Biol. Cybern. 85, 449–457 (2001)

Biological Cybernetics

Fly-like visuomotor responses of a robot using aVLSI motion-sensitive chips

Shih-Chii Liu, Alessandro Usseglio-Viretta

Institute of Neuroinformatics ETH/Universität Zürich, Winterthurerstrasse 190, 8057 Zürich, Switzerland

Received: 4 October 1999 / Accepted in revised form: 26 April 2001

Abstract. We explore the use of continuous-time analog very-large-scale-integrated (aVLSI) neuromorphic visual preprocessors together with a robotic platform in generating bio-inspired behaviors. Both the aVLSI motion sensors and the robot behaviors described in this work are inspired by the motion computation in the fly visual system and two different fly behaviors. In most robotic systems, the visual information comes from serially scanned imagers. This restricts the form of computation of the visual image and slows down the input rate to the controller system of the robot, hence increasing the reaction time of the robot. These aVLSI neuromorphic sensors reduce the computational load and power consumption of the robot, thus making it possible to explore continuous-time visuomotor control systems that react in real-time to the environment. The motion sensor provides two outputs: one for the preferred direction and the other for the null direction. These motion outputs are created from the aggregation of six elementary motion detectors that implement a variant of Reichardt's correlation algorithm. The four analog continuous-time outputs from the motion chips go to the control system on the robot which generates a mixture of two behaviors – course stabilization and fixation – from the outputs of these sensors. Since there are only four outputs, the amount of information transmitted to the controller is reduced (as compared to using a CCD sensor), and the reaction time of the robot is greatly decreased. In this work, the robot samples the motion sensors every 3.3 ms during the behavioral experiments.

1 Introduction

Many robotic systems are bulky and slow, require massive computing power, and expend a lot of energy. In addition, since silicon imagers are frequently used as the front end, the data acquisition rate of the imagers can compromise the response time of the sensorimotor control loop of the robot. Recent research has centered on sensorimotor control loop problems that use small, low-power robots with onboard sensors and some onboard processing (Webb 1995; Srinivasan et al. 1997; Huber 1999). This research has benefitted greatly from both behavioral and sensory physiological studies on data collected from animals in visual and navigation tasks (Ronacher and Wehner 1995; Srinivasan et al. 1997). Flies and bees are prototypes for examples of sensorimotor systems that work robustly in natural environments. The first attempt at solving a real-time visuomotor control problem using a robot and visual sensors that were inspired by the fly system was developed by Franceschini et al. (1992). The robot had 100 optical lenses and photodiodes covering an angle of about 360°. However, this system was quite bulky and weighed about 10 kg. With present-day technology, the size and weight of these robots can be reduced with the use of smaller onboard sensors, for example, a panoramic camera (Srinivasan et al. 1997). The speed of the system can be increased through the availability of fast onboard processors.

1.1 Neuromorphic sensors

The data acquisition rate of silicon imagers used on robotic systems is around 25–30 frames/s for a 640×480 sized image. However, the acquisition time is dominated by the onboard computational time of the visuomotor control loop. The speed of the robot is compromised by the image acquistion rate and the necessity to do frame-based computation.

One solution to this bandwidth dilemma is the use of sensors that are based on principles of neuromorphic engineering (Mead 1989). Electronic systems that model the biological function and structure of neuronal processing in invertebrates and vertebrates are developed partly to enhance the present state of technology in the

Correspondence to: S-C Liu
(Tel.: +41-1-6353047, Fax: +41-1-6353053
e-mail: shih@ini.phys.ethz.ch)

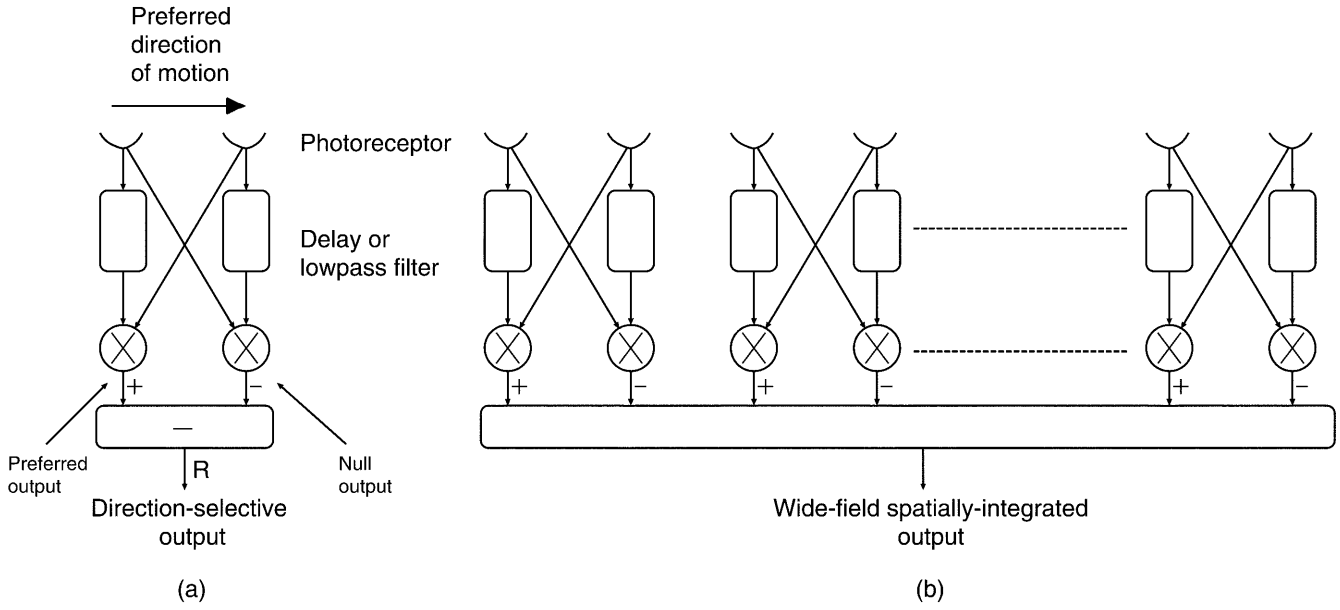

Fig. 1a,b. Block diagram of Hassenstein–Reichardt's correlation model of motion computation in insects. **a** Model of a single EMD. We have arbitrarily called the direction of a stimulus moving from left to right the preferred direction of the EMD. **b** Model of a wide-field direction-selective cell in the fly. The output is integrated from all the EMDs. This is also the model for the motion sensor used in this work

industry and partly as a form of modeling biology by building silicon circuits and systems that have to deal with problems of noise and mismatch. There has been substantial work into developing neuromorphic analog very-large-scale-integrated (aVLSI) motion sensors, since motion information is important in many different tasks (Benson and Delbrück 1992; Delbrück 1993; Kramer et al. 1995; Sarpeshkar et al. 1996; Etienne-Cummings et al. 1999), and it is natural to compute motion with continuous-time circuits. Some of these motion chips are based on a model of motion processing in fly vision (Andreou and Strohbehn 1990; Sarpeshkar et al. 1993; Moini et al. 1997; Harrison and Koch 1998).

