf, 1989; Clifford, Ibbotson, and Langley, 1997).

### Theoretical Analysis

We now describe the Reichardt motion detection algorithm, which will underlie our hardware implementation. As mentioned above, the Reichardt motion detector is a correlation-based algorithm, whose output is equivalent to the output of the spatiotemporal motion energy model proposed by Adelson and Bergen (1985) and others (van Santen and Sperling, 1985; Watson and Ahumada, 1985).

The basic idea of the Reichardt motion detector is to correlate the signal from one photoreceptor with the delayed signal from an adjacent photoreceptor (figure 2.1a). 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, strong direction selectivity is achieved (Borst and Egelhaaf, 1990).

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

$$i(x,t) = I + \Delta I \sin[2\pi f_s(x + vt)] \tag{2.1}$$

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 equation 2.1 as

$$i(x,t) = I + \Delta I \sin(\omega_t t + \omega_s x) \tag{2.2}$$

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

$$p_1(t) = |H(\omega_t)|\Delta I \sin\left(\omega_t t - \omega_s \frac{\phi}{2}\right) \tag{2.3}$$

$$p_2(t) = |H(\omega_t)|\Delta I \sin\left(\omega_t t + \omega_s \frac{\phi}{2}\right) \tag{2.4}$$

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 high-pass behavior, which eliminates the DC component of illumination $I$. We model the photoreceptor response as

$$H(\omega_t) = \frac{Kj\omega_t\tau_H}{(j\omega_t\tau_H + 1)(j\omega_t\tau_{photo} + 1)} \tag{2.5}$$

where $\tau_H$ is the time constant of the DC-blocking high-pass filter, $\tau_{photo}$ is the time constant defining the photoreceptor bandwidth, and $K$ is a constant of proportionality.

The delay required by the delay-and-correlate motion detector architecture is implemented using the phase lag inherent in a first-order low-pass filter. Low-pass filtering each photoreceptor signal yields

$$l_1(t) = \frac{|H(\omega_t)|\Delta I}{\sqrt{\tau^2\omega_t^2 + 1}} \sin\left(\omega_t t - \omega_s \frac{\phi}{2} - \tan^{-1}\,\tau\omega_t\right) \tag{2.6}$$

$$l_2(t) = \frac{|H(\omega_t)|\Delta I}{\sqrt{\tau^2\omega_t^2 + 1}} \sin\left(\omega_t t + \omega_s \frac{\phi}{2} - \tan^{-1}\,\tau\omega_t\right) \tag{2.7}$$

Correlation is accomplished by multiplying the phase-lagged signals with adjacent, nondelayed signals. The results are two "half-detector" responses:

$$m_1(t) = G[\cos(\omega_s\phi + P) - \cos(2\omega_t t - P)] \tag{2.8}$$

$$m_2(t) = G[\cos(\omega_s\phi - P) - \cos(2\omega_t t - P)] \tag{2.9}$$

where

$$G = \frac{(|H(\omega_t)|\Delta I)^2}{2\sqrt{\tau^2\omega_t^2 + 1}} \tag{2.10}$$

$$P = \tan^{-1}\tau\omega_t \tag{2.11}$$

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

$$o(t) = (\Delta I)^2|H(\omega_t)|^2\frac{\tau\omega_t}{\tau^2\omega_t^2 + 1}\sin\phi\omega_s \tag{2.12}$$

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:

$$o(t) = (\Delta I)^2|H(\omega_s v)|^2\frac{\tau\omega_s v}{\tau^2\omega_s^2 v^2 + 1}\sin\phi\omega_s \tag{2.13}$$

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. Notice that the sin $\phi\omega_s$ term predicts spatial aliasing, as it becomes negative for $1/2\phi < f_s < 1/\phi$.

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, Borst, and Reichardt, 1989.) These oscillations may be reduced by using an array of Reichardt motion detectors (figure 2.1b) and summing their responses. This has the effect of integrating over different phases of the stimulus and canceling pattern-dependent oscillations, and has been proposed as a single model of tangential cell integration (Reichardt and Egelhaaf, 1988; Single and Borst, 1998).

### Hardware Implementations

Early attempts to implement the intensity-based Reichardt architecture in silicon used translinear, current-mode circuits (Andreou, Strohbehn, and Jenkins, 1991; Harrison and Koch, 1998). As we showed in the previous section, 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 (Moini et al., 1997; Jiang and Wu, 1999). These circuits would not be expected to perform well in noisy, low-contrast environments without additional image preprocessing. Another VLSI implementation involved continuous-level signal processing after the photoreceptors, but the final motion detector output was a binary value (Liu, 1997). 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 (Pichon, Blanes, and Franceschini, 1989; Franceschini, Pichon, and Blanes, 1992).

## CIRCUIT ARCHITECTURE

We implemented the Reichardt motion detector model in continuous-time VLSI circuitry. To the best of our knowledge, this is the closest approximation to this biological motion sensor that has been built.

We measure light intensity with an adaptive photoreceptor circuit developed by Delbrück and Mead (1996). This four-transistor circuit uses a substrate photodiode and source follower ($M_1$) to convert incident light into a logarithmically encoded voltage (figure 2.2a). 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, 1996.)

The adaptive photoreceptor signal is sent to a $g_mC$ high-pass filter (figure 2.2a). 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 high-pass 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 $g_mC$ low-pass filter (figure 2.2b). The bias transistor in the circuit was made several times minimum size to improve time-constant matching across the chip. By

