Smart Pixel Technology for Digital Image Halftoning

Lieutenant Colonel Barry L. Shoop

Abstract

Smart pixel technology is a relatively new field which integrates electronic circuitry with optoelectronic devices in an effort to improve overall performance of existing applications and enable new applications not previously envisioned with either single technology. Smart pixels leverage the functionality and programmability of established electronic circuit technology and the high-speed switching and parallelism of arrays of optoelectronic devices. A smart pixel hardware implementation of the error diffusion neural network would provide the capability of real-time image halftoning for applications such as xerography, laser printing, and facsimile. The highlights from three smart pixel projects are described including the design, simulation, and experimental characterization of each.

Introduction

Smart pixel technology is a relatively new approach to integrating electronic circuitry and optoelectronic devices in a common framework. The purpose is to leverage the advantages of each individual technology and provide improved performance for specific applications. Here, the electronic circuitry provides complex functionality and programmability while the optoelectronic devices provide high-speed switching and compatibility with existing optical media. Arrays of these smart pixels leverage the parallelism of optics for interconnections as well as computation.

There are a number of different approaches to fabricating smart pixels which generally differ in the way in which the electronic and optical devices are integrated. Monolithic integration, direct epitaxy, and hybrid integration are the three most common approaches in use today. Monolithic integration allows both the electronics and the optical devices to be integrated in a common semiconductor material in a single growth process or by utilizing a re-growth process. The material of choice here would be a compound semiconductor material such as GaAs, InGaAs, or InP. Potentially, this approach would produce faster smart-pixels; however, there exist difficulties in the simultaneous optimization of both the electronic and optical circuits on the same substrate. Direct epitaxy of compound semiconductors onto silicon is another approach. The problems here are typically a result of lattice mismatch and differences between the coefficients of thermal expansion between the two materials. For example, the lattice mismatch between silicon and GaAs is 4% while the thermal expansion coefficients differ by 50%, resulting in a potentially unstable bond between the two dissimilar materials. Catastrophic failure of these devices typically occurs at this interface as a result of stress, strain, and sheer of the crystalline material. Hybrid processing is the third approach to developing smart-pixels. Here the optical devices are grown separately from the silicon electronic circuitry. Then in a subsequent processing step, the optical devices are bonded to the silicon circuitry using a variety of bonding techniques. These include flip-chip bonding, epitaxial lift-off and subsequent Van der Waals (contact) bonding, and creating a physical cavity above the silicon circuitry and flowing-in optical material such as liquid crystal material. Using this third approach, both the optical devices and the electronic circuits can be independently optimized resulting in an overall optimization of the smart-pixel.

Current Smart-Pixel Research

Over the past two years, we have become actively involved in the smart pixel technology field, designing, fabricating, and experimentally characterizing several different smart-pixel architectures focused principally on hardware implementations of the error diffusion neural network. In this report, we highlight a few of our research accomplishments in this program. Complete details of the design, simulations, and experimental results can be found in the publications listed at the conclusion of this report.

Here, we report on three specific smart pixel architectures: a 3 × 3 monolithically integrated neural array project called OPTOCHIP, two-generations of a 5 × 5 CMOS-SEED implementation of the error diffusion neural network, and a 7 × 7 diffractive optical filter which optically implements the 7 × 7 error diffusion filter of the error diffusion neural network.

3 ´ 3 OPTOCHIP Smart Pixel Array

OPTOCHIP is one of several different smart pixel realizations currently being investigated. OPTOCHIP is a monolithically integrated smart pixel which integrates GaAs/InGaP LEDs with GaAs electronics using a technique called epitaxy-on-electronics (E-O-E). The architecture for this smart pixel design was a 3×3 neural array which was designed to characterize the three functional components necessary for the implementation of an optoelectronic neuron which ultimately will be employed in the smart pixel realization of the error diffusion neural network. The three components are: 1) detection and amplification of the optical input signal, 2) a specific nonlinear sigmoidal function for thresholding, and 3) optical transmission of the resulting output signal. The solid state circuitry, which includes OPFET detectors and amplification circuitry, DCFL circuitry to produce the sigmoidal function, and driver circuitry for light emitting diodes (LEDs), was manufactured in a 0.6 μm Vitesse HGaAs III process available through the MOSIS foundry service. 50-μm × 50-μm GaAs/InGaP LEDs were then integrated into the electronics circuitry at MIT using the E-O-E post-process regrowth technique. Researchers from West Point participated as one of eight design groups sponsored by the National Center for Integrated Photonic Technology (NCIPT) under DARPA sponsorship and the direction of Professor Clif Fonstad at MIT.