In this work, we explore the use of two aVLSI motion sensors that provide global motion information to a controller system that generates two possible behaviors on the robot. Each motion sensor models the motion pathway (from the receptors to the wide-field, direction-selective cells) of the fly on a single chip. The two sensor outputs indicate whether a stimulus is moving in the preferred direction or in the null direction. As a result, only two outputs from each chip need to be sampled by the microcontroller, and the speed of the sensorimotor loop on the robot can be increased significantly. Depending on the motion outputs, the robot exhibits one of two behaviors: optomotor response and fixation. These behaviors are inspired by the behaviors of real flies.

In Sect. 2, we describe the architecture of the aVLSI motion sensor, and in Sect. 3, we discuss some of the circuit techniques used for computing a robust motion output. We show measured data from a motion chip fabricated in 1.2-μm CMOS technology in Sect. 4. Finally, we describe a controller system that generates the two behaviors using the outputs of two motion sensors mounted on the Koala (our robotic platform), and show the improvement in the sampling rate of the sensors.

2 The aVLSI motion model

The aVLSI motion sensor consists of six elementary motion detectors (EMDs) – each EMD implements a variant of the Hassenstein–Reichardt correlation algorithm (Hassenstein and Reichardt 1956; Reichardt 1961) as shown in Fig. 1a. There is no motion computation between the pairs of pixels of the six EMDs in this aVLSI sensor due to a design oversight. This error has been corrected in a new circuit implementation (Liu 2000). The outputs of the individual EMDs are summed together to produce a global motion measure, as shown in Fig. 1b. An EMD receives inputs from two neighboring pixels (or photoreceptors). One of the inputs goes through a low-pass filter. This filtered signal is multiplied with the input at the adjacent pixel. The output of the left multiplier of the EMD is sensitive to stimuli moving from left to right; the output of the right multiplier is sensitive to stimuli moving from right to left. The two outputs are subtracted to remove the common-mode dependence on the input.

The architecture of the silicon EMD (shown in Fig. 2) closely follows the anatomical layout of the fly visual system. The first layer of the motion chip models the retina layer; each photoreceptor transduces the incoming light into a voltage. The circuit has a high transient gain to changes in intensity and a low DC gain (Delbrück 1994; Liu 1999). The circuit adapts over six decades of background intensity. Because of this local gain property, the circuit primarily codes contrast information.

The second layer models the responses of the cells in the laminar layer of the fly retina. Recordings from the large monopolar cells (LMC) (Laughlin 1993; Juusola et al. 1995) in this layer show that the LMC responses are amplified versions of the receptor outputs. The DC

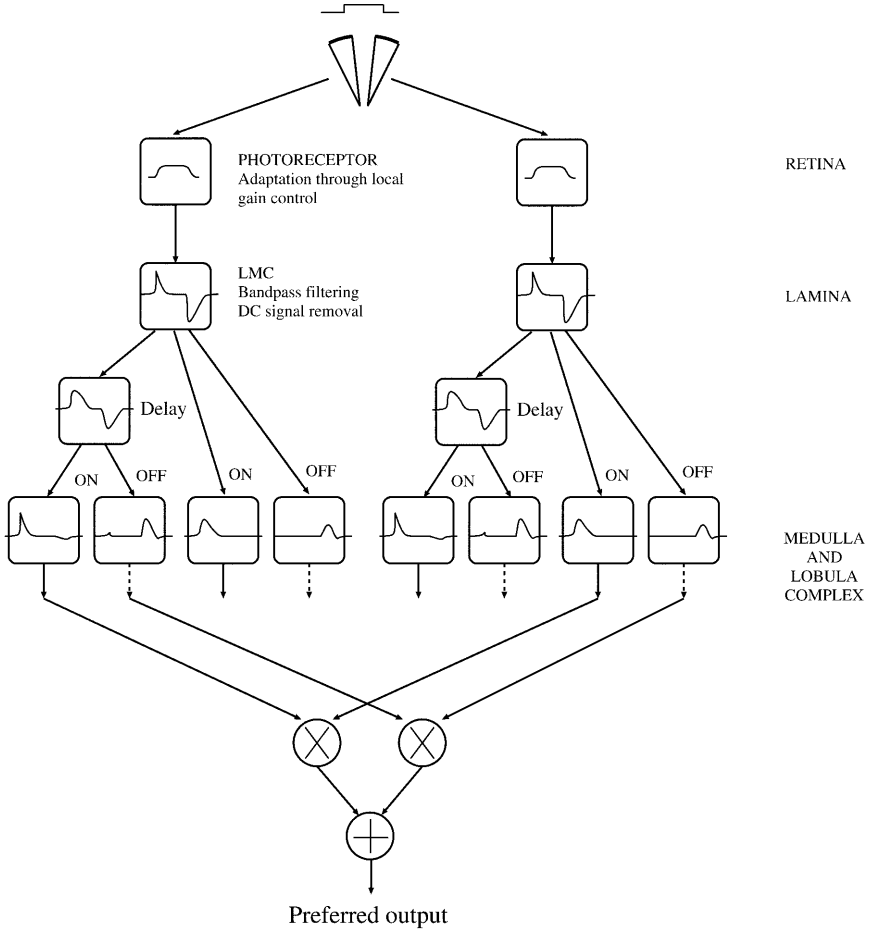

Fig. 2. Block diagram of the architecture of the silicon EMD for computing motion. The input goes through a photoreceptor circuit with local gain control. The receptor output is then band-pass filtered by the LMC circuit. The output of the LMC circuit splits into two pathways; in one pathway, the signal goes through a low-pass filter. From the filtered and unfiltered signals, we generate ON and OFF transient currents. The ON and OFF filtered signals in each pixel are correlated (or multiplied) with the unfiltered signals from the adjacent pixel. The two outputs shown in the figure are sensitive to a stimulus moving in one direction (from left to right). This figure corrects an error made in similar figures that were published in (Liu and Usseglio-Viretta (2000) and Indiveri and Douglas (2000). In these two papers, the output that was labeled with a '-' in each figure should have been labeled with a '+'

response of the cells is approximately constant, unlike the DC response of the receptors, which depends on the background intensity. In our silicon LMC, the receptor output is AC-coupled into an inverting differentiator circuit with a fixed transient gain. Only transient changes of the receptor output are amplified by this fixed gain.