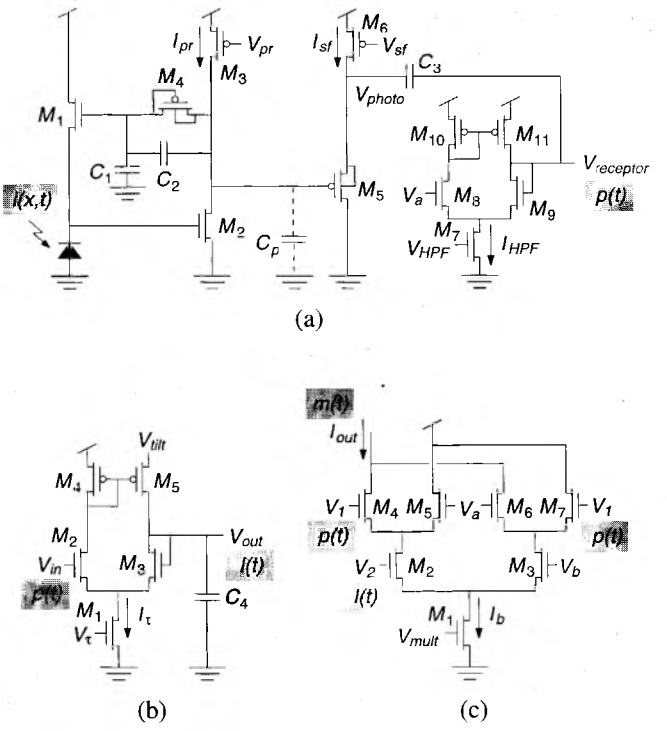


**Figure 2.2** Voltage-mode EMD subcircuits. Shaded labels indicate corresponding signals from figure 2.1a. (a) Adaptive photoreceptor ($M_1$–$M_4$, $C_1$–$C_2$) with source follower ($M_5$–$M_6$) and temporal highpass $g_mC$ filter ($M_7$–$M_{11}$, $C_3$) to remove the DC component of $V_{photo}$. (b) Temporal lowpass $g_mC$ filter. This circuit's phase lag acts as a delay. (c) Gilbert multiplier. This circuit multiplies delayed and nondelayed photoreceptor signals. The output is a current $I_{out}$, which allows for easy spatial summation.

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

Correlation is approximated by a Gilbert multiplier (figure 2.2c). The input $V_2$ comes from the low-pass filter, and $V_1$ comes from the high-pass-filtered photoreceptor from an adjacent pixel (figure 2.1b). The voltage $V_a$ is the reference voltage used by the high-pass 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

$$I_d = I_0 e^{\kappa V_g / U_T} (e^{-V_s/U_T} - e^{-V_d/U_T}) \tag{2.14}$$

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, $\kappa$ is the gate efficiency factor (typically around 0.7), and $U_T$ is the thermal voltage $kT/q$ (Mead, 1989). Subthreshold FETs exhibit exponential behavior, much like the bipolar junction transistors (BJTs) with which the Gilbert

multiplier was originally built. We take a single-ended current-mode output from the circuit, which gives us

$$I_{out} = \frac{I_b}{2} + \frac{I_b}{2} \tanh\left[\frac{\kappa(V_1 - V_a)}{2U_T}\right] \tanh\left[\frac{\kappa(V_2 - V_b)}{2U_T}\right] \tag{2.15}$$

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

$$I_{out} = \frac{I_b}{2} + \frac{\kappa^2 I_b}{8U_T^2}(V_1 - V_a)(V_2 - V_b) \tag{2.16}$$

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 low-pass filter output, we lowered the source voltage of the output FET in the current mirror of the $g_mC$ filter (see figure 2.2b). By placing the $V_{tilt}$ bias a few tens of millivolts below $V_{dd}$, we lower the DC output level by $(V_{dd} - V_{tilt})/\kappa$. This source voltage "tilt" increases the time constant of the low-pass filter, but we can compensate by raising $I_\tau$. 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 equation 2.15 that the circuit output saturates for differential inputs greater than about $4U_T \approx 100$ 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 (Egelhaaf and Borst, 1989).

Figure 2.3 shows the layout for one 1-D motion sensor, corresponding to the circuit element outlined in figure 2.1b. All experimental results shown below were measured from arrays of this circuit, which was fabricated in a 1.2 μm double-poly, double-metal *n*-well CMOS process, yielding a pixel size of 61 μm by 199 μm with 32 transistors and 4 capacitors totaling 3.0 pF.

The outputs of all EMD pairs were summed to simulate the wide-field motion-sensitive neurons found in flies. We mounted a 2.6 mm lens over the chip, which gave the photoreceptors an angular spacing of 1.3° (similar to the 1–2° angular spacing observed in fly eyes), and a total field of view of 30° (much less than the fly's eye, which sees almost an entire visual hemifield). The low-pass filter time constant was set to 50 msec, and the band-pass filters were set to pass frequencies between 0.5 Hz and 8 Hz. This circuit is characterized in detail in Harrison and Koch (2000).

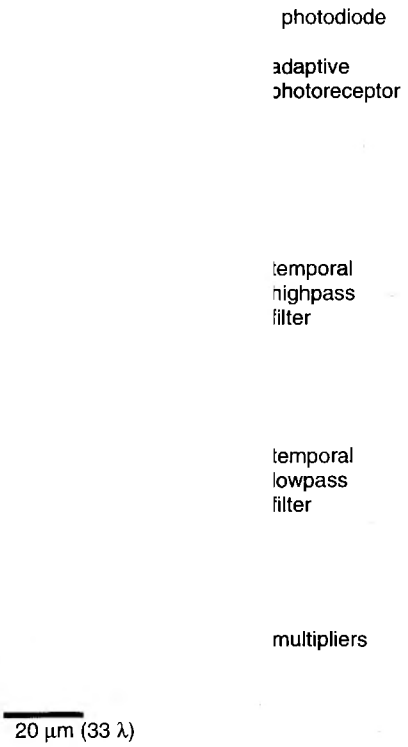


**Figure 2.3** Voltage-mode motion detector layout. Cell measures 61 μm × 199 μm in a standard 1.2 μm process with 32 transistors and 4 capacitors totaling 3.0 pF. In order to build a 2-D motion sensor, we need only add two more multiplier circuits, and some additional interpixel wiring.

## OPTOMOTOR CONTROL

Earlier, we introduced the optomotor response, where visual motion information is used as a feedback control signal to estimate and cancel self-rotation. This sensorimotor loop is perhaps the best-studied visually guided behavior of the fly. We will first describe optomotor experiments performed (by others) with flies and then describe real-time experiments where our VLSI sensor was compared directly against flies.

### Experiments Previously Performed on Flies

Warzecha and Egelhaaf (1996) recently characterized the optomotor behavior of the fly under closed-loop conditions. A female sheepfly (*Lucilia cuprina*, Calliphoridae) was rigidly attached to a meter that

measured yaw torque produced while the fly attempted to turn in response to visual stimuli (figure 2.4a), reducing the fly's behavior to a single degree of freedom. Vertical bars were presented to a large region of the fly's visual field, and could be drifted clockwise or counterclockwise. In closed-loop experiments, the fly's yaw torque was measured in real time and scaled by a constant gain term to yield angular velocity. This simulates the observed dominance of air friction in determining the instantaneous angular velocity in flies (Reichardt and Poggio, 1976). The fly's simulated angular velocity was subtracted from the angular velocity imposed by the experimenter. The resulting signal was used to control the drift rate of the visual stimulus. This simulated free-flight conditions, and allowed evaluation of the optomotor system performance.

