Presented at: International Conference on Advanced Focal Plane Arrays and Electronic Cameras (AFPAEC 1998)

Compact real-time 2-D gradient-based analog VLSI motion sensor

Rainer A. Deutschmann[a] and Christof Koch[b]

[a]Walter Schottky Institute, Am Coulombwall, 85748 Garching, GERMANY
[b]California Institute of Technology, Pasadena, CA 91125, USA

ABSTRACT

In this work we present the first working focal plane analog VLSI sensor for the spatially resolved computation of the 2-D motion field based on temporal and spatial derivatives. Using an adaptive CMOS photoreceptor the temporal derivative and a function of the spatial derivative of the local light intensity are computed. By multiplying these values separately for both spatial dimensions a vector is obtained, which points in the direction of the normal optical flow and whose magnitude for a given stimulus is proportional to its velocity. The circuit consists of only 31 MOSFETs and three capacitors per pixel. We present measurement data from fully functional prototype 2-D pixel arrays for natural stimuli of varying velocity, orientation, contrast and spatial frequency. High direction selectivity even for very low contrast input is demonstrated. As application it is shown how the pixel-parallel architecture of the sensor can favourably be used for real-time computation of the focus of expansion and the axis of rotation. Because of its compactness, its robust operation and its uncritical handling the sensor might be favourably applied in industrial applications.

Keywords: analog VLSI, motion sensor, smart vision sensor, parallel image processing, real-time computation, optical flow, machine vision, robot vision, focus of expansion, axis of rotation

1. INTRODUCTION

The analysis of moving scenes can yield a wealth of information, much beyond what can be obtained from one static image. Including the time domain into image processing allows one to tackle tasks such as figure background segmentation, object tracking and projection of its trajectory, obstacle avoidance, ego-motion estimation, recovering the 3-D velocity and 3-D structure of the viewed scene, and determining the heading direction, i.e. the focus of expansion. The first step to solving these tasks is often to compute the so called optical flow field, which is an estimate of the perspective projection on the image plane of the 3-D velocity field.

For determining the optical flow field on serial computers a tremendous body of literature exists; cf. the comparative study of Barron et. al[1] and references therein. It has proven difficult, though, to obtain the optical flow field in real-time, unless powerful computers are used. Additionally computing times scale unfavourably with the image size and usually increase with $\mathcal{O}(N^2)$, where N×N is the amount of image pixels.

In the last decade a new approach to motion computation has been paid attention to. Progress in VLSI technology for the first time allowed to implement motion detection algorithms on a custom designed chip[2–12] . Using different algorithms, these motion sensors share the following features:

- They are single-chip sensors, i.e. the photoreceptors and the motion computation circuitry sit in the focal plane. Very compact motion detection systems are therefore possible.
- They are pixel-parallel implementations, i.e. motion computation is performed in synchrony by all pixels. This means that the focal plane arrays scale favourably with image size without losing their real-time performance. For achieving higher resolutions in principle just a larger amount of pixel has to be laid out.
- No clock is required for motion computation. Since image irradiance is by nature continuous, an asynchronous circuit implementation is well suited to motion computation and avoids the negative effect of temporal aliasing otherwise encountered in computer implementations.
- In contrast to digital implementations, here transistors are used as analog computing elements. This allows for compact implementation of complex functions and filters.

Send correspondence to R.A.D. email Deutschmann@wsi.tu-muenchen.de, WWW http://maitai.wsi.tu-muenchen.de/deutschmann

- A drawback of analog VLSI motion sensors is their low accuracy, partly because there is no signal restoration as in digital circuits. Given that the input to any vision system is noisy, the additional noise introduced by the analog circuitry does not affect the motion computation seriously in many cases.

Existing analog VLSI motion sensors can be classified into *token-based* and *gradient-based* sensors. Token-based sensors generally look for features, such as edges, in the visual field and track them over time. Kramer et. al. have reported 1-D velocity sensors based on edge tracking[7,8] which operate robustly over a wide range of input parameters. Based on a similar feature extraction stage Deutschmann et. al.[10] have reported two sensors that compute the direction-of-motion vector field in a 2-D array, but are not sensitive to stimulus velocity. Etienne-Cummings et. al.[6] have reported on a 5×5 sensor, the scaling of which to large arrays might be difficult. The authors propose a multi-chip system.

On the other hand gradient-based sensors use local temporal and spatial derivatives of the light intensity to compute motion. An early design of Tanner and Mead[2] tried to implement the gradient constraint equation $dI/dt = 0$ through a feedback mechanism. The chip could only solve for one global 2-D velocity vector and showed poor performance. A similar circuit was used later to compute motion in 1-D.[13] A working redesign of the Tanner/Mead chip in 1-D has been reported recently.[14] Deutschmann et. al.[12] implemented straightforwardly the 1-D gradient model, which solves the gradient constraint equation and yields velocity independent of spatial frequency and contrast by a division of temporal and spatial derivatives. In 2-D, though, the algorithm would yield a large pixel size. A very compact 2-D pixel that makes implicit use of temporal derivatives was reported by Benson and Delbrück.[15] Their sensor, though, is tuned to a fixed velocity and is sensitive only to ON-edges.

In summary there exist a series of quite powerful 1-D motion sensors to choose from, some of which represent velocity unambiguously. The choice amongst single chip 2-D sensors is small, and dense pixel arrays have been reported only for direction-of-motion or velocity-tuned sensors. Our goal was to develop a 2-D motion sensor for applications that require the sensor output to be monotonically increasing with stimulus velocity. Additionally the motion vector orientation should be parallel to the normal optical flow. The sensor should be scalable to high density arrays on a single chip, and should be easy to use due to a small number of biasing voltages, none of them being overly critical. We achieved our goal with the motion sensor presented in this paper, called "Gradient2d" sensor, which is based on Horiuchi's work[16] on a biologically inspired oculomotor system.

The paper is organised in the following way: In the next section we will introduce the motion algorithm and its analog VLSI implementation. In the following sections we will present experimental results characterising the elementary motion detector; first the time-dependent output, then the output dependence on stimulus velocity, orientation, contrast and spatial frequency. Subsequently we will demonstrate the high direction selectivity of the sensor even for very low stimulus contrasts. In the last part of the paper we will show 15×15 motion vector fields of the entire 2-D pixel array and use its pixel-parallel architecture to compute the focus of expansion and axis of rotation of a real-world scene in real time.

2. ALGORITHM AND IMPLEMENTATION

We will derive an equation which relates the change in image brightness at a point in the focal plane to the motion of the brightness pattern. A unique solution for the optical flow will be obtained with the additional constraint of normal flow. We will show how to simplify this solution to make it suitable for implementation in analog VLSI, maintaining normal motion flow vectors.