The third layer models hypothetical cells in the medulla area or lobula complex. Here we split the LMC signal into two pathways. In one pathway, the signal goes through a low-pass filter. The original LMC signal and the filtered signal are then processed by a nonlinear differentiator circuit (Kramer et al. 1997) that generates both an ON transient current and an OFF transient current from temporal changes in the two signals. The resulting four currents (ON, delayed ON, OFF, delayed OFF) are correlated together using a simple current correlator (Delbrück 1991) (analogous to the multiplier in the Reichardt model). The ON transient current at each pixel is correlated with the delayed-ON transient current from its neighboring pixel and the same computation is performed with the OFF currents. This splitting of ON and OFF pathways in the basic process of motion detection was demonstrated via electrophysiological recordings from a wide-field, direction-selective cell in the housefly (Franceschini et al. 1989). The output currents from all four correlators are then used to charge or discharge a capacitor on a node, V_{mot}, with a leak conductance to a reference voltage. The output, V_{mot}, corresponds to the responses of the wide-field direction-selective cells in the lobula plate of the fly. The details of the circuits in this chip are discussed in Liu (2000).

3 Offsets

Silicon hardware – probably like biological neural hardware – has mismatches between neighboring pixels or cells. Offsets make robust computation difficult. The coefficient of variation (CV) – defined as the standard deviation divided by the mean – of the peak responses of blowfly photoreceptors (Juusola 1993) to a signal contrast (Michelson contrast) of 0.359 was on the order of 5 percent. The CV measured from a silicon retina (Liu 1997) fabricated in 2 μm CMOS ORBIT technology is about 6–10%. These data were obtained from the peak-to-peak responses of the photoreceptor output of the individual imager pixels in response to an LED that was modulated by a 10-Hz square wave. The effects of offsets between pixels need to be considered if comparisons or correlations between pixels are part of the circuit computation.

We have taken steps to reduce the effect of mismatches by incorporating band-pass filtering at various stages of processing in this system. Some of the methods we used to increase the robustness of the computation are as follows:

1. We band-pass filtered the signal from the photoreceptor going into the LMC circuit so that there is no common-mode dependence on the LMC output.
2. The DC level of the input from the LMC circuit to the nonlinear differentiator circuit affects the transient dynamics of the generation of the ON and OFF currents. Since the DC voltage of the LMC input to the nonlinear differentiator circuit is constant, the dynamics are not affected by the background light intensity.
3. We have chosen to use transient current inputs instead of voltage inputs to a correlator circuit so that we would not have to deal with mismatches in the DC input levels at the multiplier. A previous hardware implementation of the correlation model where voltage outputs from neighboring pixels are fed into a multiplier circuit failed to provide a robust direction-sensitive output. The signal was not robust because the input DC offset between the pixels creates different outputs at the multiplier. Instead of using a voltage-mode multiplier, we used transient currents as inputs to a current correlator circuit. The use of transient currents ensures that there will be no output current (coding zero motion) if there is no input change at the receptor output. The offsets in the voltage multiplier could be reduced via layout and sizing techniques or nulled using tunneling-injection structures (Diorio et al. 1999) or UV techniques. These solutions require additional circuitry and hence area.

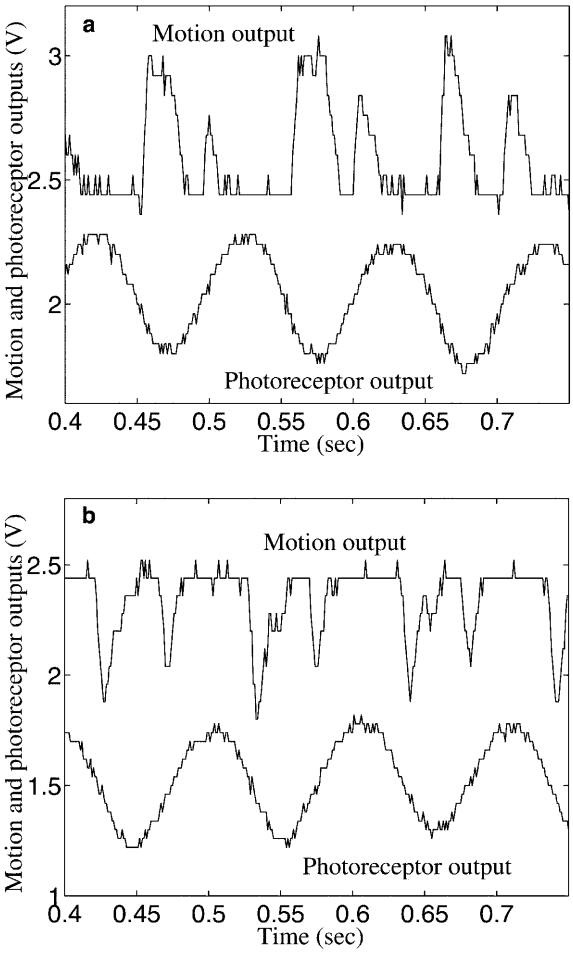

Fig. 3a,b. Responses of the photoreceptor and the silicon EMD to a sinusoidal grating of contrast 0.2 moving at a temporal frequency of 10 Hz. Notice the frequency doubling in the EMD output from the correlation circuit. **a** Responses to the grating moving in the preferred direction. The photoreceptor output has been shifted down by 0.1 V for ease of comparison. **b** Responses to the grating moving in the null direction. The photoreceptor output has been shifted down by 0.6 V

4 Responses from the aVLSI motion sensor

The motion output, V_{mot}, of the aVLSI motion sensor sits at a reference voltage of 2.5 V when there is no motion. If the stimulus moves in the preferred direction, the currents from the corresponding correlators or multipliers in Fig. 2 charge the output node, V_{mot}, above 2.5 V; if the stimulus moves in the null direction, the currents from the multipliers discharge V_{mot} below 2.5 V. Notice that we have arbitrarily called the direction of the moving stimulus as the preferred direction if the motion output is above the reference voltage. The results shown in Figs. 3–6 are obtained for a single EMD on the aVLSI chip. The temporal responses of a photoreceptor and an EMD to a sinusoidal grating moving in the preferred direction are shown in Fig. 3a, and in the null direction are shown in Fig. 3b. The asymmetry in the responses of the motion output between the ON and OFF pathways occurs in the circuit that generates the transient currents from the LMC outputs.

