

# Analog Hardware for Detecting Discontinuities in Early Vision

JOHN G. HARRIS, CHRISTOF KOCH, ERIK STAATS, AND JIN LUO

*Computation and Neural Systems Program, 216-76 California Institute of Technology, Pasadena, California 91125*

## Abstract

The detection of discontinuities in motion, intensity, color, and depth is a well-studied but difficult problem in computer vision [6]. We discuss the first hardware circuit that explicitly implements either analog or binary line processes in a deterministic fashion. Specifically, we show that the processes of smoothing (using a first-order or membrane type of stabilizer) and of segmentation can be implemented by a single, two-terminal nonlinear voltage-controlled resistor, the "resistive fuse"; and we derive its current-voltage relationship from a number of deterministic approximations to the underlying stochastic Markov random fields algorithms. The concept that the quadratic variation functionals of early vision can be solved via linear resistive networks minimizing power dissipation [37] can be extended to non-convex variational functionals with analog or binary line processes being solved by nonlinear resistive networks minimizing the electrical co-content.

We have successfully designed, tested, and demonstrated an analog CMOS VLSI circuit that contains a 1D resistive network of fuses implementing piecewise smooth surface interpolation. We furthermore demonstrate the segmenting abilities of these analog and deterministic "line processes" by numerically simulating the nonlinear resistive network computing optical flow in the presence of motion discontinuities. Finally, we discuss various circuit implementations of the optical flow computation using these circuits.

## 1 Introduction

Most early vision algorithms incorporate the generic constraint that variables such as surface orientation and reflectance, depth or optical flow vary slowly in space [7, 12, 17, 20, 27, 40, 43]. Within the standard regularization approach, this is reflected in the use of stabilizing operators corresponding to various measures of smoothness [39]. Thus, in the problem of interpolating a 2D surface through sparse and noisy depth measurement, the final surface should be as close as possible to the initial data as well as being as smooth as possible [7] or, in the problem of computing optical flow from the time-varying intensity, the final flow field should be compatible with the locally measured velocity data as well as being smooth [12, 17, 34, 48]. However, surfaces display discontinuities where the smoothness constraint is violated. Thus, the to-be-reconstructed surface may have been generated by an underlying piecewise smooth or even piecewise constant depth distribution, or the 2D velocity field induced by a rigid object moving in an otherwise stationary environment varies

smoothly across the surface of the object but is zero beyond the contours of the object.

In the last years, a number of researchers have introduced powerful algorithms to deal with the representation of such discontinuities. Geman and Geman [6] first proposed binary line processes to model discontinuities in intensity within the stochastic framework of Markov random fields. Discontinuities are subject to various constraints, such that they should form along continuous contours or should not intersect nor form parallel lines. Their approach was extended and modified to account for discontinuities in depth, texture, optical flow, and color [19, 28, 29, 38, 49]. The principal drawback of the Geman and Geman-type method is the computational expense involved in minimizing the associated nonquadratic cost functionals using stochastic optimization methods, in particular when numerous constraints are incorporated. A number of authors have used deterministic methods to find the (local) minima of the associated nonconvex variational functionals, with next-to-optimal results (e.g., [3, 24, 44]). A rigorous deterministic approach has been championed by Blake and

Zisserman [1]. Their “graduated nonconvexity” (GNC) algorithm bears many similarities to the above methods, and leads to excellent results in the case of piecewise continuous reconstruction of surfaces [2].

The idea of using analog circuits for solving vision problems was raised by Horn [14], where he proposed the use of a grid of resistors to find the inverse of the discrete approximation to the Laplacian. Poggio and Koch [37] show how standard regularization algorithms map onto simple resistive networks. Exploiting Kirchhoff’s and Ohm’s law, they proved that the minimum of the regularized, quadratic cost functional is equivalent to the state of least power dissipation in an appropriate linear resistive network, where the data are given by injecting current into certain nodes and the solution by the stationary voltage distribution. Figure 1 shows the appropriate network for membrane-type surface interpolation, where the “strength” of smoothing is given by the value of the horizontal grid conductance. For an overview of various analog circuits for implementing early vision algorithms see Koch [21] and Horn [16]




Fig. 1. Resistive network for a membrane-type surface interpolation and smoothing. The circuit minimizes the variational functional of equation (1). In the continuum limit, minimization of this functional corresponds to the Euler-Lagrange equation  $-\lambda \nabla^2 f + Gf = Gd$ . The battery supplies the measured depth data  $d_i$ , while the vertical conductance  $G$  corresponds to  $1/(2\sigma^2)$  and the horizontal conductance of the grid to  $\lambda$ . If no data are present at the particular location  $i$ ,  $G$  is set to zero. The stationary voltage distribution then corresponds to the interpolated surface  $f_i$ . The amplitude of the horizontal grid conductance,  $\lambda$ , controls the amount of smoothing. A  $48 \times 48$  pixel hexagonal network implementing this circuit has been built and tested successfully [22, 26].

The recent development of subthreshold, analog CMOS (Complementary Metal Oxide Semiconductor) VLSI circuits for various sensory tasks by Carver Mead (see in particular his recent textbook [32]) has enabled us to implement these resistive networks—together with the photo-transduction stage—using this real-time, low-power, and robust technology. Two circuit elements are particularly attractive for our purposes: a photo-transistor with a logarithmic voltage output over five orders of intensity brightness [31, 32] and a transistor circuit with a quasi-linear current-voltage relationship [32, 41]. Using this as our basic construction element, we have built and tested a number of resistive networks—with up to  $48 \times 48$  pixels—for first- and second-order (membrane and thin-plate type) smooth surface interpolation [9, 22, 26] (see figure 1). All of our experimental chips are fabricated through the MOSIS silicon foundry, which is available at most universities.

We introduce in this paper an analog, purely deterministic, approach to locating discontinuities in the case of surface interpolation and optical flow estimation. It leads to a very simple and elegant circuit implementation in terms of a two-terminal, nonlinear, voltage-controlled resistor termed “resistive fuse” [10]. We have implemented this device in analog CMOS and demonstrate its performance here.