The imposed motion schedule consisted of 3.75 sec of zero imposed motion, then 7.5 sec of clockwise rotation at 44°/sec. Figure 2.5a shows the torque data and resulting stimulus position for an individual trial. Figure 2.5b shows the averaged data over 139 trials in a total of five animals. (See Warzecha and Egelhaaf, 1996 for details on the experimental protocol.)

The fly is able to stabilize its flight and cancel out most of the imposed motion. Simulation results suggest that the nonmonotonic temporal frequency response of Reichardt motion detectors results in greater stability for the optomotor control system (Warzecha and Egelhaaf, 1996). The individual trials show an oscillatory component to the torque response around 2 Hz. This oscillation is not phase-locked to the stimulus because it is not present in the average torque trace. Oscillations are not observed under open-loop conditions, suggesting they arise from optomotor feedback (Geiger and Poggio, 1981; Warzecha and Egelhaaf, 1996). Notice that despite the large amplitude of the torque oscillations, the position trace is not dominated by this effect. This fluctuation amplitude, in terms of number of photoreceptors, is close to the amplitude observed in human microsaccades (Warzecha and Egelhaaf, 1996). Poggio and colleagues observed similar oscillations in closed-loop experiments and proposed that they arose from the 60–75 msec synaptic delay inherent in the fly visual system (Geiger and Poggio, 1981; Poggio and Reichardt, 1981).

### Duplicating Experiments with the Silicon System

We were able to directly replicate these experiments with our silicon analog of the optomotor system (figure 2.4b). To provide visual stimulation, we used an LED display with a 200 Hz refresh rate, which is currently being used to test flies in closed-loop experiments. The stimulation time schedule was identical to the fly experiments, but an

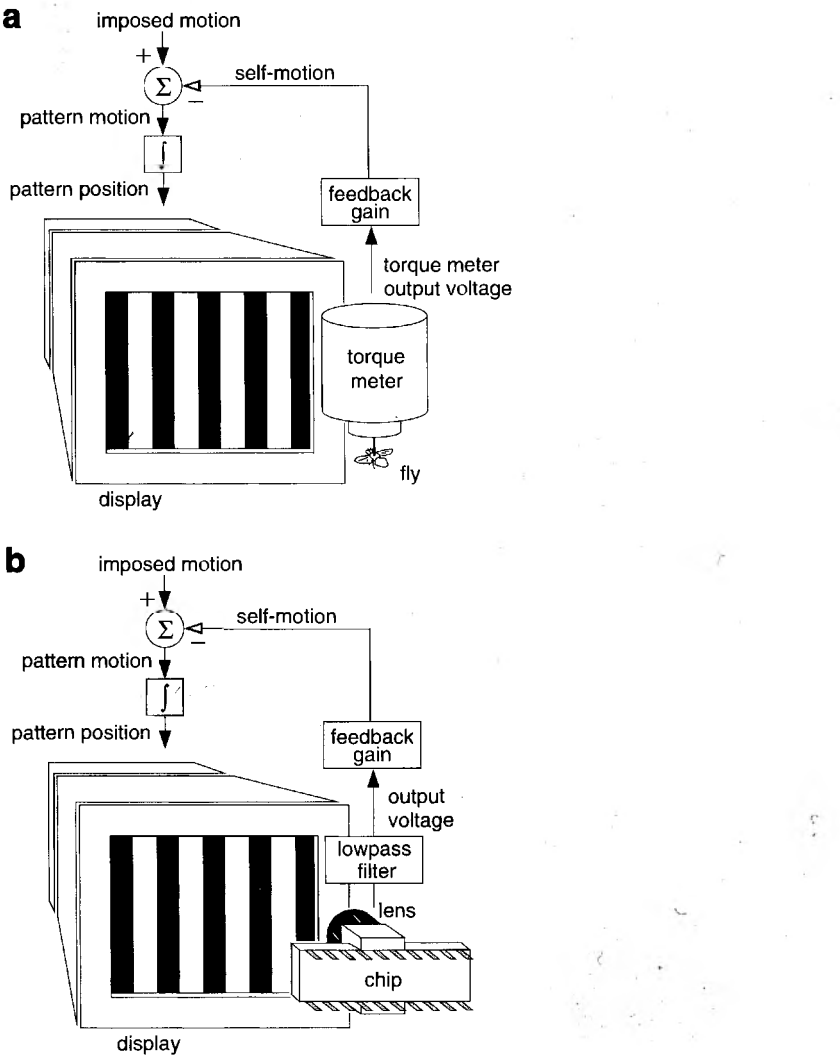


**Figure 2.4** Experimental methodology. (a) Setup used by Warzecha and Egelhaaf to measure the closed-loop torque response of the sheepfly *Lucilia* (Warzecha and Egelhaaf, 1996). The torque meter output is scaled to produce a measure of what the fly's self-motion would be if it were free to rotate. This self-motion is subtracted from the imposed motion to determine the pattern motion, creating the illusion of free flight in a room with distant walls. (Only rotation, not translation, is simulated.) (b) Setup used to replicate the closed-loop experiments with the silicon model. The output voltage from the circuit is used in place of the torque meter output voltage. The rest of the system is identical to (a).