We obtained the temporal frequency response of the aVLSI EMD (Fig. 4a) by moving a sinusoidal grating with a fixed spatial frequency in front of the sensor. The shape of the response is similar to that measured in the so called HSE cell of the fly (Fig. 4b), which is one of the wide-field direction-selective cells (Hausen 1982b). The dynamic range of the input frequency is in the range 1–20 Hz. The curves above the baseline of 2.5 V correspond to the stimulus moving in the preferred direction, and the curves below this baseline are results obtained from using a stimulus moving in the null direction. The three curves correspond to three different contrasts of the stimulus. The dependence of the motion sensor on the input contrast is dissimilar from that of the HS cell since the output of the motion sensor increases with the signal contrast. The output of the HS cell saturates around a contrast of about 0.3 (Egelhaaf and Borst 1989). This discrepancy is being addressed in a new version of the circuit.

The temporal frequency where the EMD response peaks in Fig. 4 can be controlled by the time constant of the low-pass filter. We can adjust this time constant by changing the voltage of a control signal to the filter circuit. The three curves in Fig. 5 correspond to three different delay bias settings 0.44 V, 0.45 V, and 0.46 V.

As shown by Hausen (1982a,b), the motion responses of the HS cells depend on the orientation of the stimulus moving in the preferred direction. These responses approximate a cosine tuning curve. We repeated the same

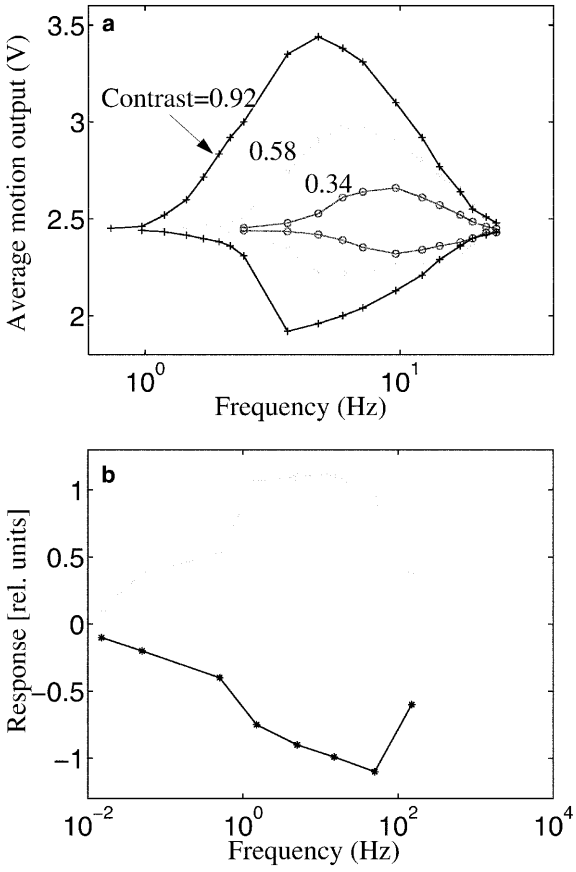

Fig. 4a,b. Frequency response of silicon EMD and a fly HSE cell. **a** Frequency response of the silicon EMD circuit. The three curves correspond to a sinusoid at three different contrasts: 0.34, 0.58, and 0.92. The contrast here is defined as the Michelson contrast, $\frac{I_{max}-I_{min}}{I_{max}+I_{min}}$. The curves above the baseline of 2.5 V are obtained for the sinusoidal grating moving in the preferred direction and the curves below the baseline are obtained for the grating moving in the null direction. **b** Data from a HSE cell replotted from Fig. 5c in Hausen (1982b)

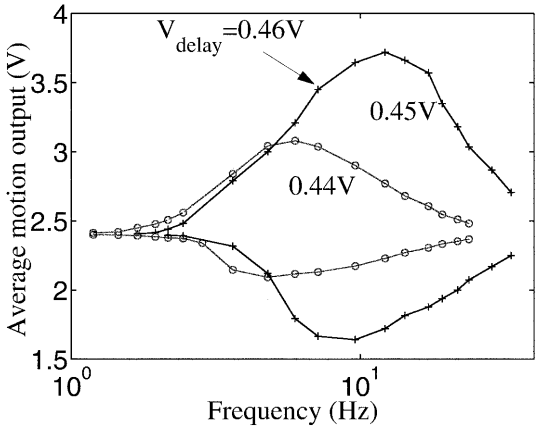

Fig. 5. Frequency response of the silicon EMD for three different time constants of the low-pass filter. The time constant is set by a bias voltage, V_{delay}. A higher V_{delay} value means a lower filter time constant. Average response of the EMD to a stimulus moving in the preferred and null directions at three different delay settings: 0.44 V, 0.45 V, and 0.46 V

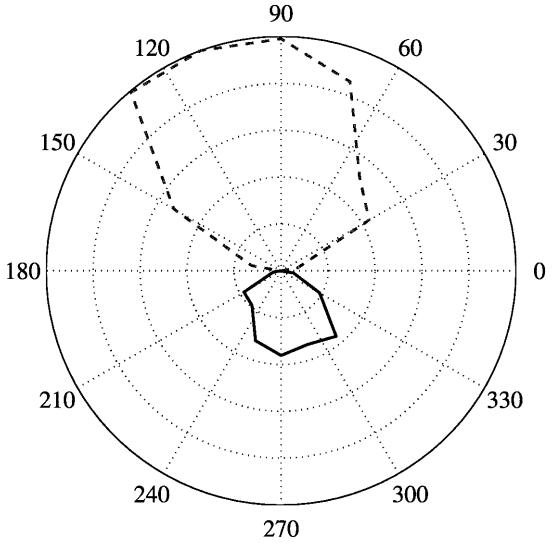

Fig. 6. Response to a sinusoidal grating at different orientations moving in the preferred direction. The contrast of the grating is 0.7 and the temporal frequency is 10 Hz. The response to the stimulus moving in the preferred direction (*dashed curve*) and in the null direction (*solid curve*) is plotted in polar coordinates. The average response of the circuit is 0.296 V

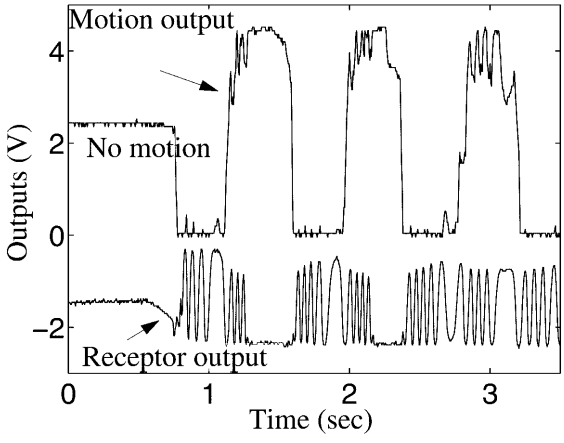

Fig. 7. Global motion output as computed from the spatial integration of the six EMD outputs. The photoreceptor signal has been shifted down by 3 V

experiments with our motion chip; the polar plot of the tuning curve from the EMD is shown in Fig. 6.