## 2 Theory

### 2.1 Smooth Surface Interpolation

Let us begin by analyzing the isomorphism between the quadratic variational functional for smooth surface interpolation and the dissipated power in linear resistive networks. Because our methodology does not distinguish between a 1D and a 2D implementation of smoothing in the presence of discontinuities, we will first consider the 1D case. The simplest possible variational functional for interpolating noisy and sparsely sampled data  $d_i$  is a first-order, membrane type of surface in interpolation:

$$J(f) = \lambda \sum_i (f_{i+1} - f_i)^2 + \frac{1}{2\sigma^2} \sum_i (d_i - f_i)^2 \quad (1)$$

where  $f_i$  is the value of the final surface  $f$  at location  $i$ ,  $\sigma^2$  is the variance of the additive Gaussian noise

process assumed to corrupt the data  $d_i$ , and  $\lambda$  is a free parameter. The first term in this functional implements the constraint that surfaces should vary smoothly. The second term in equation (1), where the sum includes only those locations  $i$  where data exist, forces the final solution  $f$  to be close to the measured data  $d$ . How close depends on the estimated magnitude of the noise, in this case on  $\sigma^2$ . Thus, the surface  $f$  minimizing (1) will be the one that best satisfies the conflicting demands of smoothness and fidelity to the measured data.

## 2.2 Quadratic Variational Functionals and Linear Resistive Networks

The quadratic variational functional can be minimized using the resistive network illustrated in figure 1, where the final smoothed surface is given by the stationary voltage distribution  $f_i$ , the observed data by the battery  $d_i$ , and the conductance  $G$  connecting the battery to the grid by  $1/(2\sigma^2)$ . If no measured surface value  $d_i$  is present at a particular location,  $G = 0$  at that location. The horizontal grid conductance  $\lambda$  controls the amount of smoothing: a large conductance value will facilitate the spread of current throughout the network while smaller conductance values will prevent this. Because the electrical power dissipated in a linear network is proportional to the square of the voltage across all resistors, the dissipated power can be identified with the variational functional of equation (1). *Maxwell's minimum heat theorem* [30] then implies that in such an idealized linear network the distribution of voltages and currents (subject to Kirchhoff's laws) minimizes the total power dissipated as heat.<sup>1</sup> In other words, the smoothest surface can be read as the stationary voltage distribution [37].

Mead has invented a saturating resistor element, whose current-voltage (I-V) relationship is linear for small voltage drops [32, 41]. The value of the slope, i.e., the conductance for small voltage drops can be varied over five orders of magnitude. The current through this circuit saturates for large values of the voltage. The measured I-V curve has the form of a hyperbolic tangent as shown in figure 2b. We have built and tested  $48 \times 48$  pixel resistive networks of the type shown in figure 1 on the basis of this circuit component

<sup>1</sup>Note that, in an *abuse de language*, the functional  $J$  of equation (1) is frequently termed "energy," while it actually represents "electrical power."




Fig. 2. Theoretical I-V curves for a linear resistor (a) and an experimentally measured I-V curve for Mead's [32] saturating resistor (b). Integrating numerically over these curves gives the co-content of the linear resistor (c) and the saturating resistor (d). Co-content is defined by equation (2) and represents generalized power for nonlinear systems. The co-content for the linear resistor is equivalent to half the dissipated power, and thus a quadratic function in  $\Delta V$  for all values of  $\Delta V$ , while the co-content for the saturating resistor is a quadratic function in  $\Delta V$  for small values of  $\Delta V$ , becoming a linear function of  $\Delta V$  as  $|\Delta V| \rightarrow \infty$ .

(for more details see [22]). But how will the implementation of ideal, linear resistors (figure 2a) with these saturating nonlinear resistors affect the final result? As pointed out by Poggio and Koch [37], the notion of minimizing power in linear networks implementing quadratic "regularized" algorithms must be replaced by the more general notion of minimizing the total resistor co-content [33] for nonlinear circuit elements. For a two-terminal voltage-controlled resistor characterized by  $I = f(V)$ , the co-content is defined as

$$J(V) = \int_0^V f(V') dV' \quad (2)$$

For a linear resistor,  $I = GV$ , the co-content is given by  $\frac{1}{2}GV^2$ , which is just half the dissipated power  $P = GV^2$  (figure 2c). For a network consisting of a collection of resistors, voltage sources, and other elements, the total network co-content  $J_{\text{total}}$  is defined as the sum of all the (linear or nonlinear) resistor co-contents. The

co-content for various resistors is plotted in figures 2 and 3. Differentiating equation (2), we have

$$f(V) = \frac{dJ}{dV} \quad (3)$$




Fig. 3. Simulated I-V curve for an infinite-gain fuse (a) and an experimentally measured I-V curve for a finite-gain resistive fuse (b). Integrating numerically over these curves gives the co-content  $J$  for the infinite-gain (c) and the finite-gain fuse (d).

It can be shown that a network of arbitrary topology consisting of strictly incrementally passive resistors (i.e., where  $dI/dV > 0$  for all  $V$ ) and ideal voltage and current sources has at most one solution, given by the unique minimum of the co-content [11]. Because the I-V curve of Mead's saturating resistor can be described quite accurately by  $I = I_0 \tanh(V/V_0)$  [32], its derivative is always positive. Thus, replacing ideal linear resistors with Mead's saturating resistors will not cause additional solutions to appear. Note that this treatment neglects the possible effect of the internal dynamics of the saturating resistor on the stability of the network.

### 2.3 Piecewise Smooth Surface Interpolation

Including binary line processes  $\ell$  into the functional of equation (1) leads to

$$J(f, \ell) = \lambda \sum_i (f_{i+1} - f_i)^2 (1 - \ell_i) + \frac{1}{2\sigma^2} \sum_i (d_i - f_i)^2 + \alpha \sum_i \ell_i \quad (4)$$

where  $\alpha$  is an additional free parameter. The first term in this functional implements the constraint that surfaces should vary smoothly for small values of the surface gradient. If all variables, with the exception of  $f_i$ ,  $f_{i+1}$ , and  $\ell_i$ , in equation (4) were held fixed and  $\lambda(f_i - f_{i+1})^2 < \alpha$ , it would be "cheaper" to pay the price  $\lambda(f_i - f_{i+1})^2$  and set  $\ell_i = 0$  than to pay the larger price  $\alpha$ . However, if the difference becomes too steep, the line process is switched on, i.e.,  $\ell_i = 1$ , and the "price"  $\alpha$  is paid. The functional of (4) is, different from (1), nonquadratic; and a large number of both stochastic and deterministic methods have been designed to find optimal or nearly optimal solutions for this and similar functionals [1, 6, 24, 29, 44].

