Real Time Image Processing: MFA in VLSI

Griff L. Bilbro Lester C. Hall Mark Clements Wentai Liu
Department of Electrical and Computer Engineering
North Carolina State University, Raleigh, NC

September 25, 1995

Abstract

We report three results in an investigation of the application of analog VLSI to real-time image processing, especially for advanced non-linear image optimization operations implementation on general purpose digital computers. First, we use SPICE to simulate parallel analog circuits that use Mean Field Annealing (MFA) to perform nonlinear signal processing. Second, we present the PCB realization of a 1-D analog deconvolution circuit using standard ICs and discrete components on PCB. Finally, we describe a MAGIC layout of a 64 pixel IC for image processing in 2-D that is being fabricated at MOSIS. We conclude that algorithms such as MFA may be useful for real time interactive imaging in analog VLSI.

1 Introduction

Real time image processing is often limited to simple operations such as sums of products, linear convolutions, and median filters. Interesting operations such as nonlinear MAP restorations[3], and important nonlinear image formation models such as in MRI[5, 4], are numerically too intensive for real time on serial processors or even on 64-node digital hypercube parallel architectures.

2 Theory

Images degraded by imperfect sensors or transmission can be modeled usefully[8] as the sum of random noise image and a convolution of the true image with a point spread function or blurring kernel. We will write

$$g = h * f + n \tag{1}$$

for the degraded image so that g_i for $i = 1, 2, ..., L$ is the i^{th} pixel of L pixels in image g. Blurring is represented by the convolution of kernel h with the true signal f by the $*$ operator, defined by the relation

$$(h * f)_i = \sum_j h_{i-j} f_j \tag{2}$$

so that the i^{th} pixel of the convolved image $h * f$ is a weighted sum of pixels of f.

The additive noise n comprises n_i where $i = 1, 2, ..., L$. We will show that n can represent missing data as well as the usual case of Gaussian noise. For subsampled images or images in which some of the pixel values are missing, $1/\sigma_i = 0$ if the value g_i for pixel i is missing. For independent, identically distributed Gaussian noise, the standard deviation $\sigma_i = \sigma, \quad \forall i$. We can restrict ourselves to the zero-mean case, $\langle n_i \rangle = 0$ for $i = 1, 2, ..., L$.

The simplest approach to estimating f given the blurred image g and the blurring kernel h is to attempt to ignore the noise n and solve

$$g = h * f \tag{3}$$

in the Fourier domain. The resulting inverse filter is unstable[8]. The Weiner filter is stable but blurs sharp edges[8]. Edges can be preserved by nonlinear Bayesian techniques, but these require global optimization such as simulated annealing[6] or mean field annealing[2, 7] The potential for implementing Mean Field Annealing (MFA) as a neural network in analog integrated circuits has been recognized previously[1] even though MFA is currently simulated on general purpose digital computers.

MFA derives from Simulated Annealing (SA) but differs from SA by analytically approximating the relevant Gibbs distribution rather than stochastically simulating it. SA works by gradually cooling an on-going stochastic simulation of a Gibbs distribution. In MFA the approximate distribution is cooled. MFA often produces much better results than other deterministic approaches and is faster than stochastic approaches, but is not guaranteed to produce the global optimum.

3 Circuit Simulations

Using the simplest spatial inverse filters for deconvolution using inverters for gain and level 2 models of 2μ CMOS transistors extracted from MOSIS fabrication runs, we simulated several circuits for parallel signal processing.

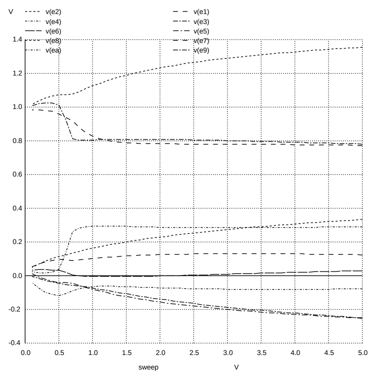

Figure 1: Plot of 10 pixel estimate of a blurred noisy image annealed from all 10 degraded measurements.