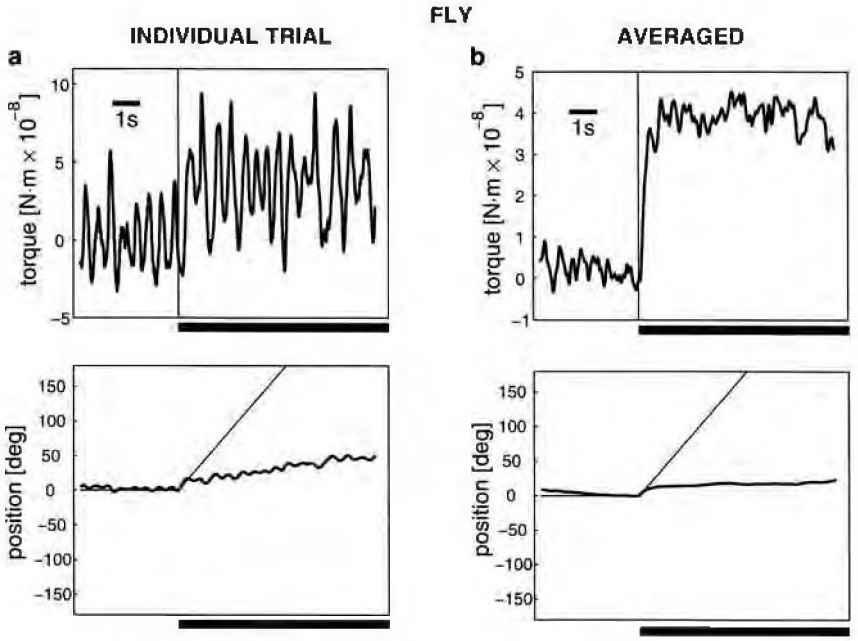


**Figure 2.5** Fly's optomotor behavior. (a) Torque (top) and angular position (bottom) vs. time for an individual closed-loop trial with a fly. Dark horizontal bar indicates experimenter-imposed rotation. Thin lines on position trace indicate position in the open-loop case. Most of the imposed rotation is canceled out by the optomotor control system. Since the position is proportional to the integral of the torque (see text for details), large torque oscillations do not cause large position oscillations. (b) Averaged torque response and angular position trace for multiple trials ($N = 139$, 5 flies). The fly showed an average drift of 9.4% of the open-loop drift velocity, with position fluctuations of 7.8° (standard deviation) about this drift. Data in (a) and (b) redrawn from Warzecha and Egelhaaf, 1996.

angular velocity of 50°/sec was used. Our chip had a much smaller field of view (10°) than the fly, so we set the stimulus distance such that the EMD array saw approximately one wavelength of the pattern. The output signal from the silicon model of the HS cell was passed through an off-chip first-order low-pass filter with a time constant of 680 msec, modeling the behavior of the thoracic motor centers (Egelhaaf, 1987; Wolf and Heisenberg, 1990; Warzecha and Egelhaaf, 1996). The filtered output of the chip was treated exactly like the signal from the torque meter in the fly experiments, and closed-loop experiments were run in real time. Figure 2.6a shows torque and position data from the chip for an individual trial, and figure 2.6b shows the averaged response over 100 trials.

The silicon system shows the same ability to greatly cancel the imposed motion. The fly showed an average drift of 9.4% of the open-loop drift velocity, with position fluctuations of 7.8° standard deviation

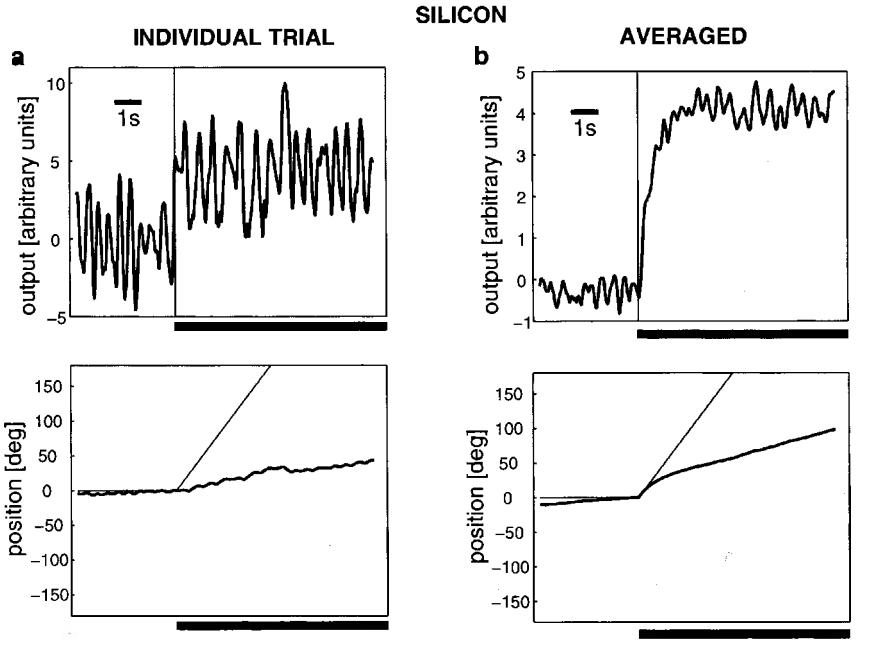


**Figure 2.6** The optomotor behavior of our silicon system. (a) Chip output signal (analogous to torque) and position vs. time for the silicon system in an individual trial. (b) Averaged torque response and angular position trace for multiple trials ($N = 100$, 1 chip). The chip showed an average drift of 22% of the open-loop drift velocity, with position fluctuations of 6.2° (s.d.) about this drift.

(SD) about this drift. The chip showed an average drift of 22% of the open-loop drift velocity, with position fluctuations of 6.2° SD about this drift. Also, we observe the same 2 Hz oscillations in the individual trials. Because we did not build any explicit delay into our system, this demonstrates that the phase lags and nonlinearities in this simple model are sufficient to produce oscillations, even in the absence of additional synaptic delays.

We believe this hardware modeling approach will prove increasingly valuable in the future, as biological models of the neural circuitry underlying more complex and sophisticated behaviors arise. To simulate a sensorimotor system in software, one must construct *two* models: a model of the biological system, and a model of the world. The physical environment is an essential element in a sensorimotor feedback loop, so this world model must increase in detail as we study more advanced behaviors. Because animals interact with their three-dimensional environment in very dynamic ways, it may not be long before software simulations of sensorimotor systems require more computational resources to model the world than to model the neural circuitry of interest.

## APPLICATION TO AUTONOMOUS VEHICLE CONTROL

Optic-flow patterns produced by self-motion are one of the richest sources of navigation information available to a mobile creature (Gibson, 1950). As an animal moves through its environment, images of the outside world move across its retina in predictable ways. Objects being approached grow larger; objects left behind grow smaller. When moving forward, images of nearby objects move across the retina faster than images of distant objects. If a creature rotates in place, the entire visual scene moves across its retina at a rate that is independent of object distance. Much information can be gained from patterns of visual motion, even if no explicit object recognition is performed (Duchon, Kaelbling, and Warren, 1998). Indeed, motion parallax information is immune to camouflage that can defeat even the most sophisticated static pattern recognition scheme when object and background have similar textures. Humans have no difficulty detecting the structure of randomly patterned objects against identically patterned backgrounds from motion cues alone.

Using egomotion-induced optic flow for robot navigation is a computationally demanding sensory task. By its very nature, it must be done in real time. Most object recognition tasks are performed on static images, and often one can tolerate latencies of several seconds. But optic flow is available only while the robot is moving, and relevant information must be extracted in real time and fed back to the motor control system to steer the robot in the right direction. The rate of computation needed depends on the rate of robot motion, but typical real-world situations require times on the order of tens or hundreds of milliseconds.