Figure 3c shows a plot of  $J(f, \ell)$  as a function of  $f_i - f_{i+1}$ . The values of the surface and of the line discontinuities are assumed to be fixed at all other locations. As long as  $\lambda(f_{i+1} - f_i)^2 \leq \alpha$ , the function  $J$  is quadratic in the gradient. However, once  $|f_{i+1} - f_i|$  exceeds  $\sqrt{\alpha/\lambda}$ ,  $J$  remains constant. The term  $\sqrt{\alpha/\lambda}$  is similar to the gradient limit of Blake and Zisserman [1].

One straightforward manner to implement line discontinuities is via binary switches, breaking the resistive connections among neighboring nodes (see figure 5b in Koch, Marroquin, and Yuille [24], and figure 7a). Each switch has an associated digital processor opening and closing the switch in either a deterministic or a stochastic fashion, depending on the value of the voltage across the switch as well as on the states of neighboring switches. Such an analog-digital implementation is quite difficult to implement within conventional two-dimensional silicon circuits, due to the high amount of connectivity among nodes and the integration of a clocked, digital processor within non-clocked analog circuit components. However, as argued by Harris, Koch, Luo, and Wyatt [11], we can replace such a hybrid network by a single analog nonlinear resistor, the "resistive fuse." The appropriate current-voltage relationship of an *infinite-gain* or *binary resistive fuse* is illustrated in figure 3a. As long as the voltage drop across this device is below a threshold, the current through the nonlinear resistor is linearly related to the voltage across it. Once past the voltage threshold,

the circuit open circuits (hence the name “fuse”), and the current is zero. Unlike conventional electrical fuses, however, the resistive fuse can “reconnect” after open-circuiting as long as the voltage across the fuse drops again below the threshold. This two-terminal device therefore implements the high-level constraint that surfaces should be smooth unless their neighboring values differ by more than  $\pm\sqrt{\alpha/\lambda}$ , at which point the surface will break (see also Perona and Malik [36]).

The I-V relationship of the device we have built is shown in figure 3b. The most salient difference from the infinite-gain fuse are the smooth flanks where the current decreases smoothly to zero for increasing values of the voltage difference. The I-V curve of this device can be directly derived from a number of different early vision algorithms. Blake and Zisserman [1] first explicitly proposed the use of an energy functional for the problem of surface interpolation and segmentation from which the line process  $\ell$  is eliminated. Its form is identical to that illustrated in figure 3c. In order to find the minimum of this nonquadratic functional, Blake and Zisserman use a continuation method [35], their GNC algorithm. The idea behind this iterative method is to map the variational functional  $J(f, \ell)$  to be minimized onto a family of functionals  $J^*(f, \ell, t)$  with  $t \in [0, 1]$ , such that  $J^*(f, \ell, t = 0)$  is given by some convex functional and  $J^*(f, \ell, t = 1) = J(f, \ell)$ . Instead of directly attempting to minimize the non-convex functional  $J^*(f, \ell, 1)$ , continuation methods involve first finding the unique solution to the convex functional  $J^*(f, \ell, 0)$  and then to apply some smooth transform (parameterized by  $t$ ) to continuously deform  $J^*$ . The minimum of the functional  $J^*(f, \ell, t)$  is then used as a starting approximation when attempting to minimize  $J^*(f, \ell, t + \Delta t)$  until the minimum to the desired function  $J^*(f, \ell, 1) = J(f, \ell)$  is reached.<sup>2</sup> For the construction of  $J^*$ , Blake and Zisserman [1; eq. (7.23)] used a piecewise polynomial of order two, whose functional dependency (for a fixed value of  $t$ ) is similar to our measured co-content in figure 3d.

The analytical form of the I-V relationship of the “resistive fuse” element used in our simulations (see below) can be rigorously derived using a deterministic *mean field approximation* to the underlying stochastic Markov random field model of piecewise smooth

surface interpolation [5]. In this approximation, common to statistical mechanics [18], the interaction among neighboring values of  $f$  and  $\ell$  is replaced by the interaction among neighboring mean values of  $f$  and  $\ell$ . If we reinterpret the results of Geiger and Girosi [5] within our electrical circuit framework, we arrive at the constitutive relationship of our *finite-gain* or *analog resistive fuse*:

$$I = f(V) = \frac{\lambda V}{1 + e^{\beta(\lambda V^2 - \alpha)}} \quad (5)$$

where  $\beta$  is a free parameter depending on the inverse of the “temperature”  $T$  of the associated MRF algorithm.  $\beta$  can also be related to the imbedding parameter  $t$  in the continuation method. We had earlier derived  $f(V)$  of equation (5) in a somewhat ad hoc manner using the Hopfield and Tank [13] analog embedding for the binary line processes proposed in Koch, Marroquin, and Yuille [24]. It is quite satisfying that all these seemingly diverse approaches arrive at very similar results, which can all be implemented via a simple two-terminal nonlinear circuit.

As discussed above, the notion of minimizing the dissipated power in linear resistive networks must be replaced by the notion of minimizing the co-content (equation (3)) in nonlinear resistive networks.  $J_{\text{total}}$  is convex if and only if all resistors are incrementally passive, i.e., have positive slope everywhere [11]. This is the case for Mead’s saturating resistor. With incrementally active resistors, i.e., with regions of negative slope such as in our fuse, there will in general exist a number of stationary network solutions for a given input image, and uniqueness cannot be guaranteed anymore. However, if the *internal* dynamics of the incrementally active resistor circuit are neglected (that is, if we neglect the fact that our fuse circuit consists of transistors with their own, relatively fast, temporal behavior), for any voltage input and any initial condition the network will *not* oscillate indefinitely but must eventually settle to *some* stationary state (for more general stability results, see Wyatt and Standley [47]). This conclusion also holds even if parasitic (nonlinear but positive) capacitors are distributed arbitrarily throughout the network, provided there are no inductors. In other words, the total co-content of all resistors  $J_{\text{total}}(t)$  will act as a Lyapunov function [11].

<sup>2</sup>Intuitively, such a method can also be thought of as “deterministic annealing” [5].