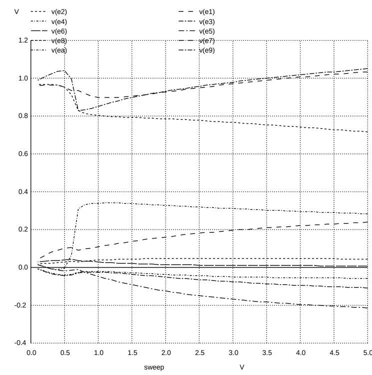

Figure 2: Plot of 10 pixel estimate from a blurred noisy image with half the data missing (every other pixel).

3.1 Inverse filter

The spatial domain analog of an inverse filter which attempts to solve $g = h * f$ for f at each pixel. We designed and simulated a circuit to do this by connecting a resistive network model of the convolution in the negative feedback path of a differential amplifier whose noninverting input is the blurred image. Simulation shows that the signal to noise can been raised enough to reduce the tails of the estimated f by about over one order of magnitude. In the next section, we describe a physical PCB realization of this circuit.

A superior restoration can be obtained using MFA. In this approach to image processing[1], we define the optimal estimate as the image that minimizes

$$H[f] = \sum_i \left(a \; |g - h * f|_i^2 + b \; R[f]_i \right) \tag{4}$$

where a and b are problem dependent constants and the sum is over all pixels. Here R is the piecewise constant regularization term that will be annealed on temperature T. In one dimension at pixel i,

$$R[f]_i = -\exp \frac{(D_L * f)_i^2}{T} - \exp \frac{(D_R * f)_i^2}{T} \tag{5}$$

where D_L is a left difference kernel, so that $(D_L * f)_i = f_i - f_{i-1}$ and D_R is a right difference kernel, so that $(D_R * f)_i = f_{i+1} - f_i$

Figure 1 shows the result for a 10 pixel image of a scene with a foreground feature of unit intensity comprising pixels 1, 2, and 3 and a larger background region of zero intensity comprising pixels 4...10 where pixel 10 is labeled with the letter a. The measured inputs were blurred by convolving with a kernel of radius 2 pixels and then randomly displaced with noise of approximate amplitude .1. Temperature decreases in steps of -.1 from T=5 on the right to T=.1 on left. The output is represented by an array of ten voltages $e_1 ... e_a$ corresponding to estimates of the ten signals $f_1 ... f_{10}$. The final estimates at low temperatures near the left of the plot show excellent agreement with the true values of unity for $i = 1, 2, 3$ and zero for the remaining $i = 4...10$.

For the case of *missing data*, this circuit can be modified to interpolate as discussed in the Theory section:

$$H[f] = \sum_{i=1,3,5,\ldots} a \; |g - h * f|_i^2 + \sum_i a \; R[f]_i \tag{6}$$

where the first sum is over only the odd pixels, but the second sum is over all pixels. Figure 2 shows the result of the same input signal, blur, and noise for the case of the additional degradation of missing pixels. The restoration is still excellent.

4 PCB Realization

We implemented a 16 pixel 1-D deconvolution analog PCB circuit with user-adjustable inputs as shown in Figure 3. The circuit consists of three PCB modules, 3 inches by 9 inches. Input power conditioning is provided by an LM317T linear voltage regulator. All other active circuit functions are performed by LM324 quad single supply low-power op-amps. A 9V battery provides power. Using 16 potentiometers, the user can specify the input signal which is displayed on the leftmost array of red LEDs. The image is blurred on the middle board. The final deblurred estimate is visually identical to the original input.

The effect of high frequency spatial noise is not evident in this circuit until the convolution kernel is aproximately 10 pixels in radius. This actual performance appears to be substantially superior to that expected from simple simulations using ideal opamps.