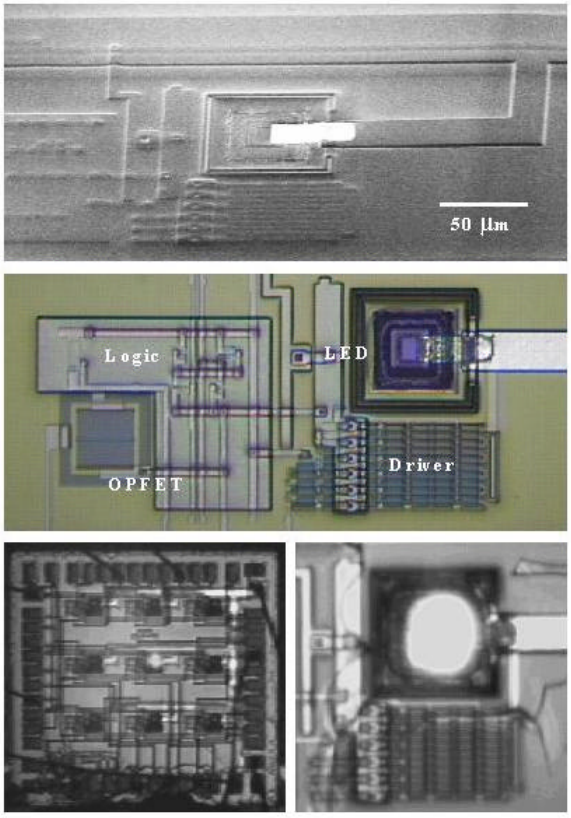

Figure 1. OPTOCHIP – SEM of regrown LED; photomicrograph of a single neural circuit; 3×3 neural array; and a single operational LED.

Figure 1 shows several images of the USMA OPTOCHIP circuitry. The top image is an SEM of the regrown LED and some of the electronic circuitry. The driver circuitry is shown below the LED with the neural circuitry shown on the left. The second image is a photomicrograph of one neuron of the 3×3 neural array. The OPFET detector, the neural logic circuitry, driver circuitry , and LED are all identified. Finally, on the bottom are photomicrographs of the 3×3 array and a single operational LED.

Experimental results have been compiled on the performance of two of three smart pixel components: the threshold circuitry and the LEDs. For the threshold circuitry, the slope of the sigmoidal function was significantly less than predicted by simulation, 3.5 compared to 41.2. This is shown in the quantizer performance of Figure 2. For our neural network application, however, this will not preclude network operation but will simply slow its convergence time somewhat. The experimental data confirmed a slight variation in performance of the circuit based on its position on the chip but this is not expected to significantly effect the performance of the neural network.

Nine LEDs were tested on each of two chips. The uniformity in the L-I curves for these LEDs, shown in Figure 3, was relatively poor. All emitted light, however, only one-third emitted sufficient light to be useful for this specific application. Spectrally, the LEDs were consistent, having a mean center wavelength of 872.5 nm with σ = 1.4 nm and a mean full-width half-maximum (FWHM) of 31.0 nm with σ = 2.5 nm. The spectral data from one representative LED is shown in Figure 4. Our MIT collaborators have recently demonstrated greater uniformity and yield subsequent to this fabrication with larger LEDs and improvements in the manufacturing process have been identified to further increase uniformity and yield.

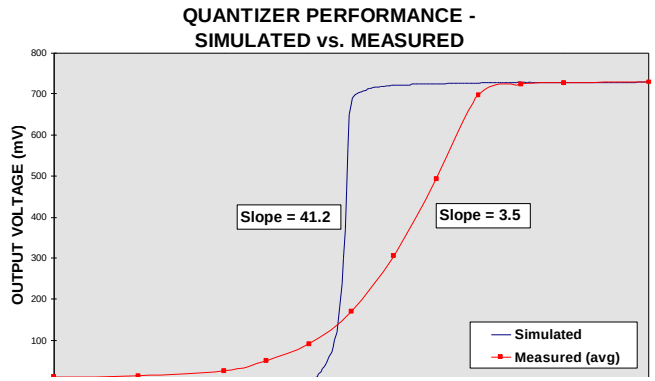

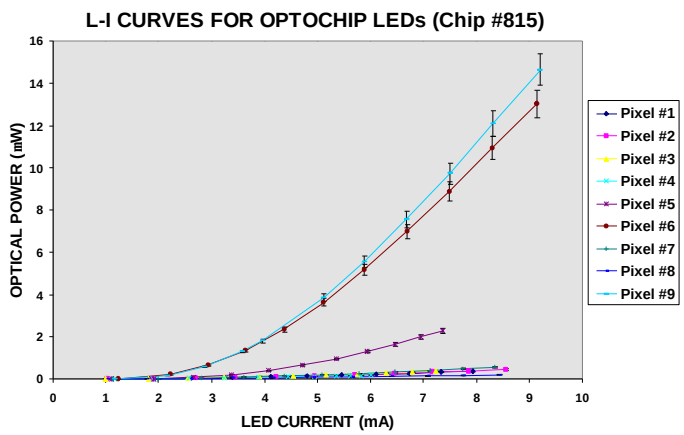

Figure 3. L-I curves for one representative OPTOCHIP.

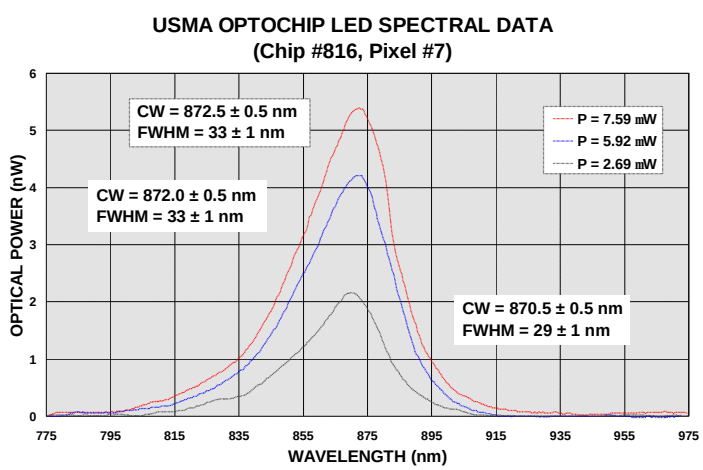

Figure 4. Spectral characteristics of a single OPTOCHIP LED.