Optic flow is also computationally demanding because, like other early vision tasks, it involves operations that must be performed identically on every pixel of an image. Local estimates of motion must be laboriously computed before the overall pattern is analyzed. This is a task that is ideally suited for parallel computation.

Measuring optic flow also involves large amounts of data. While audition involves one time-varying signal (two in the case of binaural audition), vision involves many time-varying signals. Rapid navigation requires many frames to be analyzed each second. This can tax even the most sophisticated microprocessor because it must deal with all the signals at once. If we divide the job to many processors, each dealing with one pixel and communicating with its immediate neighbors, the task becomes much easier.

In the past decade, researchers have been endowing mobile robots with biologically inspired (more specifically, insect-inspired) visual systems (Franceschini, Pichon, and Blanes, 1992; Srinivasan, Chahl, and Zhang, 1997; Huber, 1997; Lewis, 1998; Huber, Franz, and Bülthoff, 1999). These efforts have yielded promising results, but many problems

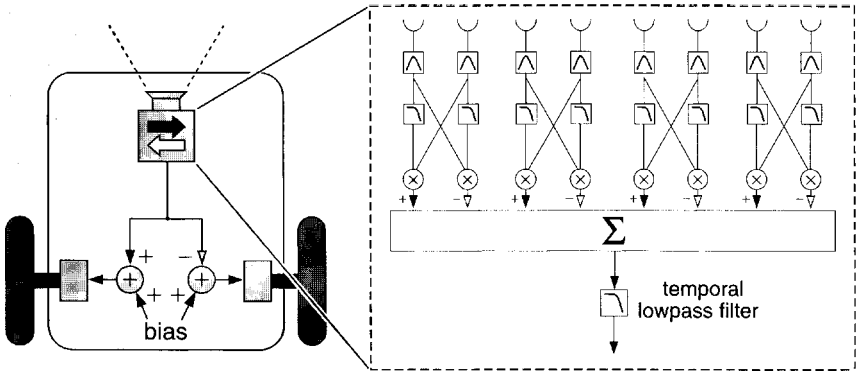


**Figure 2.7** Schematic of our optomotor system. A motion sensor chip is mounted facing forward on a two-wheeled robotic platform. Forward-facing motion sensors are largely blind to optic flow produced by forward translation, so we only measure rotation. The chip's wide-field output is used as an estimate of self rotation, then lowpass filtered ($\tau = 750$ ms) to stabilize the control loop. This signal is added to one motor and subtracted from the other, producing a compensatory rotation. A constant motor bias produces forward translatory motion.

still exist. Vision is a computationally intensive task, so powerful hardware is required to operate in real time. From an algorithmic viewpoint, the structure of visual scenes is often very complex, and it can be difficult to extract relevant information robustly.

As mentioned above, the fly uses visual motion information to stabilize its flight. Mismatch of body components or environmental disturbances may impart rotation on the animal, but sensory feedback is used to produce compensatory torque responses. This sensorimotor feedback is known as the optomotor system, and is one of the best-studied behaviors of the fly (Götz, 1975; Warzecha and Egelhaaf, 1996).

### Hardware Implementation

We constructed a hardware implementation of the optomotor system using a two-wheeled vehicle. We chose to build a physical motor system so we could evaluate our sensor's performance in the real world. Figure 2.7 shows a schematic of the system. Our wide-field motion sensor estimates self-rotation, and this signal is used to produce a compensatory rotation by the drive motors.

We constructed a simple robotic platform on which we mounted the wide-field motion sensor (figure 2.8). The robot had two large wheels driven independently by two DC motors, and a free-turning wheel in the back to maintain balance. Each drive motor was controlled with a pulse-width modulation circuit that varied the duty cycle of a constant-amplitude square-wave voltage. By changing the duty cycle of the waveform, each motor could be driven at any speed up to a maximum.

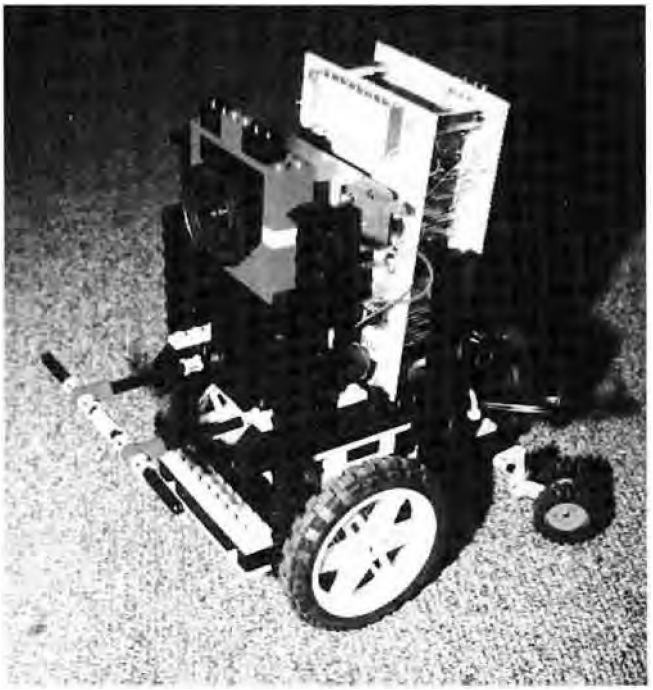

**Figure 2.8** Photograph of the optomotor system. The lens is covering the motion sensor chip. Additional off-chip electronics have been constructed to drive the dc motors. The back wheels turn freely, and merely prevent the robot from falling over. The robot measures 13 cm × 19 cm × 22 cm, and is powered by on-board batteries.

If the motors were driven at different speeds, the robot would drive in a curved trajectory.

A large asymmetry was introduced into the robot's mechanics by connecting the left and right motors to their respective wheels with different gear ratios. The left motor was connected to the left wheel with a 1:5 gear ratio, while the right motor was connected to the right wheel with a 1:1 gear ratio. This caused the robot to drive in tight circles if both motors were driven at the same speed. This asymmetry was made extreme for the sake of experiment, but perfect symmetry is impossible to achieve in any physical robot. While two actuators may match perfectly in simulation, they will never match when built and tested in the real world. This difficulty is especially pronounced in outdoor terrain, where wheels or feet may slip in sand or mud. Legged robots are especially prone to walking in curved lines due to foot slip or terrain differences, even if they have been designed and constructed with high precision.