Finally, the temporal variation of an EMD in response to a sinusoidal grating can be reduced by adding the outputs from several EMDs, as has been shown physiologically (Single and Borst 1998). In Fig. 7, we show the global output of the motion sensor from the spatial summation of the six EMDs in response to a moving sinusoidal grating. The baseline of 2.5 V corresponds to zero measured motion. At about 0.7 s, the grating moves in the null direction and the output decreases below 2.5 V. At about 1.2 s, the grating moves in the preferred direction leading to an increase in the sensor output above 2.5 V. The temporal variation in the output is reduced due to the addition of the currents

from the EMDs. However, in this case the output became saturated because of the high contrast of the grating and the linear operation used in generating the global output. A new version of this sensor (Liu 2000) aggregates the outputs of the EMDs in a nonlinear way, as proposed by Borst et al. (1995).

5 Sensorimotor system

The optomotor response (or course stabilization) behavior and the fixation behavior are two well-known behaviors of the fly. Experiments with tethered flies show that the optomotor behavior is triggered whenever the background viewed by the fly moves horizontally; the fly's response is a yaw torque that tends to decrease the speed of the retinal image (Reichardt and Poggio 1979). During free flight, this kind of response proves to be essential for correcting the flight trajectory against disturbances caused by to air movements.

The fixation behavior is evoked under controlled conditions by showing the tethered fly a stimulus which consists, for example, of a random textured stripe against a similar textured background. If the stripe oscillates out of phase with the background in the visual field of one of the two compound eyes, the fly turns towards the stimulus.

Researchers have posited the existence of two control systems with different spatial and temporal properties in controlling the yaw torque of a fly in flight (Wehrhahn et al. 1982; Egelhaaf 1987; Hausen and Egelhaaf 1989; Egelhaaf and Borst 1993). The control systems consist of a "large-field" system that controls behaviors that depend on the measurement of global motion while the "small-field" system controls behaviors that depend on measurements of the properties of a small object within the visual image (Hausen 1982a,b; Egelhaaf and Borst 1993).

We reproduced two simplified versions of these two behaviors with our aVLSI motion sensors and the Koala. The aVLSI motion sensor described here only measures the global motion of the visual field. Hence, one of the behaviors displayed by the robot (the fixation behavior) is not as sophisticated as that of the fly. In our fixation behavior, the robot turns towards the direction of the sensor that measures greater motion energy. No local information is available for the robot to perform tracking. This information can only be supplied by a "small-field" system.

5.1 Experimental setup

The motion sensors equipped with a lens (focal length 3.5 mm) are mounted on two separate boards along with a 9 V battery and some signal-shaping (peak detection, integration, and subtraction) logic. The layout of this motion sensor is shown in Fig. 8a. Each board provides two output signals (preferred direction and null direction) which swing from ground (zero motion) to 2.5 V. The boards are mounted onto a Koala (shown in Fig. 8b), a

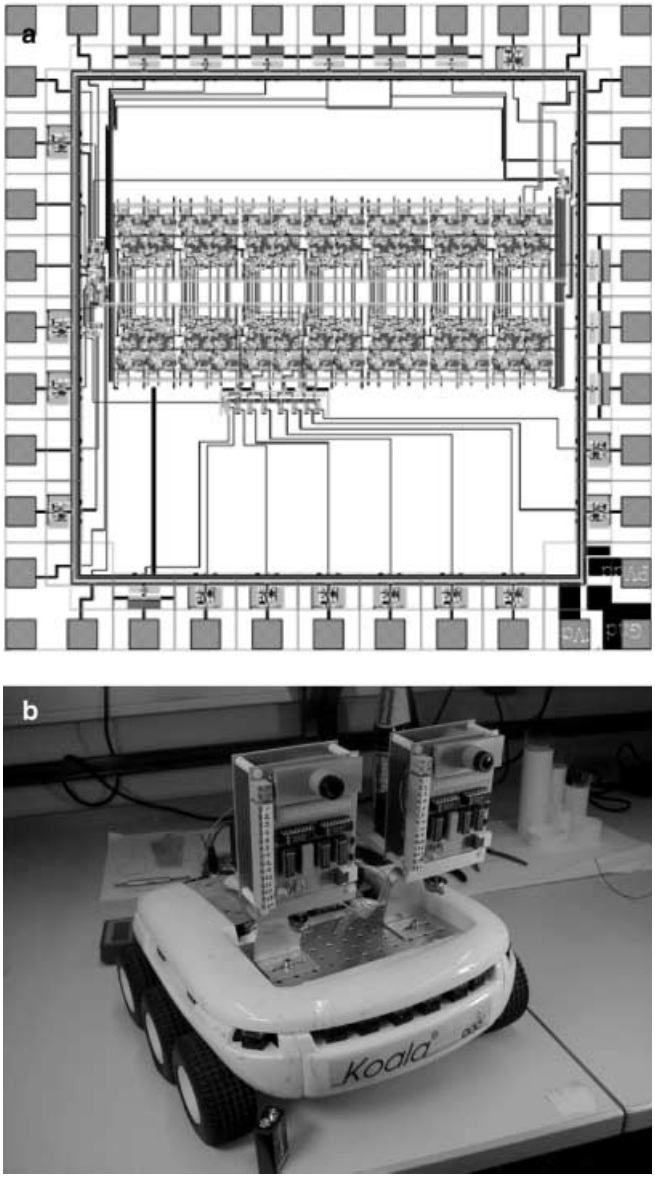

Fig. 8a,b. Koala with the two motion sensors. **a** Layout of the motion sensor used on the Koala. The chip is 2.2 mm × 2.2 mm and fabricated in a 1.2-μm CMOS process. There are twelve photoreceptors corresponding to six EMD circuits. The photodiode area is approximately 1170 μm^2. The interreceptor spacing is 77 μm; the fill factor is 5%. **b** The Koala equipped with two boards; each board has an aVLSI chip, some digital logic, and a 9 V battery

multipurpose robotic platform manufactured by the Swiss company K-Team (http://www.k-team.com). The robot has a number of features which make it suitable for the testing of sensors and sensorimotor control. It carries a 16-MHz Motorola 68331 microcontroller which can control the wheel motors, read data from 16 infrared sensors and six 10-bit analog-to-digital converters (ADCs), and write to a number of digital outputs according to the program stored in its random-access memory, or following commands from an external computer through a serial link. The outputs of the sensors are digitized by the Koala ADC and processed in real-time by the Koala's microcontroller. The experi-