Other experimental data remains to be collected on: the responsivity of the OPFET detectors, dynamic performance of each component, and a demonstration of a 1st-order oversampled modulator using all three components together. When the experimental results are concluded we anticipate providing a complete evaluation of the performance and applicability of this type of smart pixel technology to the error diffusion neural network architecture.

5´5 CMOS-SEED Smart Pixel Array

Here, we report on the results of two consecutive generations of SEED MQW modulators integrated with CMOS silicon circuitry using a flip-chip bonding hybrid integration approach. The first generation CMOS-SEED consisted of a 10 × 10 array of SEED MQW modulators integrated with 0.8 µm CMOS silicon circuitry while the second generation integrated an identical SEED array with 0.5 µm CMOS silicon circuitry. The focus of both of these designs was to demonstrate the usefulness of smart pixel technology to this specific analog neural application. Although optical weighting and interconnections are clearly the preferred method of achieving the necessary neural interconnections, neither of these designs incorporated this approach. Instead, electronic weighting and interconnections were incorporated in the design with analysis focused on optimizing the specific performance of the analog neural circuitry and the receiver and driver circuitry for the optical input and output signals, respectively. In a parallel effort within this research program, diffractive optical weighting and interconnections have recently been designed and experimentally characterized for a future smart pixel architecture. An overview of these results will be presented later in this report.

First-Generation CMOS-SEED

The functionality necessary to implement the error diffusion neural network consists of a one-bit quantizer, two differencing nodes, and the interconnection and weighting of the error diffusion filter. Figure 5 shows the first-generation circuitry for a single neuron of the error diffusion neural network using the CMOS-SEED smart pixel technology.

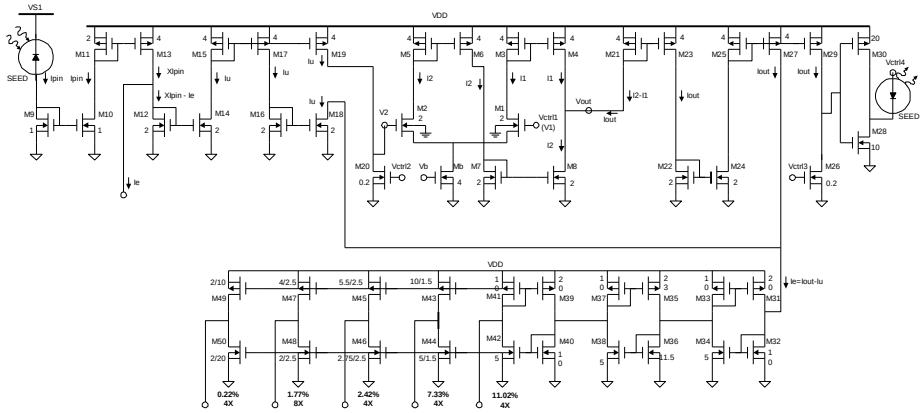

Figure 5. Circuit diagram for the first-generation smart pixel implementation of the error diffusion neural network based on a CMOS-SEED-type smart pixel architecture.

All state variables in each circuit are represented as currents. Beginning in the upper-left of the circuit and then proceeding in, the input optical signal incident on the SEED is continuous in intensity and represents the individual analog pixel intensity. The input SEED at each neuron converts the optical signal to a photocurrent and subsequently, current mirrors are used to buffer the input. The width-to-length ratio of the metal oxide semiconductor field effect transistors (MOSFETs) determine the current gain used to amplify the photocurrent. The first circuit produces two output signals: $+I_u$ which represents the state variable *u(m,n)* as the input to the quantizer and $-I_u$ which represents the state variable *-u(m,n)* as the input to the feedback differencing node. The function of the quantizer is to provide a smooth, continuous thresholding function for the neuron producing the output signal I_{out} which corresponds to the state variable *y(m,n)*. This second electronic circuit is a modified wide-range transconductance amplifier which produces a hyperbolic tangent sigmoidal function when operated in the sub-threshold regime. Additional transistors are included in the modification to provide input current to output current functionality. The third circuit takes as its input I_{out}, the state variable *y(m,n)*, produces a replica of the original signal, and drives an output optical SEED. In this case, the output optical signal is a binary quantity represented as the presence or absence of light. Here, the photoemission process is electroluminescence which results from the SEED being forward-biased. The last circuit at the bottom of the schematic implements the error weighting and distribution function of the error diffusion filter. The weighting is implemented by again, scaling the width-to-length ratio of the MOSFETs to achieve the desired weighting coefficients. The neuron-to-neuron interconnections are accomplished using the four metalization layers of the 0.8 μm silicon CMOS process. The difficulty encountered using this weighting and interconnect approach was that we were forced to constrain the size of the diffusion kernel to a 5×5 array because of the physical extent of the circuitry necessary to implement this functionality. Even with this reduced filter kernel, the weighting and interconnect circuitry consumed over 75% of the total silicon area. As a result, a smaller filter with a 5×5 region of support was designed and these coefficients were used in the smart pixel architecture. The five unique coefficients in this filter are shown at the bottom of the error diffusion circuitry along with the number of additional replications necessary to implement the 5×5 filter.