Let the image brightness at the point (x, y) in the focal plane at time t be denoted by $I(x, y, t)$. If the brightness pattern stems from a solid object and moves, the brightness of a particular point in the pattern projected on the focal plane is constant. Using the chain rule of differentiation we obtain

$$\frac{dI(x,y,t)}{dt} = 0 \qquad \Rightarrow \qquad \frac{\partial I}{\partial x} v_x + \frac{\partial I}{\partial y} v_y + \frac{\partial I}{\partial t} = 0 \tag{1}$$

for the two unknown velocity components $v_x = \partial x/\partial t$ and $v_y = \partial y/\partial t$. An additional constraint is required for a unique solution. Several constraints have been suggested in the literature, for example the smoothness of the optical flow field.[17] This approach is problematic because discontinuities in the optical flow field, on which segmentation from motion algorithms rely, are attenuated. Therefore we confine the flow vectors to be normal to the local image

brightness distribution (normal flow). In mathematical terms $I_x/v_x = I_y/v_y$, where $I_x = \partial I/\partial x$ and $I_y = \partial I/\partial y$. We thus obtain for the normal velocity field (v_x, v_y)

$$v_x = -\frac{I_x\, I_t}{I_x^2 + I_y^2} \quad , \qquad v_y = -\frac{I_y\, I_t}{I_x^2 + I_y^2} \quad . \tag{2}$$

Although these equations could directly be implemented in analog VLSI,[12] this attempt is tricky because a) the result is not defined for vanishing contrast and b) both spatial dimensions are coupled. We observe that the sum $I_x^2 + I_y^2$ is independent of the stimulus orientation and that it can be considered as confidence measure for the local velocity computation. We therefore weight (multiply) v_x and v_y by this sum and define a new motion vector field (u_x, u_y)

$$u_x = u'_o\, I_t\, I_x \quad , \qquad u_y = u'_o\, I_t\, I_y \quad ; \quad u'_o \in \mathbb{R}\ const. \tag{3}$$

Up to now we derived a normal flow field that simply consists of a product of temporal and spatial derivatives, which lends itself easily for implementation in analog VLSI. For a given stimulus the length of the motion vector (u_x, u_y) still increases linearly with velocity, but it is now dependent on the stimulus contrast.

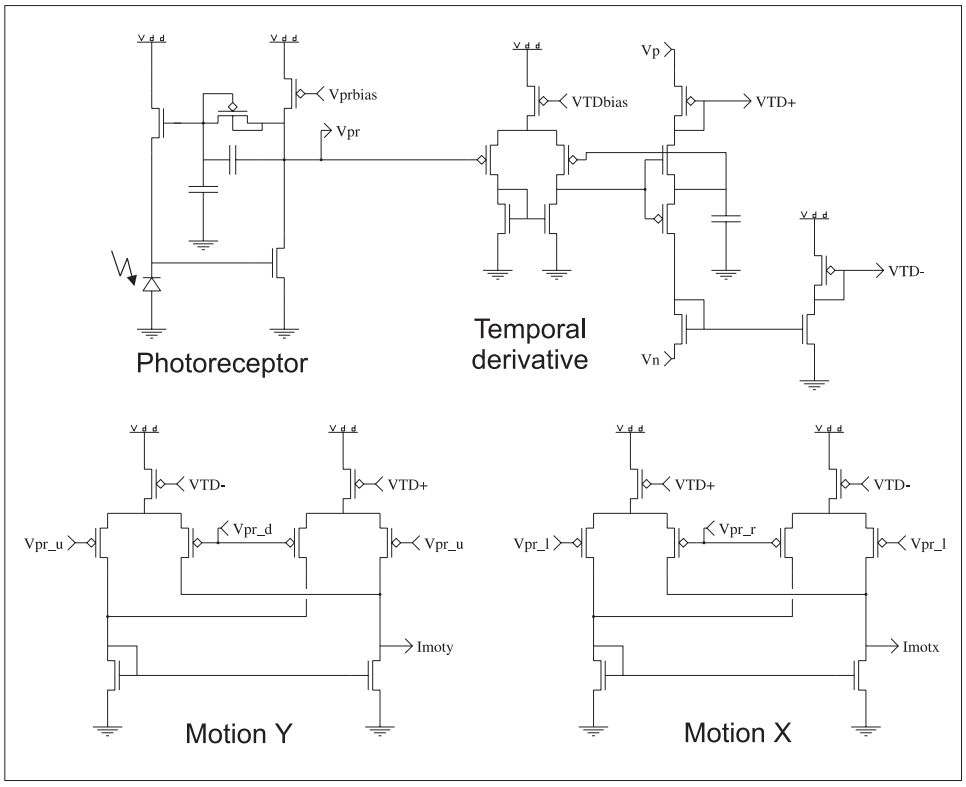

Figure 1. Schematic of the Gradient2d sensor pixel

The schematic of the actual implementation is shown in Figure 1. The light intensity is transduced to a voltage signal V_{pr} by Delbrück's adaptive photoreceptor.[18] The photoreceptor signal is transmitted to all four neighbouring pixels, and is locally used for computing the temporal derivative. The temporal derivative circuit is constructed from a transconductance amplifier with negative feedback which copies the photoreceptor voltage onto a capacitor through a bidirectional source follower. The current required to drive the capacitor voltage is amplified by a tilted current mirror (V_n slightly above GND and V_p slightly below Vdd). The temporal derivative signal is represented by the voltages V_{TD}^+ and V_{TD}^-, separately for increasing and decreasing light intensities, respectively. The circuit has the advantage of having low offset current in steady state, but it shows crossover distortion (see Section 3.1).

The spatial derivatives are computed by differential pairs. For I_x and I_y the inputs come from left and right (V_{pr_l} and V_{pr_r}), upper and lower (V_{pr_u} and V_{pr_d}) neighbour, respectively. The current sources feeding the differential