The most surprising result from this realized inverse filter was its excellent performance for wide kernels. We pessimistically expected sharp input edges to be unrecoverable after blurring by kernels wider than one or two

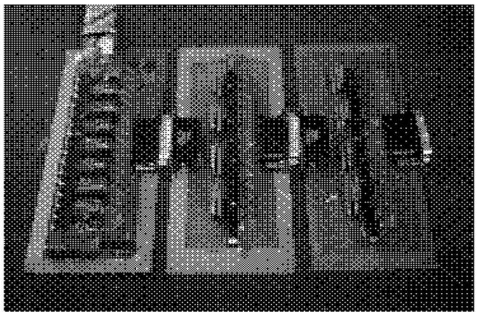

Figure 3: Photograph of a 16-pixel PCB implementation of an inverse filter deconvolution.

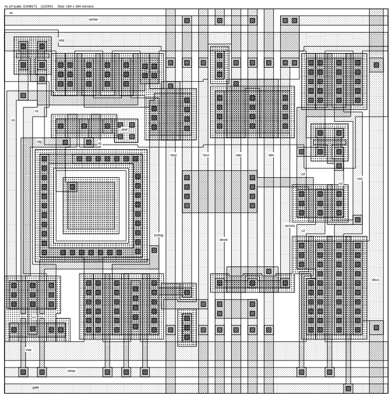

Figure 4: MAGIC layout of one pixel of user-specified low-pass image sensor.

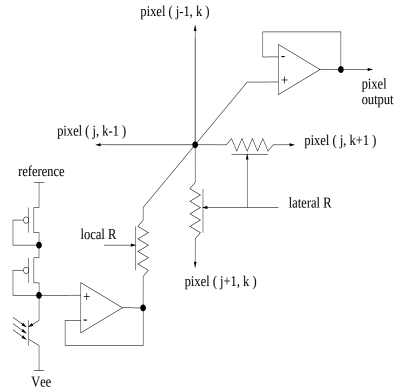

Figure 5: Schematic for the $i\ j$ pixel and connections to neighboring pixels.

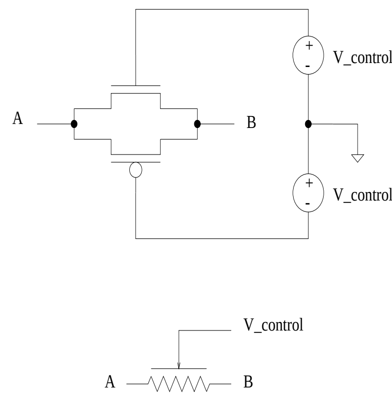

Figure 6: Schematic for the voltage controlled resistors which provide user control over the convolution kernel.

pixels, but we found that good deconvolution quality persisted for kernels with radius up to about 10 pixels in the presence of considerable noise (about ten percent of the maximum signals.

5 Analog IC Layout

We designed a chip that performs a 2-D convolution of a 2-D image using analog processing elements in each pixel shown schematically as in Figure 5. The calculations for all the pixels are performed in parallel. Voltage controlled local and lateral resistors are used to determine the convolution kernel. We laid out the chip using MAGIC as shown in Figure 4. The chip contains an 8x8 array of pixels, and is implemented in a 2 μm CMOS N-well process.

Each pixel is comprised of a buffered photodetector circuit, a local voltage controlled resistor, two lateral voltage controlled resistors, and a pixel output buffer amplifier.

The photodetector circuit[9] consists of a vertical PNP transistor which is loaded by two series diode-connected PMOS devices as shown in Figure 5. Incident photons create electron-hole pairs in the bipolar device, resulting in a photocurrent proportional to the light intensity. The photocurrent drives the two load devices, and the voltage across them is used as the output.

The photodetector output must be buffered, so that it sees too small a DC load to significantly change the output voltage. We used a simple two-stage CMOS operational amplifier in a unity-gain feedback configuration to buffer the photodetector.

The voltage controlled resistors are implemented by parallel NMOS and PMOS transistors whose gates are driven by control voltages with symmetrical values above and below ground as shown in Figure 6.

Each pixel has a local resistor connecting its buffered photodetector to its output node as shown in Figure 5. Also, there are two lateral resistors connecting the output node to the output nodes of the adjacent pixels to the right and below. Finally, the output node is connected to lateral resistors in the pixels to the left and above.