Figure 6 shows a photomicrograph of a single neuron of the first-generation error diffusion neural network. This smart pixel implementation resulted in a network with approximately 90 transistors per smart pixel and a total transistor count of nearly 1800. A total of 50 optical input / output channels are provided in this implementation. The rectangular features are the MQW modulators while the silicon circuits are visible between the modulators. The MQW modulators are approximately 70 μm × 30 μm and have optical windows which are 18 μm × 18 μm.

We have experimentally characterized some of the functionality of this first-generation smart pixel circuit.

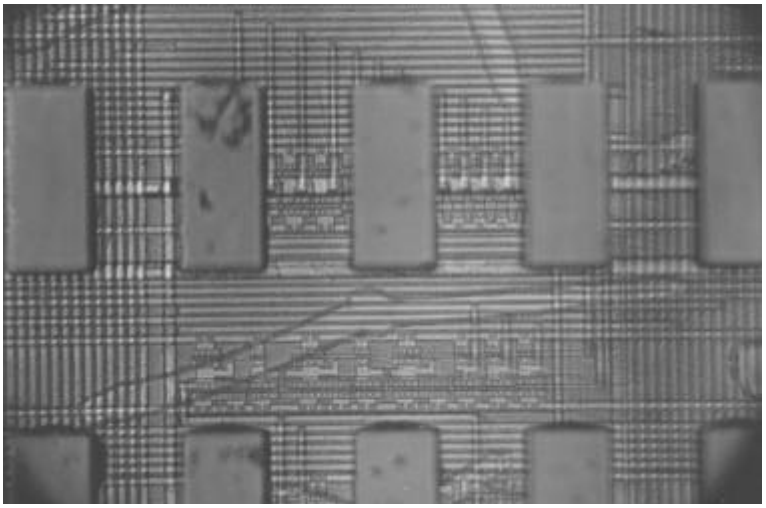

Figure 6. Photomicrograph of a single neuron of the 5 × 5 error diffusion neural network.

To date, experimental testing has confirmed electrical functionality of the quantizer operation as well as the replication features of the current mirrors within the architecture. We have also verified and characterized the electroluminescence of the SEED modulators. These experimental results demonstrated the importance of transistor matching within the wide-range transconductance amplifier and the error weighting circuitry. The results from this first-generation design and characterization were incorporated into the second-generation design in an effort to improve individual circuit operation and, as a result, improve overall network performance.

Second-Generation CMOS-SEED

In the design of the second-generation CMOS-SEED smart pixel, we concentrated on improving the performance of the neural circuitry. Specifically, we improved the error weighting circuitry and the quantizer performance. We also attempted to accurately model the operation of each of the functional elements as well as the total nonlinear network performance using SPICE simulation tools. The same 5 × 5 array size and functionality of one-bit quantization, subtraction, neuron-to-neuron weighting and interconnection, and optical input and output were retained. The electronic circuitry for the error diffusion neural network was implemented this time in 0.5 μm silicon CMOS technology. Figure 7 shows the circuitry associated with a single neuron of this second-generation error diffusion neural network. The one-bit quantizer is again implemented using a modified wide-range transconductance amplifier operated in the sub-threshold regime. The slope of this sigmoidal function was carefully designed by matching MOSFET transistors to meet the convergence criteria and the nonlinear dynamics of the error diffusion neural network. The error weighting circuitry at the bottom represents only the largest weight (11.02%) with interconnects to its four local neighbors (I_{outA} - I_{outD}).

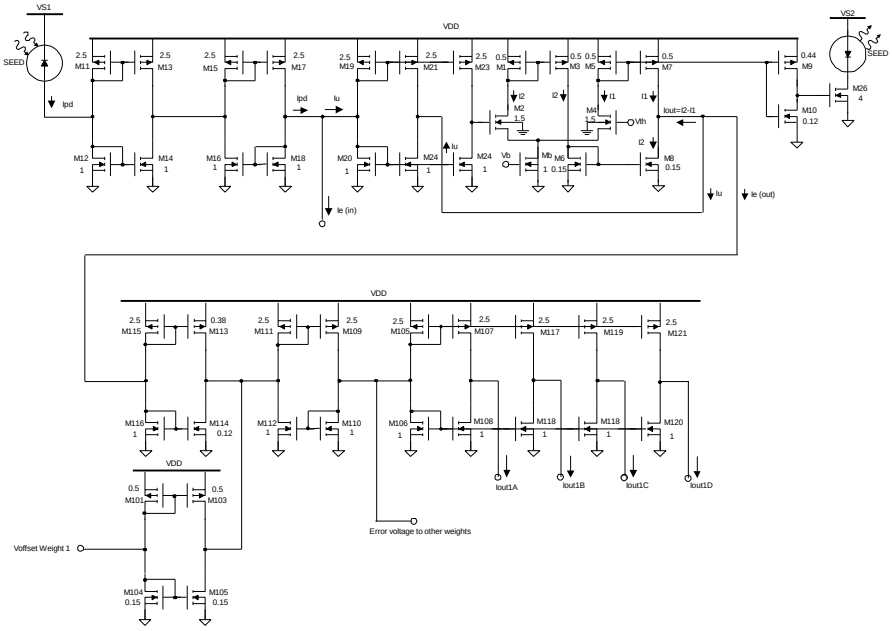

Figure 7. Circuit diagram of a single neuron and a single error weight of the 5 × 5 error diffusion neural network based on a CMOS-SEED-type smart pixel architecture.