### 3 Circuit Performance

We will not discuss the implementation of the fuse circuit (see Harris, Koch, Luo, and Wyatt [1] for details), but only demonstrate some of the capabilities of our analog resistive fuse circuit. Its measured I-V curve as well as the associated co-content is plotted in figure 3. The I-V relationship of each fuse circuit can be controlled by various voltage input controls. This allows us to change continuously from a Mead type of saturating resistor to a finite-gain fuse (figure 4). In other words, by turning the voltage control, we are




Fig. 4. Measured I-V curves that show the effect of varying an external voltage control ( $V_A$ ) from the saturating characteristic of Mead’s resistor to that of the fuse. This effectively corresponds to the implementation of a continuation method [35].

transforming the “energy” landscape in a continuous fashion from containing one unique global minimum to a landscape containing many local minima, thereby directly implementing a class of continuation methods. We also control the threshold value for which the fuse ceases to generate a current (figure 5). Because this value corresponds to the threshold  $\pm\sqrt{\alpha/\lambda}$  in an infinite-gain fuse (figure 3a), we have effective control over  $\alpha$ , determining where surface discontinuities will be introduced.

A linear network of eight analog fuses implementing 1D piecewise smooth surface interpolation was fabricated via Mosis. Figure 6 shows a segmentation result for a “noisy” 1D step edge. The network effectively smoothes out small steps without degrading large step edges<sup>3</sup>. We have experimented with several different

<sup>3</sup>We have duplicated this result using Mead’s photoreceptors to provide voltage input to the network. An edge imaged onto the chip is successfully smoothed and segmented.




Fig. 5. Measured I-V curves illustrating different line process penalties through adjustment of  $V_B$ . This is paramount to controlling the threshold  $\alpha$  in equation (4) for introducing discontinuities.




Fig. 6. Measured segmentation from an 8 pixel experimental analog resistive fuse network. The circles denote “noisy” step data that was used as the input to the network. The solid-line curve indicates measured voltages from the chip. The dotted-line curve shows the measured voltage output given by a network of Mead’s saturating resistors.

continuation methods by using the controls illustrated in figures 4 and 5. The specific continuation method used for this example was to gradually increase the amount of smoothing provided by the network, while simultaneously decreasing the threshold  $\alpha$  for breaking. The final I-V curves of the fuses in this example have the form shown in figure 3b. We have also successfully implemented a network of binary fuses (with I-V curves of the form shown in figure 3a) and are considering the computational differences between networks of binary and analog fuses.

## 4 Computing Optical Flow

We wish to demonstrate how the analog fuse will work for the problem of estimating the 2D optical flow in the presence of motion discontinuities. The optical flow is the vector field  $(u, v)$  derived from the changing image brightness  $I(x, y, t)$ . It differs, in general, from the underlying velocity field, a purely geometrical concept [15, 45]. Following Horn and Schunck [17], we assume that the brightness  $I(x, y, t)$  of the image is constant over time; in other words  $dI/dt = I_x u + I_y v + I_t = 0$ , where the velocity  $(u, v) = (dx/dt, dy/dt)$  and  $I_x, I_y$ , and  $I_t$  are the partial derivatives of the brightness  $I$ . This manner of formulating the optical flow is ill posed (“aperture problem”), due to the presence of two unknowns  $u$  and  $v$  in a single linear equation.<sup>4</sup>

### 4.1 Smooth Optical Flow

Horn and Schunck [17] “regularized” this ill-posed problem by imposing a smoothness constraint. Because objects, in general, are continuous and undergo smooth movements, neighboring points in the world will have similar velocities. The projected velocity field should reflect this fact. As the measure of smoothness Horn and Schunck chose the square of the velocity field gradient leading to a velocity field which minimizes

$$E(u, v) = \iint (I_x u + I_y v + I_t)^2 + \lambda \left[ \left( \frac{\partial u}{\partial x} \right)^2 + \left( \frac{\partial u}{\partial y} \right)^2 + \left( \frac{\partial v}{\partial x} \right)^2 + \left( \frac{\partial v}{\partial y} \right)^2 \right] dx dy \quad (6)$$

where the regularization parameter  $\lambda$  is inversely dependent on the signal-to-noise ratio. The first term describes the fact that the final solution should follow as closely as possible the measured data, whereas the second term imposes the smoothness constraint on the solution. The degree of minimization of one or the other term is governed by  $\lambda$ . Because  $E$  is quadratic in  $u$  and  $v$ , the associated Euler-Lagrange equations, discretized on a square lattice, will be linear:

$$I_x^2 u_{ij} + I_x I_y v_{ij} + I_x I_t - I_{ij} + \lambda(u_{ij} - u_{i+1j}) + \lambda(u_{ij} - u_{i-1j}) + \lambda(u_{ij} - u_{ij+1}) + \lambda(u_{ij} - u_{ij-1}) = 0 \quad (7)$$

$$I_x I_y v_{ij} + I_y^2 v_{ij} + I_y I_t - I_{ij} + \lambda(v_{ij} - v_{i+1j}) + \lambda(v_{ij} - v_{i-1j}) + \lambda(v_{ij} - v_{ij+1}) + \lambda(v_{ij} - v_{ij-1}) = 0$$

Following Poggio and Koch [37], we can map these equations directly onto a linear resistive network (figure 7a) if we identify

$$\begin{aligned} T &\rightarrow \lambda \\ G &\rightarrow -I_x I_y \\ G_x &\rightarrow I_x (I_x + I_y) \\ G_y &\rightarrow I_y (I_x + I_y) \\ E &\rightarrow \frac{-I_t}{I_x + I_y} \end{aligned} \quad (8)$$

This network will then settle into the state of least power dissipation following Kirchhoff’s current law, expressed in the Euler-Lagrange equations (7). The associated stationary voltages correspond to the solution sought:  $u_{ij}$  is equivalent to the  $x$  component and  $v_{ij}$  to the  $y$  component of the optical flow field (for more details see Hutchinson, Koch, Luo, and Mead [19]). A much simplified version of such a network has been implemented using analog CMOS [42]; see also [32]. Tanner’s chip with an array of  $8 \times 8$  phototransistors computes the constant optical flow induced by rigid body motion (i.e., it can act as an “optical mouse”) and essentially corresponds to the network in figure 7a with all the transversal conductances  $T$  set to a very high value and all switches removed.

