SIMPil: An OE Integrated SIMD Architecture for Focal Plane Processing Applications

Huy H. Cat, Antonio Gentile, John C. Eble, Myunghee Lee, Olivier Vendier, Young Joong Joo, D. Scott Wills, Martin Brooke, Nan Marie Jokerst, April S. Brown

School of Electrical and Computer Engineering
Packaging Research Center
Georgia Institute of Technology
Atlanta, Georgia 30332-0250

Abstract

Focal plane processing applications present a growing computing need for portable telecomputing and videoputing systems. This paper demonstrates the integration of digital processing, analog interface circuitry, and thin film OE devices into a compact computing package. The SIMPil architecture provides a programmable, silicon efficient SIMD processor for effective execution of early image processing applications such as edge detection, convolution, and compression. Results from a demonstration SIMPil node are presented including its microarchitecture, and performance on image processing applications.

1. Introduction

Low cost video cameras and advanced telecommunications enable many new services, such as electronic video mail and computer-based teleconferencing. Evolving compression standards (e.g., MPEG) and inexpensive disk storage allow these electronic exchanges to be stored much like e-mail is stored today. However, acquiring, transmitting, and manipulating this information presents a computational requirement beyond the capabilities of existing systems. Increasing user demand for portable *on the move* videoputing and telecomputing place additional requirements on size, weight, and power.

This paper presents a demonstration vehicle being developed at Georgia Tech to enable a new class of portable OE integrated parallel processing systems designed to efficiently process high-throughput image streams. This demonstration includes a focal plane SIMD architecture, *SIMPil*, that is currently being developed for use in videoputing systems such as high speed smart cameras. This prototype addresses issues in multidisciplinary interfacing by incorporating an integrated thin film detector, on-chip analog interface circuitry, and a powerful digital processor on a single Si CMOS chip. To illustrate the effectiveness of the SIMPil processing architecture, several image processing operations are demonstrated including edge detection, convolution, and image compression. The silicon area efficiency of this type of processing node is compared with general purpose commercial microprocessors.

This paper begins with a summary of related work, followed by a description of the SIMPil architecture, including details of the demonstration chip. Then OE integration and analog interface circuit results are presented. Finally SIMPil applications and future work are presented.

2. Related Work

Over the past decade, techniques for integrating OE devices, analog interface circuitry, and digital logic have enabled new approaches for image collection and processing. Monolithic systems incorporating focal plane arrays offer high I/O bandwidth with modest levels of dedicated analog or digital processing capability. Beginning with Mahowald and Mead's silicon retina [21], on-focal plane processing has increased in complexity from simple logic gates to latches [17] to 2-bit registers and counters in [13]. Analog processing alternatives have demonstrated even greater operational complexity using passive and active networks. The systems strive to achieve high fill factor detector arrays combined with the maximum computing capability that can effectively be incorporated nearby. These image processing solutions are compact and efficient, but lack computing power and flexibility.

An alternate system approach focuses on computation incorporating more powerful digital processors. The CM-2 [23] and MasPar [3] are examples of general purpose SIMD systems capable of performing image processing

0-8186-7591-8/96 $5.00 © 1996 IEEE
Proceedings of MPPOI '96

applications. However, these systems achieve performance and generality at the expense of focal plane I/O coupling and physical size. Other systems, including the Scan Line Array Processor (SLAP) [8], exploit frame scanning used in video cameras by operating on sequential scan lines. However, this serial loading and unloading of image data limits frame rates. The Morphological Image Processor (MIP) [7] combines dedicated processors with an on-chip focal plane array. However, its functionality is limited to morphological operations on binary images.

The SIMPil system presented in this paper combines features from both focal plane systems and image processing architectures. An integrated optoelectronic detector allows through-silicon wafer input of digital image data from a detector plane stacked above the processing plane, shown in Figure 1. By reducing the image transfer bottleneck found in decoupled detector-processor systems, high frame rates are possible without constraining processing power. Processing area does not impact detector array fill factor.

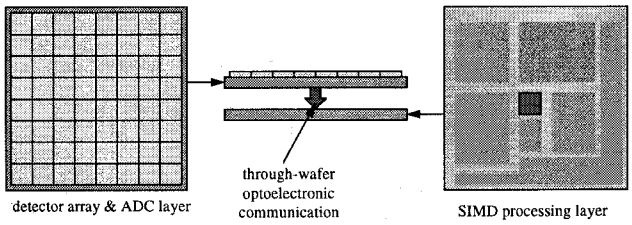

Figure 1: A Stacked Two Layer Focal Plane Processor.