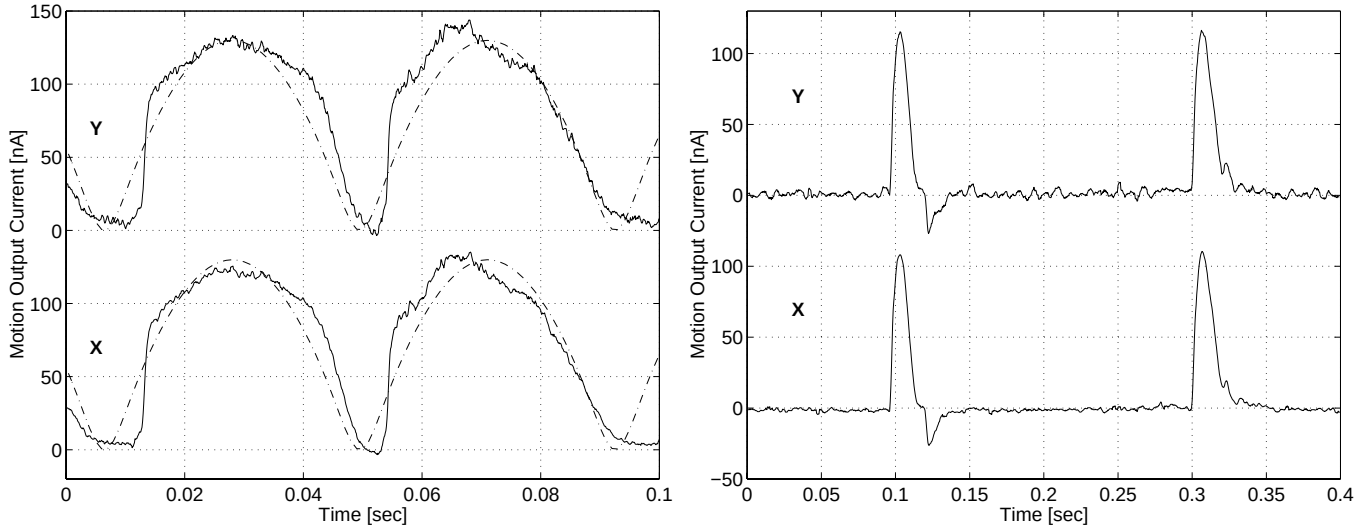

Figure 2. Time course of the motion signals for sinusoid and bar shaped stimuli. The X and Y motion vector components of one pixel were recorded over time for diagonally oriented stimuli. One period is shown. The dashed line shows the theoretical expectation.

pairs are biased by the temporal derivative voltages V_{TD}^{+} and V_{TD}^{-}, effectively making up four-quadrant Gilbert multipliers. For large spatial derivatives the output of the differential pairs saturate like a hyperbolic tangens function. In effect the actual motion vector field (u_x, u_y) computed by the Gradient2d sensor is described by

$$u_x = u_o\, I_t\, \tanh\left(\lambda\, I_x\right) \quad , \qquad u_y = u_o\, I_t\, \tanh\left(\lambda\, I_y\right) \tag{4}$$

where u_o is a scaling constant and λ describes the slope of the tanh function, which is implementation specific.

The motion vector components u_x and u_y are represented as two bidirectional currents I_{motx} and I_{moty} which are separately scanned off chip. A vertical scanner[19] is used for selecting one row of pixels, one of which is picked by a horizontal scanner. In addition to the motion vectors also the photoreceptor and temporal derivative signals can be scanned out.

In summary a total of only 31 transistors and 3 capacitors is required for the 2-D motion computation. The connectivity is easy due to only four bias voltages and the fact that only two communicating wires are required between every pair of pixels. Various prototype sensors with 8×8 up to 15×15 pixels have been fabricated with standard 2.0μm and 1.2μm CMOS processes on 2.2×2.2 mm dyes. The pixel size on a 1.2μm process is 112μm×112μm. On a chip size of 9.4×9.7 mm we expect resolutions of at least 80×80 pixels.

3. EXPERIMENTAL RESULTS

We are now characterising the single pixel output of the Gradient2d sensor pixel array. The Gradient2d sensor reports motion exactly during the time when motion is perceived, i.e. when a brightness pattern is changing in time. This is contrary to sample-and-hold type motion sensors, where the motion output is computed and then held fixed for a certain time. We therefore first look at the transient motion output of the Gradient2d sensor. Subsequently the time-integrated motion output will be used for characterisation.

3.1. TIME-DEPENDENT OUTPUT

In Figure 2 the motion output is shown as a gray value pattern is moved diagonally in front of the sensor. In the experiment displayed on the left a sinusoid was used. As predicted by Equation 4 the motion output for both X and Y component resembles a squared sinusoid; the theoretically expected response is the dashed curve. As expected for diagonally oriented stimuli both vector components are of equal size. It is remarkable how well the output is rectified. Although AC lighting was used for all experiments flicker noise in the temporal derivative could be avoided

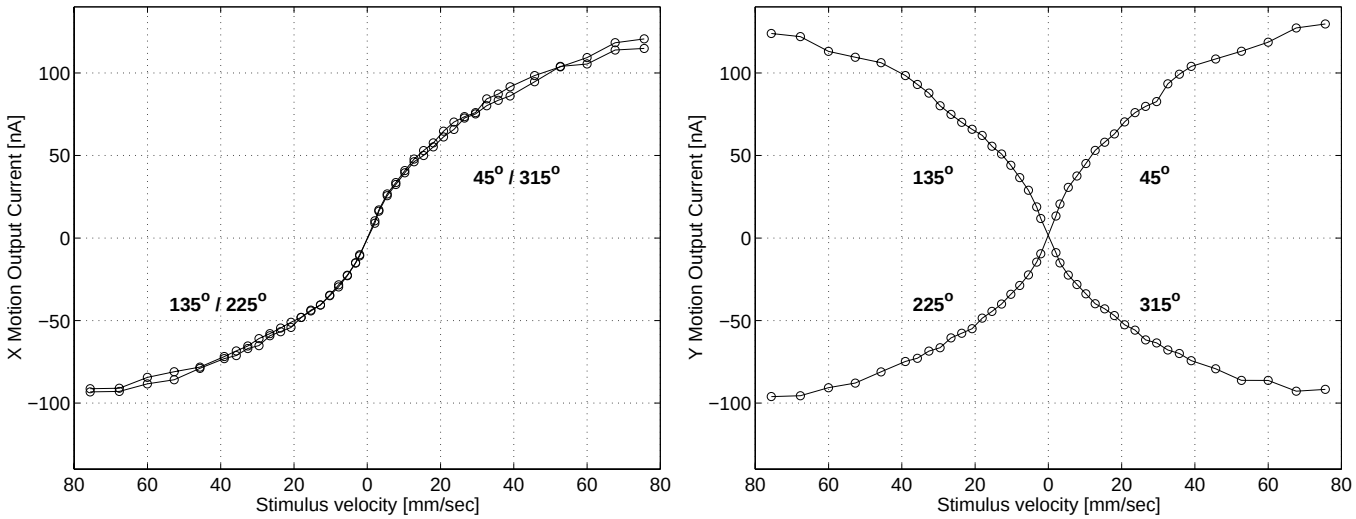

Figure 3. Velocity dependence of the Gradient2d sensor. To the left the X component, to the right the Y component of the motion output of one pixel is shown for four different stimulus orientations.

and does not show up in the motion output. The delayed onset after zero output is due to the cross-over distortion induced by the temporal derivative circuit.