When open-loop control falls short, we must introduce sensory feedback to further improve performance. Optic flow information has the potential to guide a robot in a straight path, because any deviation

involves a yaw rotation, however slight. If yaw rotation can be estimated from optic flow reliably, we can use this as an error signal in a negative feedback loop in which the motors execute a compensatory rotation to null the sensory error signal.

We constructed a feedback loop of this type using our VLSI wide-field motion sensor. The sensor was mounted facing forward on the robot, oriented so it was sensitive to horizontal motion. We oriented the sensor facing straight ahead because translatory motion by the robot produces little optic flow in the direction of travel, while rotatory (yaw) motion produces uniform optic flow around the visual field parallel to the ground. Thus the optic flow in the forward region of the visual field will be dominated by the rotatory component. The hoverfly *Syritta pipiens* uses this strategy to stabilize its flight. When moving forward, the animal uses optic flow from the forward region of the visual field to estimate self rotation. This agile creature is also capable of flying sideways, and when doing so, it uses optic flow from the lateral visual fields to estimate self-rotation (Collett, 1980). Presumably, it is attempting to measure optic flow in the regions least contaminated with optic flow produced by its own translation.

The output of our motion sensor was a continuous, time-varying voltage. This signal was filtered by a first-order low-pass filter with a time constant of 750 msec. This is a simple model of the relationship between the output of a wide-field motion-sensitive neuron in the fly and the torque response produced by the wings (Egelhaaf, 1987; Warzecha and Egelhaaf, 1996). The filtered output of the motion sensor was added to the left motor command and subtracted from the right motor command (see figure 2.7). This has the effect of adding a rotatory component to the robot's trajectory. In the absence of visual feedback, both motors turned at the same rate (so one wheel turns five times faster than the other). Visual feedback slowed one wheel and sped up the other.

### Robot Experiments

Experiments were performed indoors in our laboratory, but the visual nature of the room was not altered in any way to accommodate the motion sensor. The room was a typical cluttered laboratory environment with many shady areas under tables. The robot's position was recorded 10–20 times per second with a magnetic field tracking system that returned location and orientation in three dimensions (Polhemus, Colchester, VT). The scale of experiments was limited by the range of this system, approximately a 70 cm by 140 cm area for highest accuracy.

The optic flow feedback proved capable of nearly eliminating the effect of physical asymmetry. Figure 2.9 shows one trial without visual feedback. The line shows the robot's path, and the circle indicates the

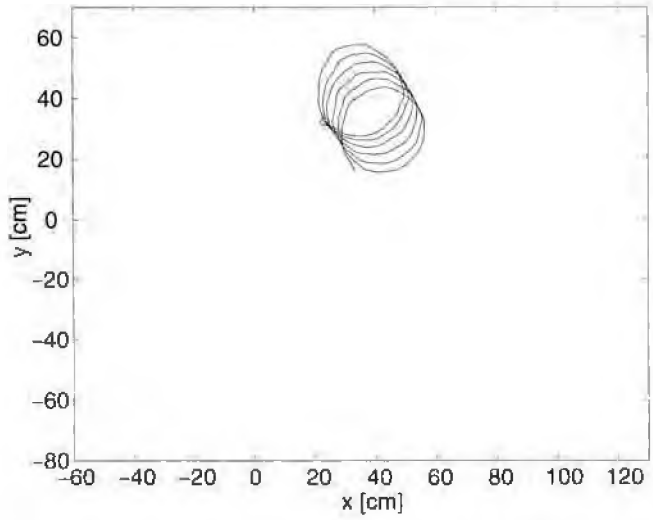


**Figure 2.9** Robot path with no sensory feedback. With the motion sensor disabled, the robot turns in circles due to the asymmetry in its mechanics. The circle denotes the ending location of the robot.

ending position. The robot is turning in tight circles. Figure 2.10 shows ten trials where visual feedback has been enabled. In general, the robot travels in straight lines. We purposely started the robot at different orientations to demonstrate that the sensor works well for general visual scenes around a room. When moving in straight lines, the robot traveled at a speed of approximately 20 cm / sec. Objects and walls were typically 0.2 to 1.5 meters away from the robot, depending on the direction.

The angular velocity of the robot (yaw rate) was computed along each path by differentiating the robot's heading as recorded by the tracking system. Figure 2.11a shows a histogram of angular velocities for the trials without feedback and figure 2.11b shows all ten trials with visual feedback. The mean angular velocity in the open-loop case is $-116°/\text{sec}$, while for the closed-loop case this decreased to $-3.7°/\text{sec}$, an improvement by a factor of 31.

Occasionally, the feedback did fail to keep the course straight. A 45° turn is visible in figure 2.10, most likely caused by the sensor being oriented toward a relatively featureless part of the room, where no motion information is available. A larger field of view would reduce the likelihood of such occurrences. Also, the magnitude of the error depends on the degree of asymmetry in the gear ratios. In a more realistic situation with higher open-loop precision, it is likely that large closed-loop errors would be rare.

Our sensor is small and extremely low power, making it easily adaptable to mobile robot applications. While the power budget on a robot is usually dominated by motors, traditional CCD imagers consume

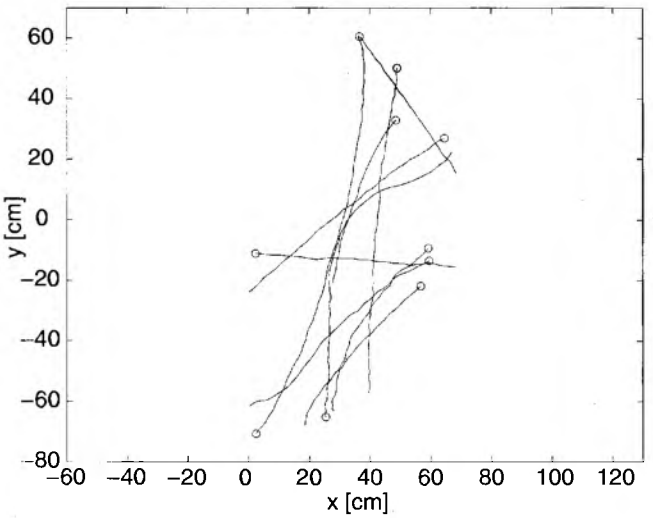


**Figure 2.10** Robot path with sensory feedback. Ten trials are shown where the motion sensor provided optomotor control to straighten the course of the robot despite a 5:1 mechanical asymmetry. Circles denote the ending location of the robot in each trial. The robot was exposed to different visual scenes during the trials.