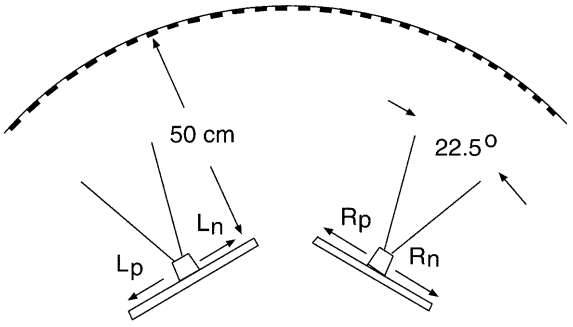

Fig. 9. Top view of the setup used for the optomotor response experiment. The background consists a striped grating with an angular wavelength of approximately 6.8°. The field of view of the sensor with a lens of focal length 3.5 mm is 22.5°. The preferred and null directions (L_p and L_n) for the left sensor and the corresponding directions (R_p and R_n) for the right sensor are shown in the figure

mental setup for one of the behaviors described below is shown in Fig. 9.

5.2 Experimental results

We obtained the two behaviors by using linear combinations of rectified and smoothed versions of the four outputs of the two aVLSI sensors: L_p, L_n, R_p, and R_n (corresponding to the preferred and null direction outputs of the left and right chips, respectively see Fig. 9). Signals from the chips were normalized at the beginning of each run in order to compensate for differences in the motion outputs over runs. The visuomotor control system on our Koala is described by the matrix **C**, see Fig. 10) given below:

$$\mathbf{C} = \begin{bmatrix} 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \end{bmatrix} \begin{bmatrix} L_p \\ R_p \\ L_n \\ R_n \end{bmatrix}$$

In this formulation, equal weight is given to the four motion signals: L_p, L_n, R_p, and R_n. The first row of **C** corresponds to the linear combination used for the course stabilization and the second row corresponds to the combination used for fixation. The criterion for activating the course stabilization behavior is that we measure similar motion in both sensors, while the criterion for activating the fixation behavior is that we measure different motion in each sensor.

The results of the matrix computations for both behaviors were digitally low-pass filtered. The output for the stabilizing behavior was integrated with a time constant of 16 ms, while the output for the fixation behavior was integrated with a time constant that is 3.5 times longer (56 ms). The temporal integration of the outputs is needed for the visuomotor loop to be stable because we do not have enough pixels on each sensor to remove the dependence of the sensor output on the spatial structure of the pattern. In the case of a sinusoidal grating pattern, there is a temporal variation in the motion output in response to the grating. Hence we cannot use the sensor outputs directly to drive the wheel motors.

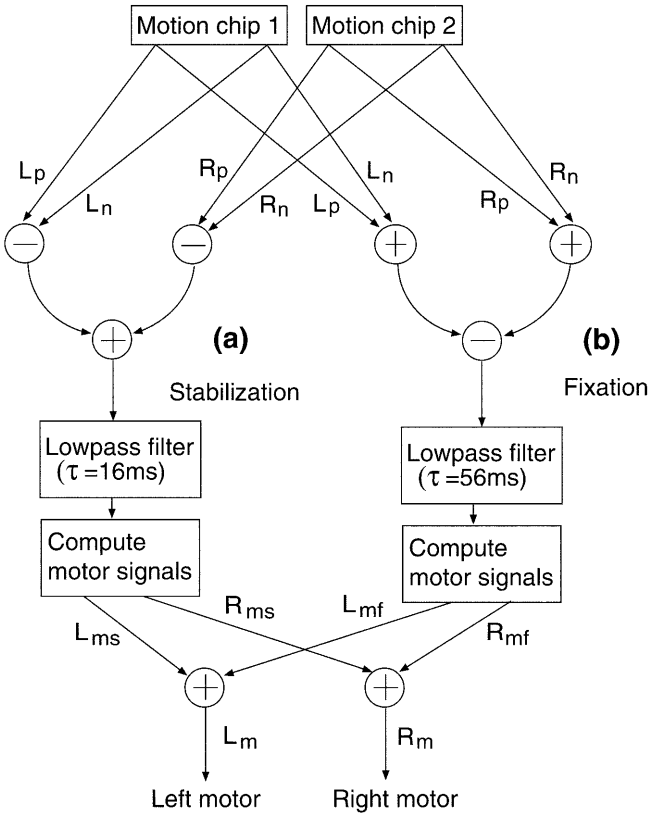

Fig. 10a,b. Block diagram of the control system for course stabilization **a** and fixation **b**. In **a**, the rectified and smoothed signals from the chips corresponding to the null and preferred directions are subtracted. The results of the subtractions are then summed and integrated over time using a low-pass filter. In **b**, the signals corresponding to the null and preferred directions are first summed. The results of these additions are then subtracted and integrated over time. The computed outputs (after they are multiplied by a gain of A) of the two behaviors are used to set the motor outputs

The digitally low-pass filtered outputs were multiplied by a constant, A, and the resulting outputs were used to drive the left motor (L_m) and the right motor (R_m). The constant, A, has been chosen empirically to suit the range of velocities to which the chips are sensitive. Each behavior generates outputs for L_m and R_m. These outputs are labeled as L_{ms} and R_{ms} for the course stabilisation behavior and L_{mf} and R_{mf} for the fixation behavior in Fig. 10. The final values for L_m and R_m are the sum of the motor signals supplied by the two behaviors.