To the right of Figure 2 a black bar on a white background with diagonal orientation was used as stimulus. From Equation 4 because of the strong temporal derivative at the edges of the bar a peak like output is expected. By defocusing the sensor lens slightly the motion output can be increased to a finite dwell time. A small overshoot can be observed at the bright-to-dark transition (OFF edge) which is caused by the temporal derivative circuit. For both OFF and ON edges the motion output again correctly reaches the same peak height.

3.2. VELOCITY DEPENDENCE

There are several ways to characterise the transient velocity dependent output of a motion sensor. Since for the experiments we were using periodic stimuli we chose to measure the average output over several periods, instead of measuring for example its peak height. Using Equation 4 it can be shown that for periodic stimuli the average motion output is linear with stimulus velocity. In the experiment a sine wave grating of 72.5% contrast and 0.068 cycles/degree spatial frequency was used. Stimulus velocity is given as on-chip speed.

In Figure 3 experimental results for stimulus velocities up to 76 mm/sec and four different stimulus orientations are shown; on the left figure the X component of the motion output of one pixel, on the right figure the Y component. For stimulus movement in 45° both X and Y component are positive, for stimulus orientation in 135° the X component is negative, the Y component is positive, correspondingly for all other orientations. In one fixed orientation the absolute value of the motion output increases monotonically with increasing stimulus velocity. As can be seen there is a slight gain mismatch between positive and negative outputs, which is caused by device mismatch in the temporal derivative and multiplication circuits.

The compressive nonlinearity observed in the motion output stems from the low pass filter behaviour of the photoreceptor. For larger stimulus velocities the photoreceptor peak-to-peak output thus decreases and causes the motion output to increase sub-linear. From a Bode plot a first order low pass filter cutoff frequency of $1/\tau = 161$Hz has been determined; cf. Figure 4 left. This cutoff frequency can be adjusted by the photoreceptor bias V_{prbias}. The observed compressive non-linearity is desired because it increases the velocity range of the motion sensor.

The Gradient2d sensor is sensitive also to a much lower range of stimulus velocities, which can be seen in Figure 4 right. Here the temporal derivative was amplified stronger by adjusting the bias voltages V_p and V_n. Stimulus velocities below 0.1 mm/sec can unambiguously be detected. Additionally is becomes obvious that the motion output is now linear with velocity, because in this range the low pass filtering in the photoreceptors is negligible.

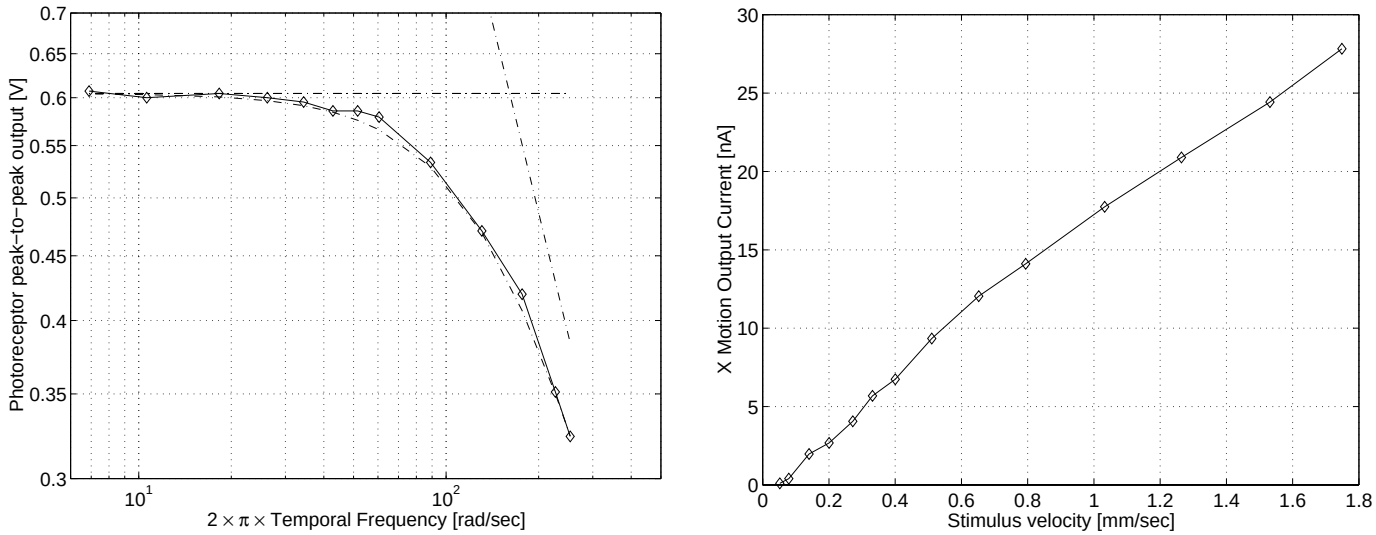

Figure 4. Left: Bode plot of the photoreceptor response for increasing flicker frequency. The dashed curve is a first order low pass filter fit with time constant τ=1/161Hz. Right: Motion output for very small stimulus velocities.