Again, all state variables are represented as currents. The error weighting and distribution circuitry for the 5×5 array was again implemented in silicon circuitry, matching individual and stage-to-stage MOSFET transistors. Bi-directional error currents were implemented to provide the circuitry and the network with fully-symmetric performance. The specific improvements to the second-generation circuitry include improving stage-to-stage isolation by tying the drain and source of selected transistors together, as seen with transistors M19 and M20 in the schematic. This configuration places a constraint on the common node which helped isolate adjoining stages. This was particularly important to accurate circuit operation since the design of the wide-range transconductance amplifier required small feature size transistors compared to the input buffer and error weighting circuitry. In the circuit layout, we also changed several of the transistors in the wide-range transconductance amplifier to ensure that small width-to-length ratios were used in an effort to minimize channel length modulation effects. The SEED driver circuit was redesigned to provide additional isolation between the large currents used to forward bias the SEED and the balance of the sensitive neural circuitry. Finally, we accurately matched transistor characteristics throughout the network using extracted transistor data from a previous 0.5 μm foundry run.

Figure 8(a) shows SPICE simulations of the input-output relationship of the wide-range transconductance amplifier as well as the error signal resulting from the difference between the input and output of the quantizer. Figure 8(b) shows the linearity of each of the five unique interconnect weights over the entire operating range of the circuit. The quantization functionality resulted in a signum function with a slope of 14 while the linearity of each of the error weights was within 2% of the design specifications. Optical input and output to each neuron was implemented using the SEED MQW modulators as in the first-generation design.

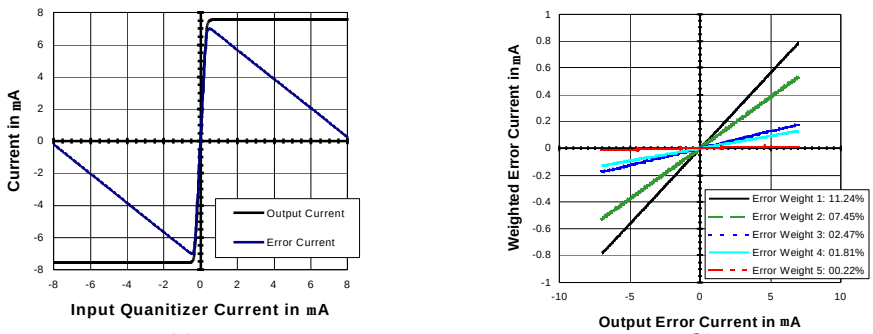

(a) (b)

Figure 8. (a) Quantization and error relationship for the wide-range transconductance amplifier, and (b) Modeled linearity performance of the error weighting circuitry.

The central neuron of this new smart pixel array consists of approximately 160 transistors while the complete 5×5 array accounts for over 3600 transistors, a nearly two-fold increase in the total transistor count over our first-generation design. The central neuron is interconnected to the surrounding 24 neurons in the 5×5 array using the same fixed interconnect and weighting scheme as in the first-generation realization. SPICE simulations of the complete 5×5 nonlinear dynamical neural network were performed using transistor parameters extracted from a previous 0.5 μm MOSIS foundry. Figure 9 shows the SPICE simulation for the performance of the complete 5×5 array to a linear input signal. As the input signal is linearly increased over the input dynamic range, the number of neurons in the

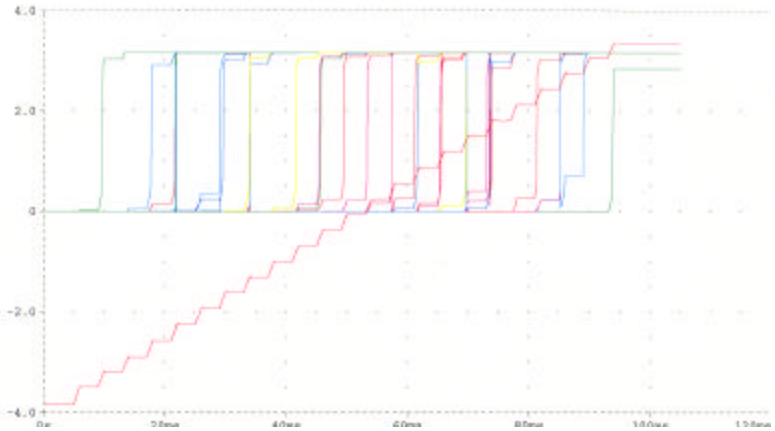

Figure 9. Dynamic performance characterization of the 5×5 error diffusion neural network to a stepped linear input.

on-state increases from 0 to 25, representing the analog input level at any given time interval as the ratio of the number of neurons in the *on*-state to those in the *off*-state. The non-uniform rate of increase in the number of neurons in the *on*-state is predictable and is a result on the artificially small network size. Figure 10 shows an expanded segment of Figure 9 which provides an analysis of the transient behavior of the network. These simulations predict individual neuron switching speeds of less than 1μs demonstrating the capability for real-time digital image halftoning.