3. The SIMD Pixel Processor (SIMPil)

Parallel architectures are well suited for image processing applications because of their ability to exploit inherent data parallelism in images. Parallel architectures can be structured in two ways. They can operate on an entire image or subimage in parallel (e.g., MPP [2] and CLIP4 [6]), or it operates on the image data in a pipelined model (e.g., Warp [12] and Geometric Arithmetic Parallel Processor (GAPP) [26]). The parallel organization has been expensive, often requiring significant resources. This has forced architectures to restrict node performance (e.g., by using bit-serial processors). Pipelined architectures reduce this problem by using a smaller number of higher performance processors. However, they are difficult to scale, and the generality of the processor is often not required.

The SIMD Pixel Processor (SIMPil) architecture maps a subarray of pixels to each processor. A bit parallel datapath supports efficient manipulation of grayscale pixel values from the subarray, while still offering large scale parallelism. Vertical coupling to the image plane allows the detector and processor arrays to be scaled while maintaining a fixed level of processing per pixel.

3.1. Processor Architecture

The block diagram of a single SIMPil node is displayed in Figure 2. The figure also illustrates how a single node interfaces to a subarray of detectors through sample and hold circuitry (S&H) and analog-to digital converters (ADC), and how each node is connected to each other in a mesh network to operate in SIMD mode. Each node includes a traditional RISC load/store datapath plus an interface to the detector array via an OE data channel. Initially, an 8-bit datapath SIMPil node was implemented. It includes an 8-word register file, an arithmetic logic unit, a shift unit, a 16-bit multiply-accumulator (MACC), and 64-word local memory. The implemented applications have executed satisfactorily with 64 words of local memory. However, additional applications may demand an increase in memory size (128 to 256 words). Keeping the silicon area devoted to memory in balance with area devoted to processing is a design goal of this architecture in order to maximize the efficiency of the system implementation.

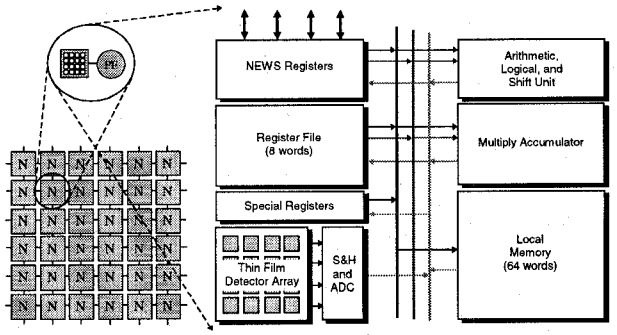

Figure 2: SIMPil Microarchitecture

The instruction set architecture (ISA) provides for arithmetic operations including addition, subtraction, multiplication, and multiply accumulation. The multiply accumulate (MACC) instruction is included because of its utility in image processing applications. For example, the MACC operation reduces the partial convolution of a 3×3 sub-image from 9 multiplies and 8 additions to only 9 multiply accumulate operations. The 16-bit accumulator in an 8-bit datapath improves precision especially when using fixed-point operands. The logic unit allows bitwise AND, OR, and exclusive-OR operations. Logical, arithmetic, and rotate shifts operations are performed in the shift unit. Register-to-register and immediate addressing modes are supported by the dyadic operations. Local memory is accessed via the load and store instructions.

Each SIMPil node interfaces to an array of thin film detectors. The instruction set architecture (ISA) allows for

up to 256 addressable detectors. Each node also includes analog to digital circuitry to convert light intensities to digitally equivalent values. The ISA has a SAMPLE instruction that synchronously captures light intensities at each detector. The SIMD execution model allows the entire image to be sampled by the system synchronously. Once the detector array has been digitized, it can be processed by the SIMPil node in data parallel fashion.

3.2. Communication

Low level image processing applications such as edge detection are usually point algorithms needing only pixel values in a small neighborhood around the data point. This pixel access locality is well supported by a nearest neighbor or mesh network. SIMPil nodes communicate through a nearest neighbor NEWS (north, east, west, and south) network using NEWS registers in the datapath. The NEWS registers are addressable as general purpose registers making the register file a total of 12 words. To communicate to a node's neighbor, an appropriate NEWS register is specified as a destination in any instruction. Reading a NEWS register accesses the locally stored value for that communicated operand.

3.3. A Typical System

The SIMPil system is an embedded, programmable, focal-plane image processing system. The processing power of the SIMPil node will surpass the computational needs of a single pixel. However, desired frame rates may not be achieved if the number of pixels assigned to a node is too large. Simulations of image processing applications suggest a good balance of 36 to 64 pixels per SIMPil node (with 50 MHz node frequencies). Our prototype target is 64 pixels per SIMPil node.