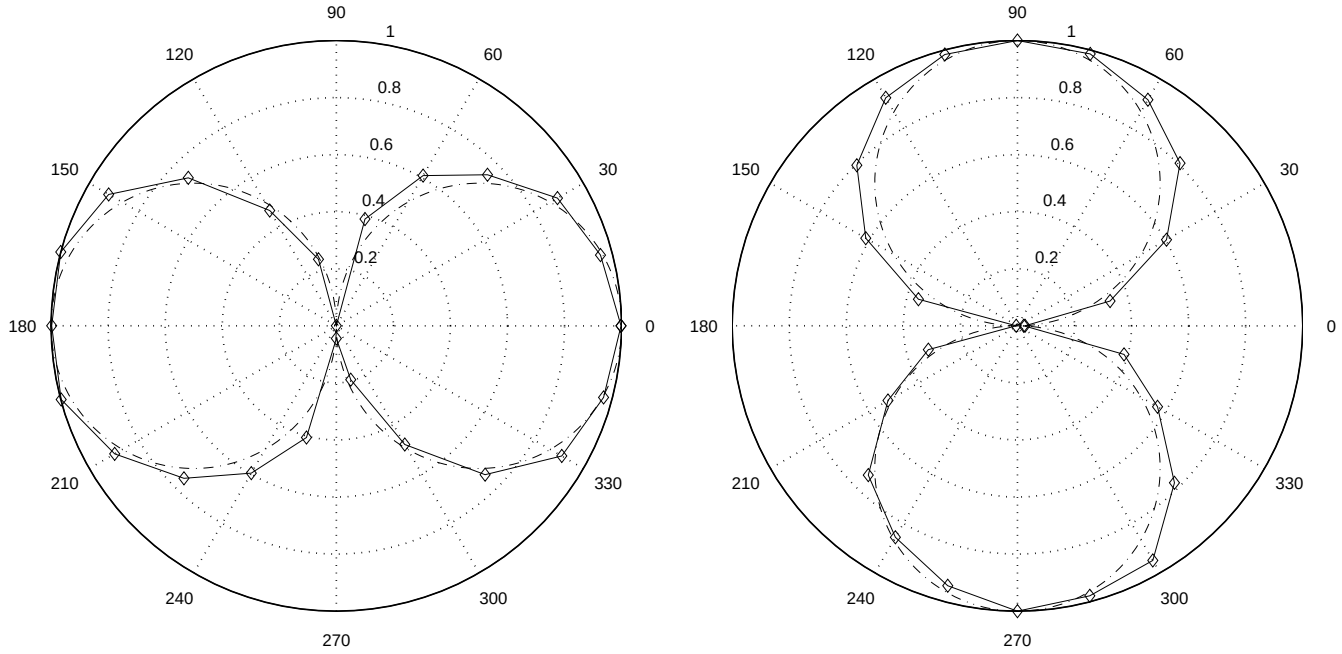

Figure 5. Measured orientation tuning curve for low contrast sine wave grating. The dashed curve represents the ideal cosine like tuning curve.

3.3. ORIENTATION TUNING CURVE

For a 2-D motion sensor a reasonable orientation tuning curve is crucial. Reasonable means that the motion output should vary smoothly with stimulus orientation. We have derived in Section 2 that the Gradient2d sensor is to produce a normal flow field, i.e. the motion output for both vector components should be cosine-like dependent on the stimulus orientation. The X component, for example, should approach zero for orientations near 90° and 270°. In an experiment a sine wave grating of 22% contrast and fixed velocity was moved in front of the sensor at orientations from 0° to 360°. The resulting motion output for both vector components is displayed in Figure 5 as polar plot. As can be seen the motion output meets the theoretical expectations and forms cosine lobes. For stronger

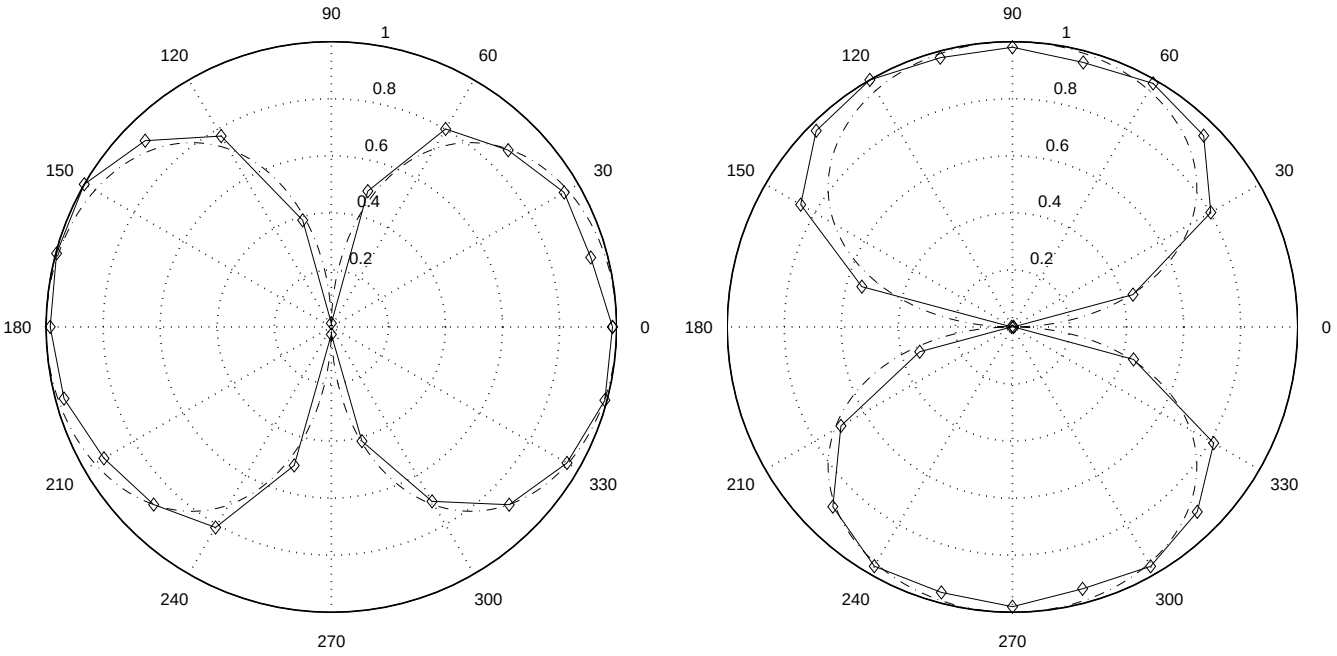

Figure 6. Measured orientation tuning curve for high contrast sine wave grating. The theoretical expectation for saturating spatial derivative is displayed as dashed line.

contrasts the saturation of the spatial derivative, which is described by the hyperbolic tangens term in Equation 4, becomes important. The effect can be observed in Figure 6, where a high-contrast sine wave grating was used. The X component of the motion output, for example, remains at a high level for orientations even around 45°. Since at this orientation the Y component also is large, in the resulting motion flow fields an emphasis of diagonally oriented vectors can be observed; cf. the flow fields in Section 4.

3.4. CONTRAST DEPENDENCE

We have shown earlier how to simplify Equation 2 for the normal velocity field in order to make it stable against the division by zero problem, and to make it suitable for a compact analog VLSI implementation. From the resulting equation it can easily be seen that the motion vectors are now contrast dependent; for low contrast the motion output should increase quadratically with stimulus contrast, for high contrasts the increase should be linear.

We have tested experimentally the contrast dependence of the motion output in that we moved sine wave gratings of contrasts ranging from 2% to 72.5% in front of the sensor, while their spatial frequency was 0.068 cycles/degree and their velocity was kept at 8 mm/sec. Results for one vector component are shown in Figure 7 left. It can be seen that for low contrasts the increase of the motion output is first quadratic, then continues on linearly.

This dependence of the motion output on the stimulus contrast might exclude the Gradient2d sensor from some applications. A slow object of high contrast can produce a similar motion output as a faster object of lower contrast. On the other hand the motion output can now be considered as reliability measure for a local direction of motion computation. Low contrast objects produce a smaller motion output than high contrast objects and thus signal the presence of a less reliable motion cue.