We show experimental data taken from the sensors and the Koala's motor outputs during the two behaviors (see Fig. 11). To duplicate a similar experimental setup in Reichardt and Poggio (1979) for tethered flies, we set the motor signals $R_m = -L_m$ so that the Koala only rotates. These data were collected in the open-loop configuration; the Koala did not physically move in response to the stimulus. The stimulus in the optomotor response experiment consists of a moving background consisting of black and white stripes with an angular wavelength of 6.8°. The field of view of the sensor with a lens of focal length 3.5 mm is 22.5° (shown in Fig. 9). Hence the sensor sees about four spatial wavelengths of the background. The fields of view of the two sensors do not overlap.

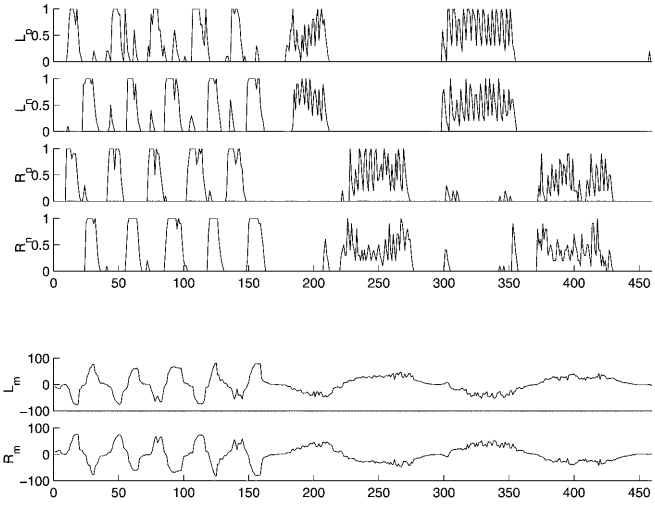

Fig. 11. Measured values of rectified and smoothed outputs of the chips and the speeds of the motors in an open-loop experiment. The *top four plots* correspond, from *top* to *bottom*, to the signals digitized from the left chip and right chip – that is, the preferred and null directions (L_p, L_n, R_p, and R_n) – in arbitrary units. The *bottom two plots* are the speeds of the left motor (L_m) and right motor (R_m) in arbitrary units. The abscissa are in units of 3.3 ms. During the first 160 time units, course stabilization is performed: similar motion in front of both sensors in the preferred (L_p and R_p become positive) or null (L_n and R_n become positive) direction triggers a quick response from the Koala, which rotates in the same direction as the stimulus. From time unit 170 onwards, a stimulus is oscillated in front of one of the two motion sensors (first the left sensor, then the right sensor, then the left sensor, and back to the right sensor). The Koala reacts by rotating towards the stimulus. The ordinate units of the bottom two plots correspond to a rotation speed of approximately 1.3° s^{-1}

The outputs of the aVLSI sensors and the motor signals for the two behaviors are shown on the left part of the six plots in Fig. 11. The top four plots show the rectified and smoothed outputs of the sensors L_p, L_n, R_p, and R_n. The lower two plots show the left and right motor outputs, L_m and R_m. When the background moves in one direction, both the left and the right sensors measure the same motion, as shown by L_p and R_p becoming positive or by L_n and R_n becoming positive. The robot either rotates to the left (R_m is positive and L_m is negative) or to the right (R_m is negative and L_m is positive). The units used in controlling the Koala's motors correspond to 3 mm s^{-1} (K-Team 1999). If the motors move in opposite directions at a speed of 30 mm s^{-1} (corresponding to 10 units, see also Fig. 11), the Koala would rotate at an angular speed of about 13° s^{-1}.

To elicit the fixation-like response, we oscillated a black strip on a white background in front of one of the sensors. In this case, one of the two motion sensors measured a higher motion output and the robot turned accordingly towards the direction of this sensor. From Fig. 11, we see that the Koala samples the motion sensor every 3.3 ms, corresponding to a sampling rate of about 300 Hz. In obtaining both the stabilization and fixation measurements, we did not allow the Koala to physically move.

6 Conclusion

We have described the architecture and algorithmic details of a low-power, continuous-time, aVLSI motion sensor that captures most of the recorded properties of cells in the motion-processing pathway of the fly visual system. The circuits use photoreceptors that adapt over six decades of background intensity. To ensure that we obtain a robust output, we reduced the effects of background intensity, offsets, and mismatches on the motion computation by temporal band-pass filtering of the output of each processing stage. We split the signals into ON and OFF channels so that only ON edges are correlated together and OFF edges are correlated together. We also used transient currents as inputs to the multiplier stage of the model so that we would not be subjected to pixel mismatches in the DC level of voltage inputs to the multiplier. The current output of the multiplier goes to zero when there is no motion. In the case of *voltage* inputs, the output of the multiplier would vary between different pixels because of pixel mismatches.

We have demonstrated a visuomotor controller system that generates a mixture of two behaviors using a linear combination of the outputs of neuromorphic motion sensors. Huber et al. (1999) recently described a visuomotor controller that generates the same two behaviors using a video camera and a Khepera (robot) system. The sampling rate of the camera was 12 Hz. Harrison and Koch (1999) also described a model of the optomotor response behavior of a fly using aVLSI motion sensors and a robot.

The Koala robot in our work has a capability of rotating at about 100° s^{-1} with no visual feedback. The control system that we implemented samples the output from the sensors every 3.3 ms. Hence we can improve the time constant of the visuomotor control loop and the performance of the robot by using small, low-power, and cheap aVLSI preprocessors instead of silicon imagers. Future comparisons between the speeds of different systems will be performed when we incorporate improved motion sensors that generate a better analog representation of the stimulus speed.

While we have obtained a fixation-like behavior using the outputs of the motion sensors that measure global motion, we are not yet able to perform behaviors like tracking and smooth pursuit. We will need to include a sensor that measures the properties of a small-field stimulus in the visual image, just as in the small-field system of the fly. This type of local sensory information along with the global motion information available from the silicon motion sensors will be used to generate more behaviors in the robot.

We believe that by building neuromorphic circuits that extract different visual representations and by mounting these circuits on a robotic platform, we can evaluate the viability of models of biological visuomotor controllers in artificial systems. These circuits give concise, analog, continuous-time information to the robot about its external visual environment. The interaction of these continuous-time sensors with the motor system provides a powerful method for explor-

ing real-time interaction of the agent (robot) with its environment.

Acknowledgements. We thank Rodney Douglas for supporting this work, and the MOSIS foundation for fabricating the circuit. We thank Tobias Delbrück, Andre van Schaik, and Adrian Whatley for critically reading this article. We also thank the anonymous referee who provided detailed comments. This work was supported in part by a Swiss National Foundation SPP Research grant and the U.S. Office of Naval Research. The Koala has been purchased using a NASA grant.