Using current VLSI technology, between 16 and 64 SIMPil nodes can be fabricated on a single Si VLSI chip. By tiling an array of 16 chips each containing 16 nodes, a 128×128 pixel resolution is achieved. The aggregate total for this system is 16,384 pixels and 256 SIMPil nodes. Operating at 50 MHz, SIMPil can perform 781 Kops/sec for each pixel. Eight bits is the minimum datapath width for pixels supporting 256 gray scale levels.

This system can be scaled in two ways to meet different needs. More chips can be tiled to achieve higher resolution for operation on an entire image at the same processing rate. For example, with 64 SIMPil nodes on a single VLSI chip, an 8×8 array of chips fits on a 4" square MCM, yielding a 4096 SIMPil nodes system with a resolution of 512×512 pixels. Scanned arrays can delivery higher resolutions at slower frame rates.

3.4. Digital Test Results

The first prototype of the SIMPil system consists of a single SIMPil node processing data received from an optical data image stream. This prototype includes the thin film detector integration site, analog receiver amplifier, and a prototype SIMPil node. The communication (NEWS) registers are not included in this implementation due to pad limitations. The prototype, shown in Figure 3, was fabricated in Hewlett Packard's 0.8μm CMOS n-well, 3-metal process through the MOSIS fabrication service. The single-ended receiver and device integration site are located on the right side of the chip.

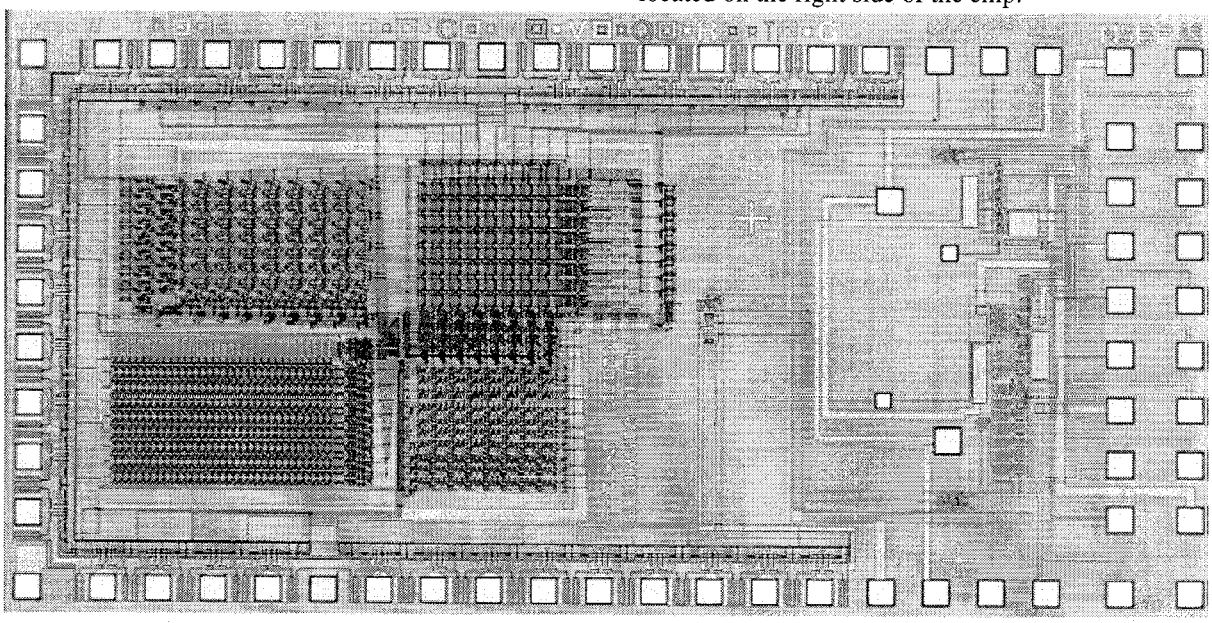

Figure 3: Photomicrograph of a SIMPil Node and Receiver Amplifier

The prototype was digitally tested using a Tektronix LV500 chip tester. The test setup is shown in Figure 4a. The prototype sits on the circular printed circuit board. The LV500 provides stimulus to 27 digital control signals for the SIMPil node, emulating the instruction unit by executing SIMPil programs. It is also used to read and write to the 8-bit data bus. Functional tests were performed on each functional unit using test programs simulated on the simulator. The register file can be loaded with values from the off-chip bus. The ALU, barrel shifter, and multiply-accumulate unit were stimulated with operands from the register file, and its results were written back to the register file. The results are then read from the register file to the off-chip bus for evaluation. The memory was tested using a similar technique. Every word in the memory was alternately written with zeros and ones and then read back. Figure 4b illustrates a code fragment that is currently being exercised. The fragment is part of a convolution program.