3.5. SPATIAL FREQUENCY DEPENDENCE

It has been mentioned earlier that due to the continuous-time operation of the Gradient2d sensor, temporal aliasing in the motion computation cannot occur. Spatial aliasing, though, can occur in all motion sensors with pixel based architecture. In the Gradient2d sensor the spatial derivative at one pixel location is computed from the difference of the photoreceptor voltages of the two adjacent pixels. For a sine wave grating of ever higher spatial frequency, for example, the local spatial derivative is therefore more and more underestimated, until the wavelength equals twice

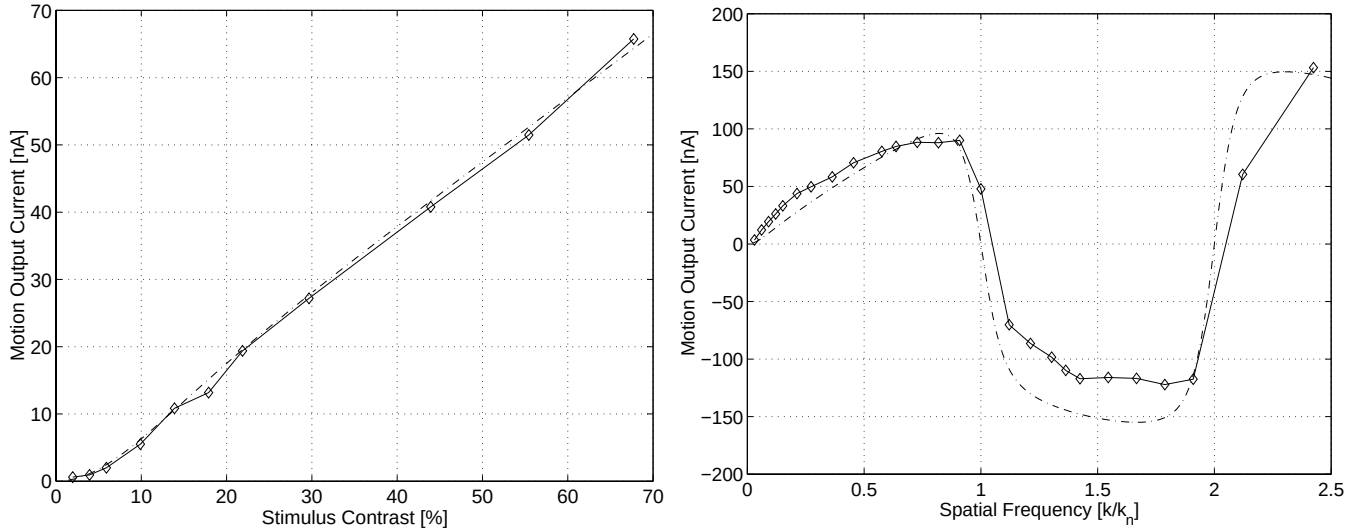

Figure 7. Left: Measured contrast dependence of the motion output. For low contrasts the motion output is quadratically dependent on the stimulus contrast, whereas for high contrasts the dependency is linear. The dashed line shows the theoretical expectation. Right: Measured spatial frequency dependence of the motion output. The motion output for sine wave stimuli of equal velocity (5.5 mm/sec), contrast (72.5 %) and orientation is plotted for spatial frequencies between 0.01 cycles/degree and 1 cycle/degree. The dashed line was computed considering the discrete spatial derivative and low pass filtering in the photoreceptors.

the pixel spacing. At this point the computed spatial derivative turns negative. The spatial frequency where aliasing occurs is called *Nyquist frequency* k_n.

We have experimentally tested this prediction by using sine wave gratings of equal velocity and contrast, but of spatial frequencies ranging from 0.01 cycle/degree to 1 cycle/degree. In Figure 7 right the motion output is plotted versus the spatial frequency, which is given in units of the Nyquist frequency. For this particular sensor the Nyquist frequency was $k_n = 0.32$ cycles/degree. The dashed line represents the theoretical prediction considering the discreteness of the spatial derivative and low pass filtering in the photoreceptors. As can be seen the measured data is well approximated by the theoretical curve. The motion output increases with increasing spatial frequency until shortly before the aliasing point, where the output crosses zero and reverses sign. At even higher spatial frequencies a second zero crossing can be observed.

As in the discussion about the contrast dependence it can be argued here that the motion output dependence on the spatial frequency hampers the use of the Gradient2d sensor. But first of all it is obvious that as with any other pixel based motion sensor spatial aliasing has to be avoided, e.g. by defocusing the lens slightly. Secondly the fact that the motion output approaches zero for long wavelengths can again be seen as indicator for low local contrast and a less reliable motion signal.

3.6. DIRECTION SELECTIVITY

For some applications the local direction of motion of very low contrast objects might be of interest. For that reason we have tested the Gradient2d sensor with a stimulus of only 4 % contrast. In an experiment the 4% stimulus was moved in front of the sensor in one direction for 4 seconds, then the direction of motion was reversed. The output of the sensor is of course very small, and relatively noisy, cf. Figure 8 left. But if the output is integrated over a short time, for example 100 ms, then the direction of motion of the stimulus can unambiguously be determined, as shown in Figure 8 right. Here a histogram of the motion output, averaged over 100 ms, is shown. Clearly leftward motion can be distinguished from rightward motion.