We simulated the behavior of this linear optical flow network on a 16 node hypercube using natural imagery acquired via a video camera. Figures 8a and b show a sequence of several moving people. Figure 8c illustrates the solution of the linear resistive network: the smoothness assumption of Horn and Schunck leads to the qualitatively correct optical flow field. However, at the occluding boundary between the two rightmost people, the estimated flow field is zero due to the averaging effect of the smoothness assumption. Furthermore, the algorithm smoothes over the motion discontinuities along the contour of the people.

<sup>4</sup>Recently, Verri, Girosi, and Torre [46] have proposed a rather elegant solution to the aperture problem. However, for demonstrating the use of “fuses,” the formulation of the optical flow problem presented here is entirely sufficient.




Fig. 7a. Resistive network for computing optical flow with binary switches. The conductances  $G$  as well as  $G_x$  and  $G_y$  (for clarity, only two such elements are shown) depend on the measured spatial and temporal image gradients, as well as the battery  $E$  (see equation (8)). The conductance  $T$  (corresponding to  $\lambda$ ) controls the amount of smoothing. On-chip photoreceptors compute  $I_x$ ,  $I_y$ , and  $I_t$  and control the values of  $G$ ,  $G_x$ ,  $G_y$ , and  $E$  according to equation (8). The  $x$  component of the velocity,  $u$ , is given by the voltage in the top network, while the  $y$  component of the velocity,  $v$ , is given by the voltage in the bottom network. Digital processors, not shown, control the opening and closing of the switches. The parasitic capacitances  $C$  govern the dynamic behavior of the circuit.

#### 4.2 Piecewise Smooth Optical Flow

In a previous publication [19] we showed how binary motion discontinuities can be incorporated into the optical flow estimation functional of equation (6). We placed a number of constraints on the formation of line processes, most notably that motion discontinuities should form along continuous contours, should not intersect, and should coincide in general with intensity discontinuities [4]. The state of the binary line processes was evaluated using the steepest descent rule. The network implementation we envisaged at that time is shown in figure 7a, where the binary line processes

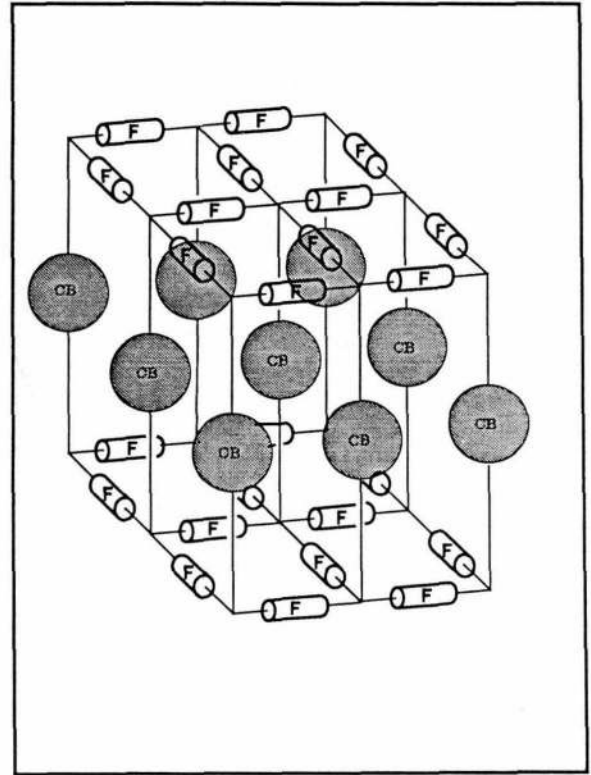


Fig. 7b. Proposed resistive network for computing optical flow with resistive fuses. The binary switches and horizontal resistors of figure 7a have been replaced by resistive fuse elements (shown as cylinders with a schematic "F"). Constraint boxes (indicated with shaded circles) have replaced the data-dependent resistors and batteries of figure 7a. As described in the text, every constraint box minimizes the local square error of the image brightness constraint equation  $dI/dt = 0$  [8].

are implemented via binary switches. Digital processors (not shown) control the state of these switches (see figure 4b in Koch, Marroquin, and Yuille [24]). Such a scheme requires a sort of hybrid technology, exploiting both analog as well as digital circuits. However, we will now replace the resistance  $T$ , the binary switch as well as its digital "controller," by our nonlinear fuse circuit. Following the above discussion for piecewise smooth surface interpolation, we simply replace the current through the linear resistor (of the form  $\lambda(u_{ij} - u_{i+1j})$ ) by the current through the non-linear fuse (of the form  $f(u_{ij} - u_{i+1j})$ ), where  $I = f(V)$  is the constitutive relationship of the fuse (equation (5)). We then have