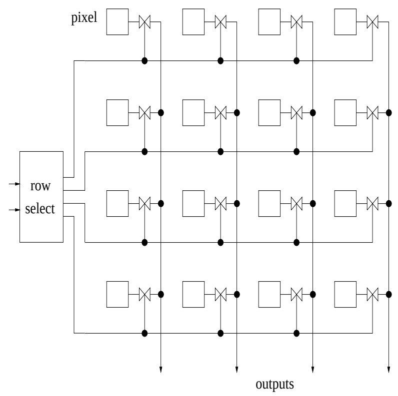

Figure 7: Row select logic for output multiplexing (only lower left 4x4 corner is shown here).

All local resistors are controlled by a single local control voltage, and all lateral resistors are controlled by a lateral voltage. The output node is buffered by an opamp that is essentially identical to the photodetector buffer.

The chip contains an 8x8 array of pixels that are connected as described above. At each edge of the array, all eight nodes are connected, through lateral resistors, to an externally accessible bus. Sixteen extra lateral resistors were added for the top and left edges. This arrangement allows great flexibility in determining boundary conditions for the convolution.

Each pixel has an area of about 27,000 square microns. The array covers about 1.7 square millimeters.

Since it is not possible to give all 64 pixels a dedicated output pin, the outputs are multiplexed off the chip a row at a time. A 3-to-8 decoder is used to select a single pixel row to be gated onto the eight output lines. The decoder is implemented in static CMOS logic as is shown for one corner of the chip in Figure 7.

The second metal layer of the 2 μm CMOS process is used as a light shield over all circuitry except the vertical PNP. This layer also serves as the conductor for the photodetector reference voltage.

6 Conclusion

We have explored three aspects of applying analog VLSI to real-time, interactive image processing, especially for advanced medical imaging applications. Simulation, PCB realization, and MAGIC layout indicate that analog circuits can perform advanced image processing interactively and in real-time.

Acknowledgements

This research was supported in part by The Piedmont Triad Engineering Research Center at Winston-Salem State University, Winston-Salem, NC, and by The Center for Advanced Computation and Communication at North Carolina State University, Raleigh, NC

References

[1] Griff L. Bilbro and Wesley E. Snyder. Applying mean field annealing to image noise removal. *Journal of Neural Network Computing*, Fall:5–17, 1990.

[2] Griff L. Bilbro, Wesley E. Snyder, Stephen J. Garnier, and James W. Gault. Mean field annealing: A formalism for constructing GNC-like algorithms. *IEEE Transactions on Neural Networks*, 3(1), 1992.

[3] Griff L. Bilbro, Wesley E. Snyder, and Reinhold C. Mann. The mean field minimizes relative entropy. *Journal of the Optical Society of America A*, 8(2):290–294, February 1991.

[4] Stephen J. Garnier, Griff L. Bilbro, James W. Gault, and Wesley E. Snyder. The effects of various basis image priors on mr image map restoration. *Journal of Mathematical Imaging and Vision*, 5(1), 1995. 1 of 2 articles in issue.

[5] Stephen J. Garnier, Griff L. Bilbro, James W. Gault, and Wesley E. Snyder. Magnetic resonance image restoration. *Journal of Mathematical Imaging and Vision*, 5(1), 1995. 1 of 2 articles in issue.

[6] D. Geman and S. Geman. Stochastic relaxation, Gibbs Distributions, and the Bayesian restoration of images. *IEEE Transactions on PAMI*, PAMI-6(6):721–741, November 1984.

[7] H. P. Hiriyannaiah, G. L. Bilbro, W. E. Snyder, and R. C. Mann. Restoration of piecewise constant images via mean field annealing. *Journal of the Optical Society of America A*, 6(12):1901–1912, December 1989.

[8] Anil K. Jain. *Fundamentals of Digital Image Processing*. Prentice Hall, 1989.

[9] Carver Mead. *Analog VLSI and Neural Systems*. Addison-Wesley, 1989.