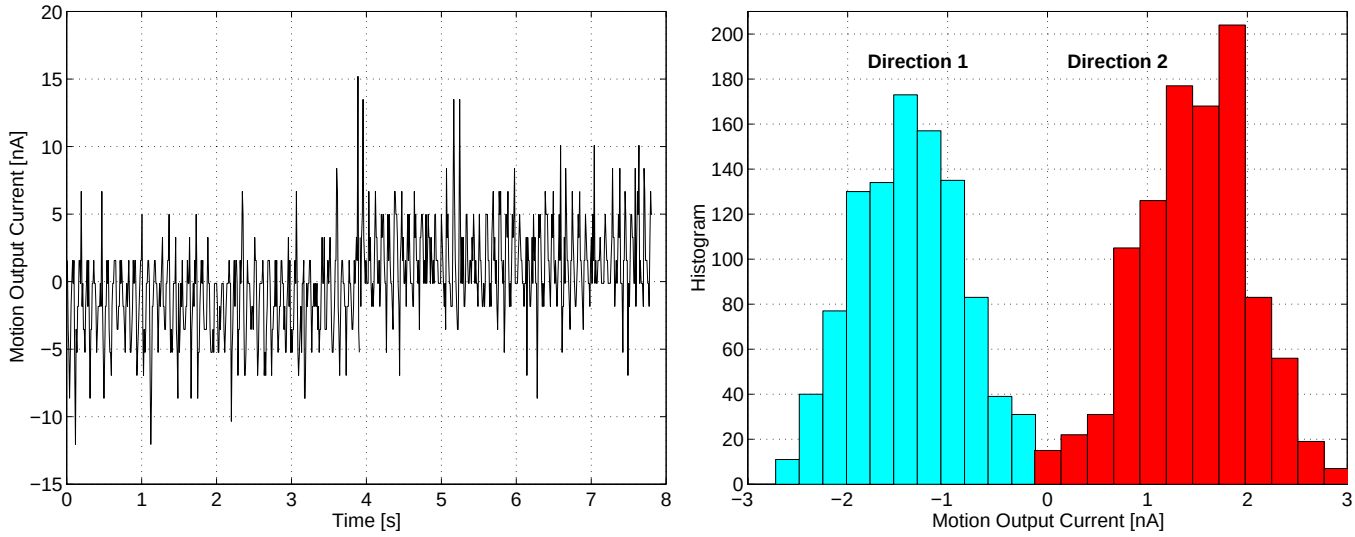

Figure 8. Direction selectivity for 4 % contrast stimulus. Left: Motion output during leftward motion (t <4 sec) and rightward motion (t >4 sec). Right: Histogram of the 100 ms time averaged motion output.

4. APPLICATION: DETERMINING FOCUS OF EXPANSION

In this section we briefly present examples of how the Gradient2d sensor can favourably be used for real world tasks. Consider an application where the focus of expansion (FOE) or the axis of rotation (AOR) of a moving scene is of interest. Then the X component of the FOE can be computed by finding the position of the zero crossing (ZC) of the column average of u_x, the X component of the corresponding motion vectors. If there are multiple ZCs, the one with the maximal slope is to be taken. The FOE Y component and the AOR location are given accordingly:

$$x^{FOE} = ZC_x\left(\sum_y u_x(x,y)\right) \qquad y^{FOE} = ZC_y\left(\sum_x u_y(x,y)\right) \tag{5}$$

$$x^{AOR} = ZC_x\left(\sum_y u_y(x,y)\right) \qquad y^{AOR} = ZC_y\left(\sum_x u_x(x,y)\right) \tag{6}$$

The Gradient2d sensor can be used not only to compute the motion vector field $u(x,y)$ in real time, but also the row and column averages: The pixel array is addressed by one row and one column scanner, and the motion vector components are represented as bidirectional currents. Thus by gating out one entire row or column, the vector sum is obtained automatically. All N row sums and N column sums are read into a computer, where the ZCs and thus the FOE or AOR is determined. Since only 2N instead of N^2 operations are required, where N×N is the array size, the FOE and AOR can be determined very fast. We achieve frame rates beyond 400Hz. In Figures 9 and 10 we present snapshots of a 15×15 Gradient2d sensor whose motion output was set to saturate earlier than in Figure 3, effectively computing a direction of motion flow field. Clearly the motion of the scene is captured by the sensor, and the row and column averages computed on-chip indicate correctly the axis of rotation and focus of contraction.

5. SUMMARY

In this contribution after a brief review of existing motion sensors we have derived a 2-D motion computation algorithm that is simply based on a multiplication of temporal and spatial derivatives. We have demonstrated how this algorithm can very compactly be implemented using analog VLSI technology. Subsequently we have characterised the single pixel output regarding its time-dependent response, and its mean motion output for stimuli of different velocities, orientations, contrasts and spatial frequencies. We have argued that the motion output dependence on contrast and spatial frequency can favourably be used to asses the reliability of the motion cue. The high sensitivity of the Gradient2d sensor to low contrast objects has been demonstrated. Finally in the last section we have presented an algorithm for real time computation of the focus of expansion and the axis of rotation of a moving scene using

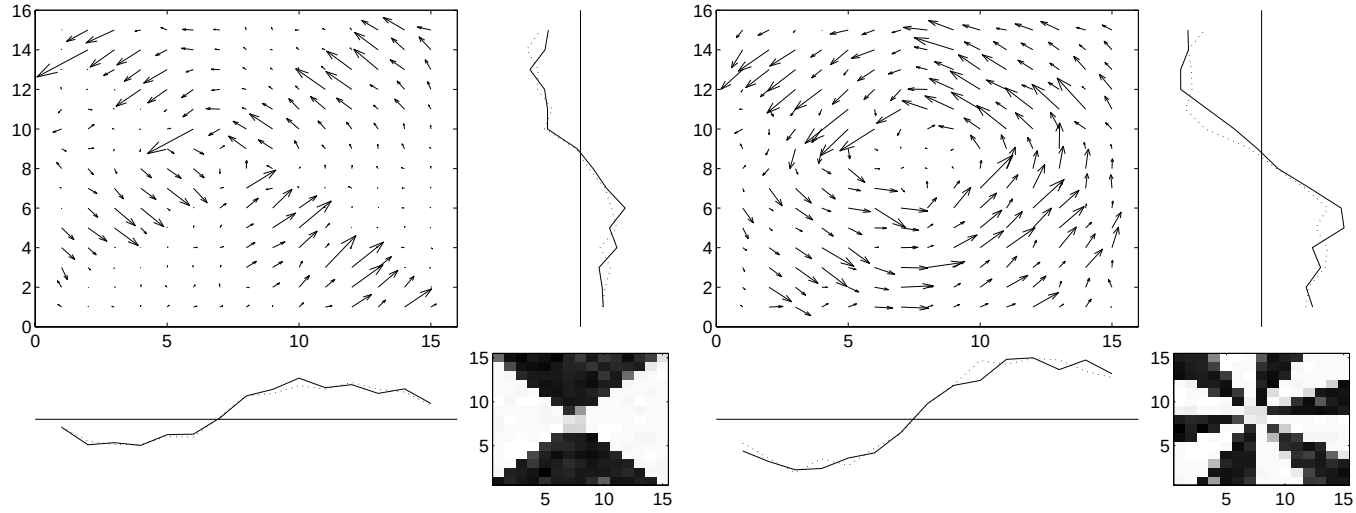

Figure 9. Determining the axis of rotation from two different moving scenes. The motion vector fields as computed by the sensor, and the visual input as seen by the on-chip CMOS imager are displayed for one instant in time. The column sum of the Y component and the row sum of the X component of the motion vectors are drawn below the flow fields and to their right, respectively. The solid curve was computed on-chip, the dotted curve was computed from the flow field external of the sensor for comparison. The axis of rotation for both stimuli is correctly marked by the zero crossing.