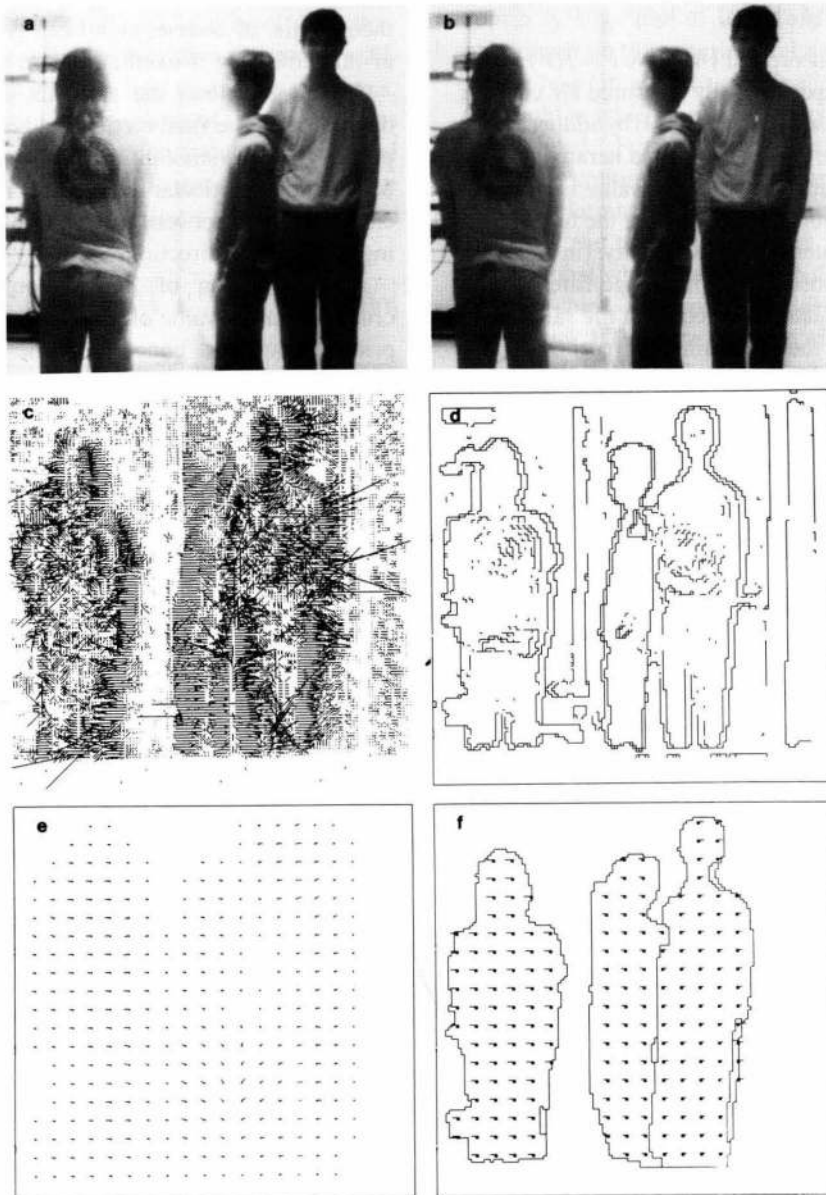


Figure 8. Computed optical flow of several moving people. (a) and (b) show the two  $128 \times 128$  pixel images captured by a video camera. The two leftward people move toward the left while the rightward person moves to the right. (c) The initial, local velocity data prior to smoothing. (d) the zero-crossings of the Laplacian of a Gaussian of both images; the zero-crossings are thresholded to remove noise. (e) Smooth and subsampled optical flow obtained by numerically solving equation (7b). Two undifferentiated “blobs” move to the left and one moves to the right. (f) The subsampled optical flow computed with analog resistive fuses. In order to visualize the behavior of the fuses, we indicated their state with a solid line if the voltage difference across them exceeds the value for which their conductance is half of the conductance for the zero voltage case. The fuses are prevented from breaking at locations where no zero-crossings are present.

$$\begin{aligned}
 & I_{x-ij}^2 u_{ij} + I_{x-ij} I_{y-ij} v_{ij} + I_{x-ij} I_{t-ij} \\
 & + f(u_{ij} - u_{i+j}) + f(u_{ij} - u_{ij+1}) \\
 & + f(u_{ij} - u_{i-j}) + f(u_{ij} - u_{ij-1}) = 0 \\
 & I_{x-ij} I_{y-ij} u_{ij} + I_{y-ij}^2 v_{ij} + I_{y-ij} I_{t-ij} \\
 & + f(v_{ij} - v_{i+j}) + f(v_{ij} - v_{ij+1}) \\
 & + f(v_{ij} - v_{i-j}) + f(v_{ij} - v_{ij-1}) = 0
 \end{aligned} \tag{9}$$

Figure 9 shows the theoretical I-V curve  $I = f(V)$  superimposed onto the experimentally measured I-V curve of our analog fuse. We solve equations (9) by adding a fixed parasitic capacitance at every node and iteratively computing the voltage until a stationary value is reached.

We use the additional constraint that the fuse should only break if an intensity discontinuity (in our case thresholded zero-crossings of the image filtered with the Laplacian of a Gaussian) occurs at the same location, implementing the constraint that motion discontinuities usually coincide with intensity discontinuities [4]. We achieve this by increasing the threshold  $\alpha$  by a factor 100 if no zero-crossing is present, effectively preventing the fuse from “breaking” in the absence of intensity edges. The fusion of motion and intensity discontinuities could be accomplished at the circuit level by biasing the voltage control for the current cutoff on the fuse circuit,  $V_B$ , with the output of a second chip computing zero-crossings of the  $\nabla^2 G$  operator.




Fig. 9. Comparison of the I-V curve used in our simulations (dotted line given by equation (5)) with the experimentally measured curve of our fuse circuit (solid line; see also figure 3b).

Traditionally, additional constraint terms have been used embodying the fact that, in general, motion discontinuities occur along extended contours and rarely intersect. However, because zero-crossings are located on continuous contours and do not intersect, we have found these terms to be superfluous and do not use them. This, of course, significantly reduces the cost of implementing discontinuities in analog hardware.

Figure 8f shows the resultant stationary voltage distribution. The final motion field is substantially improved over the smooth optical flow shown in figure 8e. Notice in particular how the algorithm finds the occluding contour between the two rightward people moving in opposite directions.

The formation of line discontinuities depends crucially on the value of  $\alpha$ . Because we assume no a priori knowledge concerning the distribution and amplitude of any motion discontinuities, we adopt the heuristic procedure, illustrated in figure 10, of multiplying the penalty for a line process,  $\alpha$  (see equation (4)), with  $1/K(t)$ , where  $K(t)$  increases from a small value to a large one. This has the effect that the system initially detects only very large discontinuities in the optical flow. As time goes on, the effective  $\alpha$  decreases, enabling the detection of smaller motion discontinuities. In the final iterations, we reduce  $K(t)$  to remove spurious or isolated motion discontinuities. As illustrated in figure 5, we can vary  $\alpha$  on our fuse circuit.




Fig. 10. Variation in the effective threshold  $\alpha/K$  for determining when a line process should be introduced as a function of the iteration time. This schedule corresponds to requiring a very large gradient to locate a discontinuity initially, subsequently facilitating detection of discontinuities. A cleanup cycle at the end eliminates unwanted discontinuities.