References

Andreou A, Strohbehn AK (1990) Analog VLSI implementation of Hassenstein–Reichardt–Poggio models for vision computation. In: Proc IEEE Int Conf Syst Man Cybern, Los Angeles, CA, 4–7 November, pp 707–710

Borst A, Egelhaaf M, Haag J (1995) Mechanisms of dendritic integration underlying gain control in fly motion-sensitive interneurons. J Comput Neurosci 2: 5–18

Benson R, Delbrück T (1992) Direction selective silicon retina that uses null inhibition. In: Moody J, Hanson S, Lippmann R (eds) Advances in neural information processing systems, vol. 4. Morgan Kaufmann, San Mateo, Calif., pp 756–763

Delbrück T (1991) Bump circuits for computing similarity and disimilarity of analog voltages. In: Proc Int Joint Conf Neural Netw, Seattle, Wash., 8–12 July, pp 475–479

Delbrück T (1993) Silicon retina with correlation-based velocity-tuned pixels. IEEE Trans Neural Netw 4: 529–541

Delbrück T (1994) Analog VLSI phototransduction by continuous-time, adaptive, logarithmic photoreceptor circuits. CNS Memo no. 30, California Institute of Technology, Pasadena, Calif

Diorio C, Minch BA, Hasler H (1999) Floating-gate MOS learning systems. Proc Int Symp Future Intell Integr Electron, Sendai, Japan, 14–17 March pp 515–524

Egelhaaf M (1987) Dynamic properties of two control systems underlying visually guided turning in houseflies. J Comp Physiol A 161: 777–783

Egelhaaf M, Borst A (1989) Transient and steady-state response properties of movement detectors. J Opt Soc Am A 6: 116–126

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

Etienne-Cummings R, van der Spiegel J, Mueller P (1999) Hardware implementation of a visual motion pixel using oriented spatiotemporal neural filters. IEEE Trans Circuits Syst II 46: 1121–1136

Franceschini N, Riehle A, Le Nestour A (1989) Directionally selective motion detection by insect neurons. In: Stavenga DG, Hardie RC (eds) Facets of vision. Springer, Berlin Heidelberg New York, pp 360–390

Franceschini N, Pichon JM, Blanes C, Brady JM (1992) From insect vision to robot vision. Philos Trans R Soc B 337: 283–294

Harrison RR, Koch C (1998) An analog VLSI model of the fly elementary motion detector. In: Jordan M, Kearns M, Solla S (eds) Advances in neural information processing systems, vol. 10. Morgan Kaufmann, San Mateo, Calif, pp 283–294

Harrison RR, Koch C (1999) A robust analog VLSI motion sensor. Auton Robots 7: 211–224

Hassenstein B, Reichardt W (1956) Systemtheoretische Analyse der Zeit-, Reihenfolgen- und Vorzeichenauswertung bei der Bewegungsperzeption des Rüsselkäfers *Chlorophanus*. Z. Naturforsch 11: 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, Egelhaaf M (1989) Neural mechanisms of visual course control in insects. In: Stavenga DG, Hardie RC (eds) Facets of vision. Springer, Berlin Heidelberg New York, pp 391–424

Huber SA, Franz MO, Bulthoff H (1999) On robots and flies: modeling the visual orientation behavior of flies. Robot Auton Syst 29: 227–242

Indiveri G, Douglas R (2000) Neuromorphic vision sensors. Science 288: 1189–1190

Juusola M (1993) Linear and non-linear contrast coding in light-adapted blowfly photoreceptors. J Comp Physiol A 172: 511–521

Juusola M, Uusitola RO, Weckström M (1995) Transfer of graded potentials at the photoreceptor-interneuron synapse. J Gen Physiol 105: 117–148

Kramer J, Sarpeshkar R, Koch C (1995) An analog VLSI velocity sensor. In: Proc IEEE Int Symp Circuits Syst, Seattle, Wash., April 30 – May 3, pp 413–416

Kramer J, Sarpeshkar R, Koch C (1997) Pulse-based analog VLSI velocity sensors. IEEE Trans Circuits Syst II 44: 86–101

K-Team (1999) Koala user manual version 1.1. Lausanne

Laughlin SB (1993) Coding efficiency and visual processing. In: Blakemore C (ed) Vision: coding and efficiency. Cambridge University Press, Cambridge, pp 25–31

Liu SC (1997) Neuromorphic models of visual and motion processing in the fly visual system. Ph.D. thesis, California Institute of Technology, Pasadena, Calif

Liu SC (1999) Silicon retina with adaptive filtering properties. Analog Integr Circuits Signal Process 18: 1–12

Liu SC (2000) A neuromorphic aVLSI model of global motion processing in the fly visual system. IEEE Trans Circuits Syst II 47: 1458–1467

Liu SC, Usseglio-Viretta A (2000) Visuo-motor fly-like responses of a robot using aVLSI motion-sensitive chips. Second Int ICSC Symp Neural Comput, Berlin, Germany, 23–26 May

Mead CA (1989) Analog VLSI and neural systems. Addison-Wesley, Reading, Mass

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

Reichardt W (1961) Autocorrelation: a principle for the evaluation of sensory information by the central nervous system. In: Rosenblith WA (ed) Sensory communication. Wiley, New York, pp 303–317

Reichardt W, Poggio T (1979) Figure-ground discrimination by relative movement in the visual system of the fly, Part I: Experimental results. Biol Cybern 35: 81–100

Ronacher B, Wehner R (1995) Desert ants *Cataglyphis fortis* use self-induced optic flow to measure distances traveled. J Comp Physiol A 177: 21–27

Sarpeshkar R, Bair W, Koch C (1993) An analog VLSI chip for local velocity estimation based on Reichardt's motion algorithm. In: Hanson S, Cowan J, Giles L (eds) Advances in neural information processing systems, 5. Morgan Kaufmann, San Mateo, Calif, pp 781–788

Sarpeshkar R, Kramer J, Indiveri G, Koch C (1996) Analog VLSI architectures for motion processing: from fundamental limits to system applications. Proc IEEE 84: 969–982

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

Srinivasan MV, Chahl JS, Zhang SW (1997) Robot navigation by visual dead-reckoning: inspiration from insects. Int J Pattern Recog Artif Intell 11: 35–47

Webb B (1995) Using robots to model animals: a cricket test. Robot Auton Syst 16: 117–134

Wehrhahn C, Poggio T, Bulthoff H (1982) Tracking and chasing in houseflies (*Musca*). Biol Cybern 45: 123–130