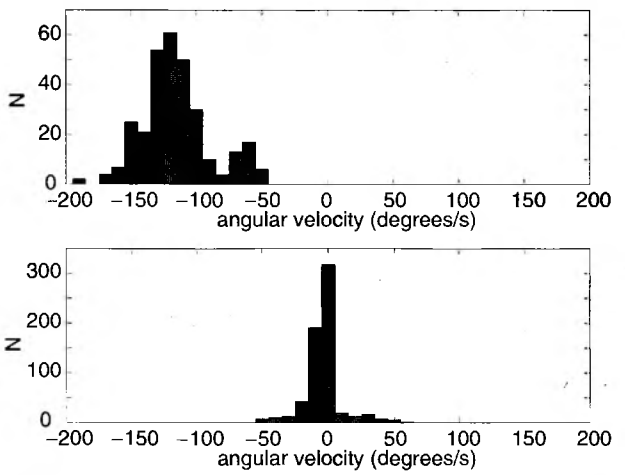


**Figure 2.11** Histogram of angular velocities. (a) No visual feedback. The turning behavior of the robot is obvious. The mean angular velocity was −116 deg/s. (b) Compilation of all ten trials with visual feedback. The mean angular velocity was greatly reduced to −3.7 deg/s.

significant amounts of power, and digital microprocessors or digital signal processors (DSPs) capable of processing real-time video consume even more. On the Sojourner rover of the recent Pathfinder Mars mission, the CCD imagers alone consumed 0.75 W—5% of the total power budget at peak solar cell output. The CPU system consumed an additional 24%, and much of the CPU's time was devoted to processing static images while the rover was not moving (Matthies et al., 1995). By comparison, our EMD array consumed less than 5 μW of power. Traditional imaging and image processing is expensive in terms of time, size, and power. Biologically inspired analog VLSI approaches to this problem can bring down the cost and make robot vision more practical.

## REFERENCES


Adelson, E. H., and Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. *J. Optical Soc. Am.* [*A*] 2: 284–299.

Andreou, A. G., Strohbehn, K., and Jenkins, R. E. (1991). Silicon retina for motion computation. In *Proceedings of the 1991 IEEE International Symposium on Circuits and Systems.* pp. 1373–1376.

Autrum, H. (1958). Electrophysiological analysis of the visual systems in insects. *Exp. Cell Res. Supplement* 5: 426–439.

Borst, A. (1990). How do flies land? *BioScience* 40: 292–299.

Borst, A., and Bahde, S. (1988). Spatio-temporal integration of motion. *Naturwissenschaften* 75: 265–267.

Borst, A., and Egelhaaf, M. (1989). Principles of visual motion detection. *Trends Neurosci.* 12: 297–306.

Borst, A., and Egelhaaf, M. (1990). Direction selectivity of blowfly motion-sensitive neurons is computed in a two-stage process. *Proc. Natl. Acad. Sci. U.S.A.* 87: 9363–9367.

Clifford, C. W. G., Ibbotson, M. R., and Langley, K. (1997). An adaptive Reichardt detector model of motion adaptation in insects and mammals. *Vis. Neurosci.* 14: 741–749.

Collett, T. S. (1980). Some operating rules for the optomotor system of a hoverfly during voluntary flight. *J. Comp. Physiol.* [*A*] 138: 271–282.

Delbrück, T., and Mead, C. A. (1996). Analog VLSI phototransduction by continuous-time, adaptive, logarithmic photoreceptor circuits. CNS Memo 30, California Institute of Technology.

DeVoe, R. D. (1980). Movement sensitivities of cells in the fly's medulla. *J. Comp. Physiol.* [*A*] 138: 93–119.

DeVoe, R. D., and Ockleford, E. M. (1976). Intracellular responses from cells in the medulla of the fly, *Calliphora erythrocephala*. *Biol. Cybern.* 23: 13–24.

Duchon, A. P., Kaelbling, L. P., and Warren, W. H. (1998). Ecological robotics. *Adaptive Behav.* 6: 473–507.

Egelhaaf, M. (1985). On the neuronal basis of figure-ground discrimination by relative motion in the visual system of the fly: II. Figure-detection cells, a new class of visual interneurones. *Biol. Cybern.* 52: 195–209.

Egelhaaf, M. (1987). Dynamic properties of two control-systems underlying visually guided turning in house-flies. *J. Comp. Physiol.* [*A*] 161: 777–783.

Egelhaaf, M., and Borst, A. (1989). Transient and steady-state response properties of movement detectors. *J. Optical Soc. Am.* [*A*] 6: 116–127.

Egelhaaf, M., and Borst, A. (1993). A look into the cockpit of the fly: Visual orientation, algorithms, and identified neurons. *J. Neurosci.* 13: 4563–4574.

Egelhaaf, M., Borst, A., and Reichardt, W. (1989). Computational structure of a biological motion-detection system as revealed by local detector analysis in the fly's nervous system. *J. Optical Soc. Am.* [*A*] 6: 1070–1087.

Egelhaaf, M., Hausen, K., Reichardt, W., and Wehrhahn, C. (1988). Visual course control in flies relies on neuronal computation of object and background motion. *Trends Neurosci.* 11: 351–358.

Franceschini, N., Pichon, J. M., and Blanes, C. (1992). From insect vision to robot vision. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 337: 283–294.

Geiger, G., and Nässel, D. R. (1981). Visual orientation behaviour of flies after selective laser beam ablation of interneurones. *Nature* 293: 398–399.

Geiger, G., and Nässel, D. R. (1982). Visual processing of moving single objects and wide-field patterns in flies: Behavioural analysis after laser-surgical removal of interneurons. *Biol. Cybern.* 44: 141–149.

Geiger, G., and Poggio, T. (1981). Asymptotic oscillations in the tracking behaviour of the fly *Musca domestica*. *Biol. Cybern.* 41: 197–201.

Gibson, J. J. (1950). *The Perception of the Visual World*. Boston: Houghton-Mifflin.

Götz, K. G. (1975). The optomotor equilibrium of the *Drosophila* navigation system. *J. Comp. Physiol.* [*A*] 99: 187–210.

Haag, J., and Borst, A. (1997). Encoding of visual motion information and reliability in spiking and graded potential neurons. *J. Neurosci.* 17: 4809–4819.

Harrison, R. R., and Koch, C. (1998). An analog (VLSI) model of the fly elementary motion detector. In M. I. Jordan, M. J. Kearns, and S. A. Solla (eds.), *Advances in Neural Information Processing Systems 10*. Cambridge: MIT Press, pp. 880–886.