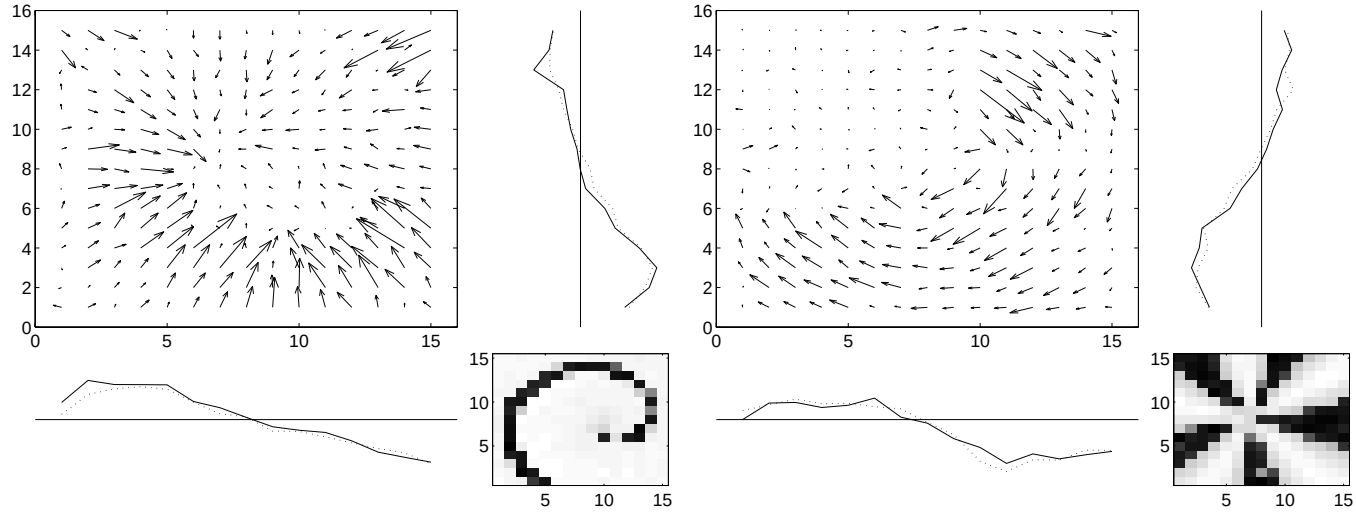

Figure 10. Left: A receding object was simulated by rotating a spiral in the field of view of the sensor. The focus of contraction is correctly indicated by the zero crossings of the averages. Similarly the focus of expansion can be determined. Right: Even for an occluded input the correct axis of rotation is determined by the zero crossings.

the pixel-parallel architecture of the Gradient2d sensor. Using a 15×15 pixel array the algorithm was shown to work reliably. In conclusion we have demonstrated a compact 2-D gradient based motion sensor that due to its easy handling and robust operation might be used in some industrial applications.

ACKNOWLEDGMENTS

Many thanks to Timothy Horiuchi for laying the grounds of this work. This work was supported by the Center for Neuromorphic Systems Engineering as a part of the National Science Foundation's Engineering Research Center program.

REFERENCES

1. J. Barron, D. Fleet, and S. Beauchemin, "Systems and experiment: Performance of optical flow techniques," *Intern. J. Comp. Vis.* **12**, pp. 43–77, 1994.
2. J. Tanner and C. Mead, "An integrated analog optical motion sensor," *VLSI Signal Processing, IEEE Press* **2**, pp. 59–76, 1986.
3. A. Andreou, K. Strohbehn, and R. Jenkins, "Silicon retina for motion computation," *Proc. IEEE Int. Symp. on Circuits and Systems* **3**, pp. 1373–1376, 1991.
4. T. Horiuchi, W. Bair, B. Bishofberger, A. Moore, and C. Koch, "Computing motion using analog VLSI vision chips: an experimental comparison among different approaches," *Intern. Journal of Computer Vision* **8**, pp. 203–216, 1992.
5. T. Delbrück, "Silicon retina with correlation-based, velocity-tuned pixels," *IEEE Trans. on Neural Networks* **4**, pp. 529–541, 1993.
6. R. Etienne-Cummings, J. V. der Spiegel, and P. Mueller, "A focal plane visual motion measurement sensor," *IEEE Trans. Circuits and Systems 1* **44**, pp. 55–66, 1997.
7. J. Kramer, "Compact integrated motion sensor with three-pixel interaction," *IEEE Trans. Pattern Anal. Machine Intell.* **18**, pp. 455–560, 1996.
8. J. Kramer, R. Sarpeshkar, and C. Koch, "Pulse-based analog VLSI velocity sensors," *IEEE Trans. Circuits and Systems II* **44**, pp. 86–101, 1997.
9. A. Moini, A. Bouzerdoum, K. Eshraghian, A. Yakovleff, X. Nguyen, A. Blanksby, R. Beare, D. Abbott, and R. Bogner, "An insect vision-based motion detection chip," *IEEE J. of Solid-State Circuits* **32**(2), pp. 279–284, 1997.
10. R. A. Deutschmann, C. Higgins, and C. Koch, "Real-time analog VLSI sensors for 2-D direction of motion," in *Proc. Int. Conf. on Artificial Neural Networks ICANN'97*, vol. 1327 of *Lecture Notes in Computer Science*, pp. 1163–1168, Springer Verlag, 1997.
11. R. R. Harrison and C. Koch, "An analog VLSI model of the fly elementary motion detector," in *Advances in Neural Information Processing Systems*, M. I. Jordan, M. J. Kearns, and S. A. Solla, eds., vol. 10, MIT Press, Cambridge, MA, 1998.
12. R. A. Deutschmann and C. Koch, "An analog VLSI velocity sensor using the gradient method," in *Proc. International Symposium on Circuits and Systems ISCAS'98*, 1998.
13. A. Moore and C. Koch, "A multiplication based analog motion detection chip," *Proc. SPIE, Visual Information processing: From neurons to chips* **1473**, pp. 66–75, 1991.
14. A. Stocker, "Smooth optical flow computation by a constraint solving circuit," *INNS/ENNS/JNNS Newsletter* **21**, pp. 5–6, 1998.
15. R. G. Benson and T. Delbrück, "Direction-selective silicon retina that uses null inhibition," in *Advances in Neural Information Processing Systems*, J. E. Moody, S. J. Hanson, and R. P. Lippmann, eds., vol. 4, pp. 756–763, Morgan Kaufmann, San Mateo, CA, 1992.
16. T. Horiuchi, *Analog VLSI-Based, Neuromorphic Sensorimotor Systems: Modeling the Primate Oculomotor System*. PhD thesis, California Institute of Technology, 1997.
17. B. Horn and B. Schunck, "Determining optical flow," *Artificial Intelligence* **17**, pp. 185–203, 1981.
18. T. Delbrück and C. A. Mead, "Analog VLSI phototransduction," *Caltech CNS Memo* **30**, pp. 139–161, 1994.
19. C. A. Mead and T. Delbrück, "Scanners for visualizing activity of analog VLSI circuitry," *Caltech CNS Memo* **11**, 1991.