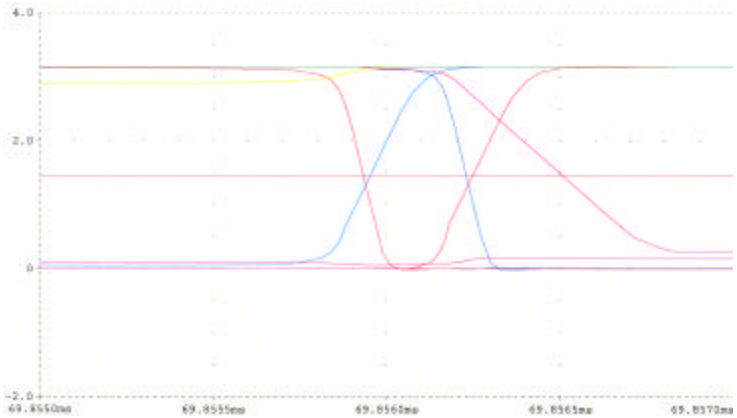

Figure 10. Transient analysis of 5 × 5 network. performance.

Both the SPICE simulations and the experimental results to date demonstrate that this approach to a smart pixel implementation of the error diffusion neural network provides sufficient accuracy for the digital halftoning application. The individual neuron switching speeds also demonstrate the capability for this smart pixel hardware implementation to provide real-time halftoning of video images.

7 ´ 7 Diffractive Optical Filter

As mentioned previously, optical weighting and interconnection is the preferred method for providing the error diffusion filtering functionality in the error diffusion neural network. In this specific project, a 7 x 7 non-uniform array generator was fabricated and experimentally characterized for this filtering function. The diffractive optical filter will be used in a future smart pixel architecture to spatially redistribute the quantization error generated in the error diffusion neural network.

The diffractive optical filter was designed using an iterative Fourier transform algorithm (IFTA) with the resulting spot distribution being formed in the back focal plane of a Fourier transforming lens. The objective of the design was to implement a discrete-level representation of a continuous-phase surface profile that generates a specific 2-D weighting and spatial distribution which is a non-uniform array with unequal spot spacing in the two spatial dimensions, shown in Figure 11. The surface profile is optimized to balance the error from the desired weights and the diffraction efficiency.

0.0003	0.0019	0.0051	0.0068	0.0051	0.0019	0.0003
0.0019	0.0103	0.0248	0.0328	0.0248	0.0103	0.0019
0.0051	0.0248	0.0583	0.0766	0.0583	0.0248	0.0051
0.0068	0.0328	0.0766	●	0.0766	0.0328	0.0068
0.0051	0.0248	0.0583	0.0766	0.0583	0.0248	0.0051
0.0019	0.0103	0.0248	0.0328	0.0248	0.0103	0.0019
0.0003	0.0019	0.0051	0.0068	0.0051	0.0019	0.0003

Figure 11. Two-dimensional weighting and spatial distribution of the diffractive optical filter.

The diffractive filter was designed to be incorporated into a CMOS-SEED smart pixel implementation of the error diffusion neural network which dictated the wavelength and spot spacing to be 850 nm and 160 μm x 80 μm, respectively. A 50 mm lens was selected as the Fourier transforming lens. The interconnect and weighting distribution shown in Figure 11 represents the impulse response of a 2-D error diffusion filter with a circularly symmetric frequency response and a cutoff frequency of 0.25π.

The IFTA was used to compute the optimal eight-level surface relief to produce the coefficient distribution in Figure 11. The number of features in the unit cell were 128 × 128 with the individual feature size being 2.1 μm × 4.2 μm. The unit cell was replicated 37 × 18 times to produce a filter approximately 1cm × 1 cm in size. The elements were fabricated by Honeywell Advanced Technology Center on 3 inch diameter fused silica substrates 0.5 mm thick with a refractive index of 1.4525. Ten separate 1 cm × 1 cm elements were delivered for characterization along with a detailed fabrication process error summary.

The etch depth errors were measured at several test locations on the 3 inch wafer substrates to an accuracy of 0.2 nm. For our specific samples, deviations from ideal etch depths for the three etches were: –1.7 nm for the π/4 etch (234.8 nm), –0.8 nm for the π/2 etch (469.6 nm), and –6.1 nm for the π etch (939.2 nm). These

measured errors were subsequently used as phase substitutions in the IFTA and the coefficient weighting was recomputed and normalized.

The measured spatial distribution of the diffractive optical filter can be compared with the desired distribution through an RMS error analysis. In general, the RMS error can be defined as

$$e_{rms} = \sqrt{\frac{\sum_{i=1}^{M}\sum_{j=1}^{N}\left(p(i,j)_{design} - p(i,j)_{measured}\right)^2}{M \cdot N}}, \tag{1}$$

where M × N represents the spatial extent of the diffusion filter. In our application, M=N=7. The RMS error between the design coefficients and the coefficients resulting from etch depth errors is 0.040%.

The mask alignment errors were measured using a two-dimensional vernier grid to a precision of 0.125 μm. The composite mask alignment error, defined as the root sum square of the two orthogonal measurement errors for mask 1-to-2 was 0.18 μm while for mask 1-to-3 was 0.15 μm. These alignment errors constitute approximately 8% of the feature size of the element and result in a reduction in the theoretical diffraction efficiency by a factor of 0.92.

For filter characterization, a tunable Ti: sapphire laser was coupled into a fiber and the output beam was then collimated. The diffractive filter was placed in the front focal plane of a 75 mm Fourier transforming lens. The back focal plane or Fourier plane was then imaged through an adjustable, rectangular aperture to provide individual access to each spot in the array. Using a beam splitter, the array was viewed on one axis with a CCD camera and the optical power was simultaneously measured on the other axis using an Anritsu ML9001A optical power meter with a MA9411A power sensor.