Harrison, R. R., and Koch, C. (2000). A robust analog VLSI Reichardt motion sensor. *Analog Integrated Circuits and Signal Processing* (in press).

Hassenstein, B., and Reichardt, W. (1956). Systemtheoretische Analyse der Zeit-, Reihenfolgen-, und Vorzeichenauswertung bei der Bewegungsperzeption des Rüsselkäfers *Chlorophanus*. *Z. Naturforch.* 11b: 513–524.

Hausen, K. (1982a). Motion sensitive interneurons in the optomotor system of the fly: I. The horizontal cells: Structure and signals. *Biol. Cybern.* 45: 143–156.

Hausen, K. (1982b). Motion sensitive interneurons in the optomotor system of the fly: II. The horizontal cells: Receptive field organization and response characteristics. *Biol. Cybern.* 46: 67–79.

Hausen, K. (1984). The lobula-complex of the fly: Structure, function, and significance in behaviour. In M. A. Ali (ed.), *Photoreception and Vision in Invertebrates*. New York: Plenum, pp. 523–559.

Hausen, K., and Egelhaaf, M. (1989). Neural mechanisms of visual course control in insects. In D. G. Stavenga and R. C. Hardie (eds.), *Facets of Vision*. Berlin: Springer-Verlag.

Hausen, K., and Wehrhahn, C. (1990). Neural circuits mediating visual flight control in flies: II. Separation of two control-systems by microsurgical brain-lesions. *J. Neurosci.* 10: 351–360.

Heisenberg, M., and Buchner, E. (1977). The rôle of retinula cell types in visual behavior of *Drosophila melanogaster*. *J. Comp. Physiol.* [*A*] 117: 127–162.

Hengstenberg, R. (1982). Common visual response properties of giant vertical cells in the lobula plate of the blowfly *Calliphora*. *J. Comp. Physiol.* [*A*] 149: 179–193.

Huber, S. A. (1997). Studies of the Visual Orientation Behavior in Flies Using the Artificial Life Approach. Ph.D. thesis, Universität Tübingen, Germany.

Huber, S. A., Franz, M. O., and Bülthoff, H. H. (1999). On robots and flies: Modeling the visual orientation behavior of flies. *Robotics Autonomous Syst.* 29: 227–242.

Jiang, H.-C., and Wu, C.-Y. (1999). A 2-D velocity- and direction-selective sensor with BJT-based silicon retina and temporal zero-crossing detector. *IEEE J. Solid State Circuits* 34: 241–247.

Kimmerle, B., Warzecha, A.-K., and Egelhaaf, M. (1997). Object detection in the fly during simulated translatory flight. *J. Comp. Physiol.* [*A*] 181: 247–255.

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

Land, M. F. (1997). Visual acuity in insects. *Annu. Rev. Entomol.* 42: 147–177.

Laughlin, S. B. (1994). Matching coding, circuits, cells, and molecules to signals—general principles of retinal design in the fly's eye. *Prog. Ret. Eye Res.* 13: 165–196.

Lewis, M. A. (1998). Visual navigation in a robot using zig-zag behavior. In M. I. Jordan, M. J. Kearns, and S. A. Solla (eds.), *Advances in Neural Information Processing Systems 10*. Cambridge: MIT Press, pp. 822–828.

Liu, S.-C. (1997). Neuromorphic Models of Visual and Motion Processing in the Fly Visual System. Ph.D. thesis, California Institute of Technology.

Matthies, L., Gat, E., Harrison, R., Wilcox, B., Volpe, R., and Litwin, T. (1995). Mars microrover navigation: Performance evaluation and enhancement. *Autonomous Robots* 2: 291–311.

Mead, C. (1989). *Analog VLSI and Neural Systems*. Menlo Park, Calif.: Addison-Wesley.

Moini, A., Bouzerdoum, A., Eshraghian, K., Yakovleff, A., Nguyen, X. T., Blanksby, A., Beare, R., Abbott, D., and Bogner, R. E. (1997). An insect vision-based motion detection chip. *IEEE J. Solid-State Circuits* 32: 279–284.

Pichon, J.-M., Blanes, C., and Franceschini, N. (1989). Visual guidance of a mobile robot equipped with a network of self-motion sensors. In *SPIE Mobile Robots IV*, vol. 1195, pp. 44–53.

Poggio, T., and Reichardt, W. (1981). Visual fixation and tracking by flies: Mathematical properties of simple control systems. *Biol. Cybern.* 40: 101–112.

Reichardt, W., and Egelhaaf, M. (1988). Properties of individual movement detectors as derived from behavioural experiments on the visual system of the fly. *Biol. Cybern.* 58: 287–294.

Reichardt, W., and Poggio, T. (1976). Visual control of orientation behaviour in the fly. *Q. Rev. Biophys.* 9: 311–375.

Single, S., and Borst, A. (1998). Dendritic integration and its role in computing image velocity. *Science* 281: 1848–1850.

Srinivasan, M. V., Chahl, J. S., and Zhang, S. W. (1997). Robot navigation by visual dead-reckoning: Inspiration from insects. *Int. J. Pattern Recognition Artif. Intell.* 11: 35–47.

Strausfeld, N. J. (1976). *Atlas of an Insect Brain.* Berlin: Springer.

van Santen, J. P. H., and Sperling, G. (1985). Elaborated Reichardt detectors. *J. Optical Soc. Am.* [*A*] 2: 300–321.

Wagner, H. (1986). Flight performance and visual control of flight of the free-flying housefly *Musca domestica* L.: II. Pursuit of targets. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 312: 553–579.

Wandell, B. A. (1995). *Foundations of Vision.* Sunderland, Mass.: Sinauer Associates.

Warzecha, A.-K., and Egelhaaf, M. (1996). Instrinsic properties of biological motion detectors prevent the optomotor control system from getting unstable. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 351: 1579–1591.

Watson, A. B., and Ahumada, A. J. (1985). Model of human visual-motion sensing. *J. Optical Soc. Am.* [*A*] 2: 322–342.

Weckström, M., Juusola, M., and Laughlin, S. B. (1992). Presynaptic enhancement of signal transients in photoreceptor terminals in the compound eye. *Proc. R. Soc. Lond.* [*B*] 250: 83–89.

Wolf, R., and Heisenberg, M. (1990). Visual control of straight flight in *Drosophila melanogaster*. *J. Comp. Physiol.* [*A*] 167: 269–283.

Zanker, J. M. (1990). On the directional sensitivity of motion detectors. *Biol. Cybern.* 62: 177–183.