### 4.3 Constraint Boxes and Optical Flow

The network depicted in figure 7a is extremely difficult to implement in analog VLSI. First of all, the battery voltages and resistance values are relatively complicated functions of the input data. Furthermore, because the resistance values are functions of the input data, they may be positive or negative. Negative resistors are difficult to implement and may lead to stability problems. We have developed a methodology of constraint boxes for implementing most of the constraint terms required in early vision problems [8]. In general, suppose we would like to minimize a quadratic constraint on the voltages  $V_k$  of the form

$$F^2(V_1, \dots, V_k, \dots, V_n) \quad (10)$$

where  $F$  is an expression linear in the  $V_k$ 's. A constraint box is an  $n$  terminal, fully reciprocal device with the current  $I_k$  at every node (with an associated voltage  $V_k$ ) satisfying

$$i_k = -F \frac{\partial F}{\partial V_k} \quad (11)$$

It is straightforward to show that this choice of currents will minimize the constraint term  $F^2$ . Furthermore, a network of constraint boxes will minimize the sum of the squares of the constraints [8]. Constraint boxes have been successfully used to implement the coupled depth/slope method for thin plate surface reconstruction [9].

For the problem of estimating the optical flow, we have

$$F = I_x u + I_y v + I_t \quad (12)$$

Thus, we need to design a constraint box injecting the following current into the  $u$  and  $v$  resistor meshes:

$$\begin{aligned} i_u &= -F \frac{\partial F}{\partial u} = -I_x(I_x u + I_y v + I_t) \\ i_v &= -F \frac{\partial F}{\partial v} = -I_y(I_x u + I_y v + I_t) \end{aligned} \quad (13)$$

Note that no additional circuitry beside a single constraint box per node is required (figure 7b). Thus, constraint boxes significantly reduce the circuit complexity and ameliorate the stability problem (because no negative resistors are required).

### 5 Conclusion

We have successfully demonstrated a simple and elegant analog circuit implementation of the line discontinuities of Geman and Geman [6] and of the graduated non-convexity algorithm of Blake and Zisserman [1]. We only report on the experimental data for an 8-pixel 1D analog fuse circuit. Our circuit allows us to modify the threshold at which the fuse breaks and a discontinuity "is detected." Furthermore, we have implemented one particular continuation method for smoothly changing the I-V relationship from that of a saturating resistor to that of the nonlinear resistive fuse. We found this to be critical when attempting to detect discontinuities with our chips. We also successfully designed and tested an 8 pixel 1D binary fuse circuit as well as a two-dimensional, 20×20 pixel version of this network. Our current implementation medium is subthreshold analog CMOS, using a 2μm process. One of the key advantages of this technology is that a significant number of computational primitives are available for our use [32]. Note that the best technology in which to ultimately implement these circuits will critically depend upon the requirements of the vision application in terms of power, speed, system size, cost, and required accuracy. It is not conceptually difficult, for example, to map our circuit ideas to a bipolar process where our circuits would gain the advantages of higher speed and better accuracy at the cost of lower circuit density and much higher power dissipation.

In agreement with other researchers in the field, we have found empirically that deterministic implementations of the stochastic Markov random field algorithms of early vision lead to solutions little different from those obtained with the more time-consuming stochastic optimization techniques. Furthermore, the fact that we require no additional constraints on the formation of discontinuities, except that they should coincide with intensity discontinuities, makes their implementation into dedicated hardware feasible. Given the provisional nature of our first chips, we have not yet attempted to characterize their accuracy, although we expect their resolution to be within 5 to 7 bits. We have demonstrated already, however, that these circuits are inherently robust and accurate enough to allow for simple navigation tasks such as following edges or tracking moving light sources when mounted onto small, highly mobile toy cars operating in a laboratory environment [23].

We previously demonstrated chips implementing smooth two-dimensional first-order surface interpolation [26] and a one-dimensional second-order or thin-plate surface interpolation [9]. Computer simulations have shown that detection of discontinuities in surface orientation, such as occurring along creases, is feasible in problems such as edge detection and surface interpolation [1, 25] and can be incorporated into our thin-plane interpolation circuits [9].

We thus have all the elementary circuit elements in hand—phototransistors for on-chip image acquisition [32], resistive networks for smoothing, and resistive fuses for detecting discontinuities—to build analog chips to compute optical flow fields in the presence of motion and intensity discontinuities.

### Acknowledgements

We foremost wish to thank Carver Mead for laying the framework upon which we have built our research. Our theoretical ideas would never have been ported into silicon without him. All chips were fabricated through MOSIS with DARPA's support. Research in the laboratory of C.K. is supported by a National Science Foundation grant IST-8700064, Office of Naval Research Young Investigator and National Science Foundation Presidential Young Investigator Awards, the James S. McDonnell Foundation, Rockwell International Science Center and the Hughes Aircraft Artificial Intelligence Center. J.G.H. is funded by a Hughes Aircraft Fellowship.