In the characterization, we considered three performance metrics: the accuracy of the spatial distribution, the diffraction efficiency, and the power uniformity across the diffractive filter. By collecting data on all ten samples, we were able to analyze process uniformity as well as determine an average spatial distribution across the ten samples. A 75 mm lens was used in place of the 50 mm Fourier transforming lens resulting in a simple scaling of the theoretical spot spacing from 160 μm x 80 μm to 240 μm x 120 μm according to the equation d=λf/W, where d, λ, f, and W are the spot spacing, wavelength, focal length of the transforming lens, and grating period, respectively. To measure the spot spacing, a knife edge controlled with a differential micrometer was inserted into the Fourier plane. By observing the image with the CCD camera, the spacing between each row in both the *i*- and *j*- dimensions was measured. The average spot spacing over this ensemble was 239.0 μm × 118.4 μm resulting in an error of 0.42% × 1.33% with standard deviations of 1.34 × 1.22, respectively. Assuming a diffraction-limited spot size and a 50 mm Fourier transforming lens, these spacing errors will not adversely effect the performance of the network.

To experimentally measure the power distribution, the rectangular aperture was used to isolate each spot in the array. The optical power at each location was measured with this process being repeated for each spot in the array for all 10 diffractive optical filters. The power measurements for each array were normalized and the power at each position averaged across all 10 samples. Figure 12 shows the average of the measured power distribution for each of the 49 coefficients in the array.

The RMS error was again computed, this time comparing the designed and measured coefficients in each array position. The RMS error for the measured distribution is 0.079% yielding extremely good correlation. The RMS error between the measured coefficients and the coefficients resulting from the etch

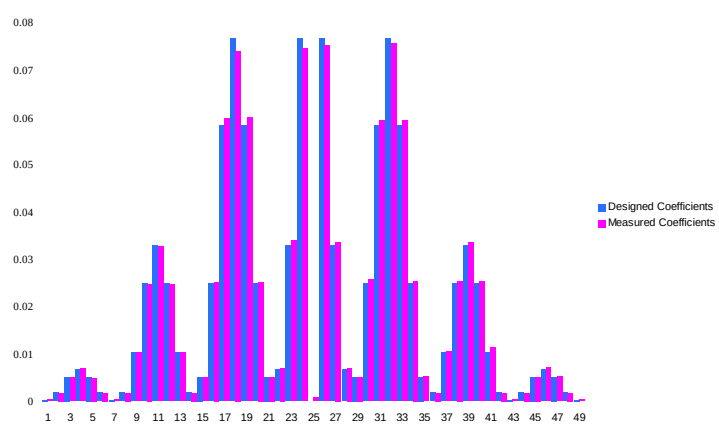

Figure 12. Comparison of designed and measured coefficients

depth analysis was 0.059%.

The average value of each of the nine unique, non-zero coefficients from the measured distribution was next compared with the design value for the corresponding coefficient. With the exception of the extreme corner positions and the central coefficient (zero diffraction order) in the array, the total error for each coefficient in the array was bounded by ±8% when compared to the design coefficients. When compared to the IFTA predicted coefficients, which accounted for fabrication errors, the total error drops to ±4%. In the extreme corner coefficient positions, we note considerable coefficient errors for both the measured distribution and for the distribution accounting for fabrication etch depth errors. These errors are 26.6% and 10.0 % respectively.

The ±4% error between the IFTA coefficients and the measured coefficients is well within instrumentation errors. The power meter accuracy is specified as ±5% at a power level of −23 dBm. Experimental power levels ranged from −18 dBm to −70 dBm at the extreme corner positions, accounting for the additional errors encountered in those positions.

Due to the requirement for zero-order suppression, we define the diffraction efficiency as the ratio of the total power in the 7×7 array to the total power incident from the source. For a continuous profile phase-only element, the upper bound on diffraction efficiency is 93.72%. An eight-level phase only element has an upper bound of 91.15%. The fabrication errors discussed earlier reduce this by a factor of 0.92 to 83.8%. The average measured diffraction efficiency for the ten filters was 83.7% with a standard deviation of 3.65 resulting in extremely good correlation to that predicted by theory.

The power uniformity across the surface of the filter is another important parameter that was measured by translating the filter in both the *i*-and *j*- dimensions while monitoring the power in one single array position. The array position selected had a power of −34 dBm and the maximum power fluctuation across the entire element was ±4%, again well within the instrumentation error of the power meter. This measurement demonstrates the feasibility of using a diffractive filter to process an entire image and all of its associated quantization errors simultaneously.

To determine how the filter would perform in an error diffusion network, we used two approaches: a computer simulation in which we compared halftone representations of a gray-scale image processed through the filter using several different sets of coefficients and a uniform gray-scale comparison. A gray-scale image was halftoned using both the ideal and the measured coefficients. No noticeable differences between the two images were found and no visual artifacts or distortions were introduced as a result of coefficient errors.

Additional simulations were conducted to determine the impact of the large corner coefficient error and the incomplete suppression of the central coefficient. The same image was processed with the corner measured coefficient error reduced to 7.5% and again with the central coefficient set to zero. In a detailed analysis of these images, no differences between the corrected images and the actual measured coefficient image could be found.

In the uniform gray-scale comparison, we considered two measures: average gray-scale accuracy of the output and anisotropy. A 25% gray image was first processed using both the design and measured coefficients. Ten images, each 256×256 were processed and yielded an output average gray-scale of 0.248772 with a standard deviation of 0.02 for the designed coefficients. For the measured coefficients, the average gray-scale was 0.247742 with a standard deviation of 0.04 resulting in a 0.4% error from the designed filter's performance. Anisotropy is a quantitative metric for judging the quality of a halftoning algorithm. Using the measured coefficients and fixed gray-scales of ½, ¼, 1/8, 1/16 and 1/32, the anisotropy curves were consistently −10 dB (lower bound) and indistinguishable from those produced by the design coefficients.

Conclusions

Through this research, we have gained significant experience with and insight into the advantages and disadvantages of the different types of smart pixel technology. We have developed the ability to design and fabricate analog smart pixel circuitry with reproducible precision and have characterized all of the necessary components of a smart pixel implementation of the error diffusion neural network.

Future Directions

Future efforts will continue to improve the analog neural circuitry performance and integrate the optical error weighting and distribution into a fully-functional architecture. Smart pixel technology incorporating vertical cavity surface emitting lasers (VCSELs) will also be investigated. Other image processing applications, beyond digital image halftoning, may also be investigated.

Publications and Conference Presentations

B. L. Shoop, A. H. Sayles, and D. M. Litynski, "New devices for optoelectronics: smart pixels," Chapter 20 in *Handbook of Fiber Optic Data Communications*, C. DeCusatis, D. Clement, E. Maass, and R. Lasky, Ed., Academic Press, San Diego, California (1997), pp. 705-758.

B. L. Shoop, R. W. Sadowski, Invited Paper in *Special Issue of Optical Engineering on Optoelectronic Data Communication, (to be published in December 1998)*.

B. L. Shoop, A. H. Sayles, D. A. Hall, and E. K. Ressler, "A smart pixel implementation of an error diffusion neural network for digital halftoning," Invited Paper in *Special Issue of the International Journal of Optoelectronics on Smart Pixels*, vol. 11, No. 3, pp. 217-228, 1997.

B. L. Shoop, "Image processing algorithms for smart pixel architectures," Invited Paper, in *OSA Annual Meeting Technical Digest*, 1997 (Long Beach, CA) October 1997.

D. A. Hall, B. L. Shoop, J. R. Loy, G. B. Tait, E. K. Ressler, J. F. Ahadian, and C. G. Fonstad, "Experimental demonstration of a 3x3 monolithically integrated smart pixel neural array based on epitaxy-on-electronics technology," Postdeadline Paper in *OSA Annual Meeting Technical Digest*, 1997 (Long Beach, CA) October 1997.

G. R. Kilby, B. L. Shoop, and J. N. Mait, "Performance of diffractive elements for use in optoelectronic analog-to-digital conversion," in *OSA Annual Meeting Technical Digest*, 1997 (Long Beach, CA) October 1997.

B. L. Shoop, A. H. Sayles, D. M. Litynski, and P. K. Das, "The role of smart pixel technology in future signal processing applications," Invited Paper, at *Gordon Research Conference on Optical Signal Processing and Holography*, Kimball Union Academy, NH, June 1997.

A. H. Sayles and B. L. Shoop, "Application of smart pixels to digital image halftoning," Invited Paper in *Thirtieth Annual Asilomar Conference on Signals, Systems, and Computers*, (Institute of Electrical and Electronics Engineers, Signal Processing Society), pp. 491-495, 1996.

J. N. Mait, D. W. Prather, and B. L. Shoop, "Design of diffractive elements for photonic processors," Invited Paper in *Thirtieth Annual Asilomar Conference on Signals, Systems, and Computers*, (Institute of Electrical and Electronics Engineers, Signal Processing Society), pp. 496-500, 1996.

Cadet Publications and Presentations

B. J. Smith and B. L. Shoop, "Wavelength dependent, multi-functional operation in a single multiple quantum well modulator," Proceedings of the Tenth National Conference on Undergraduate Research, (Asheville, North Carolina), pp. 1138 - 1142, 1996.

Research, (Asheville, North Carolina), April 1996.

J. H. Cuneo and B. L. Shoop, "Optically controlled nonlinearity for an error diffusion neural network," *Proceedings of the 1998 National Conference on Undergraduate Research*, (Salisbury, MD), April 1998.

S. A. Bradley and B. L. Shoop, "Dynamic optical interconnect weighting for a neural network," *Proceedings of the 1998 National Conference on Undergraduate Research*, (Salisbury, MD), April 1998.

Cadet Involvement

Brian J. Smith ('96) Jason H. Cuneo ('98) Scott A. Bradley ('98)

Collaboration

Department of Electrical Engineering and Computer Science

Major Dirk A. Hall	Major Gregory R. Kilby
Major Robert W. Sadowski	Colonel Andre H. Sayles
Lieutenant Colonel Glen P. Dudevoir	Colonel Daniel M. Litynski

Rensselaer Polytechnic Institute

Professor Pankaj K. Das

Army Research Laboratory, Adelphi, MD

Dr. Joseph Mait	Dr. George Simonis

Massachusetts Institute of Technology

Professor Clifton G. Fonstad, Jr.	Joseph F. Ahadian