### References

1. A. Blake and A. Zisserman. *Visual Reconstruction*. MIT Press: Cambridge, MA, 1987.
2. A. Blake. "Comparison of the efficiency of deterministic and stochastic algorithms for visual reconstruction," *IEEE Trans. PAMI* 11:2-12, 1989.
3. P.B. Chou and C. Brown. "Multimodal reconstruction and segmentation with Markov random fields and HCF optimization," *Proc. Image Understanding Workshop*, pp. 214-221, Cambridge, MA, February, 1988.
4. E. Gamble and T. Poggio. "Integration of intensity edges with stereo and motion," *Artif. Intell. Lab Memo* No. 970, MIT, Cambridge, 1987.
5. D. Geiger and F. Girosi. "Parallel and deterministic algorithms for MRFs: surface reconstruction and integration," *AI Memo* No. 1114, MIT, Cambridge, 1989.
6. S. Geman and D. Geman. "Stochastic relaxation, Gibbs distribution and Bayesian restoration of images," *IEEE Trans. PAMI* 6:721-741, 1984.
7. W.E.L. Grimson. *From Images to Surfaces*. MIT Press: Cambridge, MA, 1981.
8. J.G. Harris. "Solving early vision problems with VLSI constraint networks," *AAAI Neural Architectures for Computer Vision Workshop*, Minneapolis, August 20, 1988.
9. J.G. Harris. "An analog VLSI chip for thin plate surface interpolation." In *Neural Information Processing Systems*, D. Touretzky (ed.). Morgan Kaufmann: Palo Alto, 1989.
10. J.G. Harris and C. Koch. "Resistive fuses: circuit implementations of line discontinuities in vision," *Snowbird Neural Network Workshop*, April 4-7, 1989.
11. J.G. Harris, C. Koch, J. Luo, and J. Wyatt. "Resistive fuses: Analog hardware for detecting discontinuities in early vision." In *Analog VLSI Implementations of Neural Systems*, pp. 27-56. C. Mead and M. Ismail (eds.). Kluwer: Norwell, MA, 1989.
12. E.C. Hildreth. *The Measurement of Visual Motion*. MIT Press: Cambridge, MA, 1984.
13. J.J. Hopfield and D.W. Tank. "Neural computation in optimization problems," *Biological Cybernetics* 52:141-152, 1985.
14. B.K.P. Horn. "Determining lightness from an image," *Comput. Graphics Image Process.* 3(1): 277-299, 1974.
15. B.K.P. Horn. *Robot Vision*. MIT Press: Cambridge, 1986.
16. B.K.P. Horn. "Parallel networks for machine vision," *Artif. Intell. Lab. Memo*, 1071, MIT, Cambridge, 1988.
17. B.K.P. Horn and B.G. Schunk. "Determining optical flow," *Artificial Intelligence* 17:185-203, 1981.
18. K. Huang. *Statistical Mechanics*. Wiley: New York, 1963.
19. J. Hutchinson, C. Koch, J. Luo, and C. Mead. "Computing motion using analog and binary resistive networks," *IEEE Computer* 21:52-63, 1988.
20. K. Ikeuchi and B.K.P. Horn. "Numerical shape from shading and occluding boundaries," *Artificial Intelligence* 17:141-184, 1981.
21. C. Koch. "Seeing chips: analog VLSI circuits for computer vision." *Neural Computation* 1:184-200, 1989.
22. C. Koch, J. Luo, J. Hutchinson, and C. Mead. "Optical flow and surface interpolation in resistive networks: Algorithms and analog VLSI chips." In *Neural Networks*, B. Shriver (ed.). IEEE Computer Science Press Book, in press, 1990.
23. C. Koch, J. Harris, T. Horiuchi, A. Hsu, and J. Luo. "Real-time computer vision and robotics using analog VLSI circuits," *Neural Inform. Process. Systems Conf.*, Denver, November, 1989.
24. C. Koch, J. Marroquin, and A. Yuille. "Analog neuronal networks in early vision," *Proc. Natl. Acad. Sci. USA* 83:4263-4267, 1986.
25. S.C. Liu and J.G. Harris. "Generalized smoothing networks in solving early vision problems," *Comput. Vision Pattern Recog. Conf.*, 1989
26. J. Luo, C. Koch, and C. Mead. "An experimental subthreshold, analog CMOS two-dimensional surface interpolation circuit," *Neural Inform. Process. Systems Conf.*, Denver, November, 1988.
27. D. Marr and T. Poggio. "Cooperative computation of stereo disparity," *Science* 194:283-287, 1976.
28. J.L. Marroquin. "Probabilistic solution of inverse problems," *AI Lab Memo* No. 860, MIT, Cambridge, 1985.

29. J. Marroquin, S. Mitter, and T. Poggio. "Probabilistic solution of ill-posed problems in computational vision," *J. Amer. Statistic. Assoc.* 82:76-89, 1987.
30. J.C. Maxwell. *A Treatise on Electricity and Magnetism*, 3rd ed., vol. I, pp. 407-408, 1891. Republished by Dover Publications: New York, 1954.
31. C.A. Mead. "A sensitive electronic photoreceptor." In *1985 Chapel Hill Conference on Very Large Scale Integration*, pp. 463-471, 1985.
32. C.A. Mead. *Analog VLSI and Neural Systems*. Addison-Wesley: Reading, MA, 1989.
33. W. Millar. "Some general theorems for non-linear systems possessing resistance," *Philosophical Magazine* 42:1150-1160, 1951.
34. H.H. Nagel. "On the estimation of optical flow: Relations between different approaches and some new results," *Artificial Intelligence* 33:299-324, 1987.
35. J.M. Ortega and W.C. Rheinboldt. *Iterative Solution of Nonlinear Equations in Several Variables*. Academic Press: New York, 1970.
36. P. Perona and J. Malik. "A network for multiscale image segmentation," *Proc. 1988 IEEE Intern. Symp. on Circuits and Systems*, pp. 2565-2568, Espoo, Finland, June, 1988.
37. T. Poggio and C. Koch. "Ill-posed problems in early vision: From computational theory to analogue networks," *Proc. Roy. Soc. Lond. B* 226:303-323, 1985.
38. T. Poggio, E.B. Gamble, and J.J. Little. "Parallel integration of vision modules," *Science* 242:436-440, 1988.
39. T. Poggio, V. Torre, and C. Koch. "Computational vision and regularization theory," *Nature* 317:314-319, 1985.
40. T. Poggio, H. Voorhees, and A. Yuille. "A regularized solution to edge detection," *Artif. Intell. Lab Memo No. 833*, MIT, Cambridge, 1986.
41. M.A. Sivilotti, M.A. Mahowald, and C.A. Mead. "Real-time visual computation using analog CMOS processing arrays." In *1987 Stanford Conf. VLSI*, pp. 295-312. MIT Press: Cambridge, 1987.
42. J.E. Tanner. "Integrated optical motion detection." Ph.D. Thesis, Dept. of Computer Science, Caltech, 1986.
43. D. Terzopoulos. "Multilevel computational processes for visual surface reconstruction," *Comput. Vision Graph. Image Process.* 24:52-96, 1983.
44. D. Terzopoulos. "Regularization of inverse problems involving discontinuities," *IEEE Trans. PAMI* 8:413-424, 1986
45. A. Verri and T. Poggio. "Motion field and optical flow: Qualitative properties," *IEEE Trans. PAMI* 11:490-498, 1989.
46. A. Verri, F. Girosi, and V. Torre. "Mathematical properties of the two dimensional motion field: From singular points to motion parameters," *J. Opt. Soc. Amer. A* 6:469-712, 1989.
47. J.L. Wyatt, Jr. and D.L. Standley. "Criteria for robust stability in a class of lateral inhibition networks coupled through resistive grids," *Neural Computation* 1:58-67, 1989
48. A. Yuille and N.M. Grzywacz. "A computational theory for the perception of coherent visual motion," *Nature* 333:71-73, 1988.
49. Y.T. Zhou and R. Chellappa. "Computation of optical flow using a neural network." *Proc IEEE Intern. Conf. Neural Networks* 2:71-78, San Diego, CA, July, 1988.