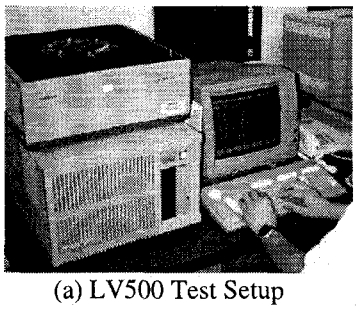

(a) LV500 Test Setup

```
loadi r0, 12
load r1, r0
loadi r0, 24
load r2, r0
mul rw, r1, r2
macc r3, re, r7
;; c11 = (36) = r3
loadi r0, 36
add r3, r3, r6
store r0, r3
```

(b) Convolution Code Fragment

Figure 4: SIMPil Digital Testing

3.5. Area Efficiency Comparison

One system evaluation metric is the efficient use of available silicon area for processing. Specifically, we look at area efficiency by considering the number of operations executed per second per area (Mop/s·mm^2). As a benchmark for comparison we have evaluated this metric for three high-performance processors -- the Alpha 21164a, Pentium P54, and PowerPC 603ev (Table 1). In these systems, we consider fixed point (integer) performance only.

This comparison is of course not totally equivalent because these architectures devote a large area to floating point units and caches and have wider word widths. However, these are attributes of such general purpose processors. SIMPil is a small, special-purpose processor targeted for image processing applications. Its data path is only 8 bits for moderate intensity resolution. Therefore, we expect it to be very efficient for the set of applications it targets. All of these commercial processors are fabricated in state-of-the-art 0.35 μm processes. Using generalized scaling theory [1], we have projected the performance and area of a SIMPil node fabricated with 0.35 μm and 3.3 Volt technology. Power efficiency comparisons are reported in [5].

Table 1: Area Efficiency Comparison

Processor	Technology (μm)	Vdd (V)	Area (mm^2)	Clock Freq. (MHz)	Mop/s·mm^2
Alpha 21164a[1]	0.35	3.3	209	417	**2.4**
Pentium P54[1]	0.35	3.3	90	150	**2.0**
PowerPC 603ev[1]	0.35	2.5	81	166	**2.0**
SIMPil	0.8	5.0	2	50	**16.8**
SIMPil (scaled)	0.35	3.3	.38	176	**310.0**

These results assume an effective throughput of 0.67 instructions per cycle for SIMPil. These numbers express the efficiency of the SIMPil architecture for the targeted set of applications.

4. Thin Film Detector Integration and Design Issues

Three critical factors in the receiver design are input capacitance, resistance, and detector dark current, which determine the speed and sensitivity of the receiver. The operating speed is set in this circuit by the RC time constant. Since the detector is integrated directly onto the Si circuit, the input resistance is not set by the package interface (usually 50 Ω), but rather, can be treated as a design variable, and should be maximized to minimize the noise in the receiver. The input capacitance, set by the detector and detector pads, constitutes a trade-off between a minimum value (to maximize the input resistance) and a maximum value (to maximize the detector size for alignment tolerance). A selection of detector which has low capacitance per unit area, yet high responsivity and low dark current is a choice of vital importance in the design of the photoreceiver. A MSM meets all of these requirements save the responsivity, (which is poor, generally 0.2 to 0.4 A/W) and the dark current, (which is relatively large in InP material systems [15]). The inverted metal-semiconductor-metal-semiconductor photodetector (I-MSM), with the electrodes defined on the bottom of the device, eliminates the shadowing effect of the electrodes, producing high responsivity. An InAlAs cladding layer minimizes the dark current [22] thus yielding a, low capacitance, large area detector for low noise, high speed, alignment tolerant integrated receiver operation.

[1] Information taken from [4]

The Si CMOS circuit and the I-MSM were independently optimized and fabricated, and the hybrid integration was performed in post-processing steps. The I-MSM structure consisted of an InP (substrate) / InGaAs (100 nm, etch layer) / InAlAs (40 nm) / InGaAs (1000 nm) / InAlAs (40 nm), with all layers nominally undoped [18]. The cladding layer also eliminates low frequency gain [16], and enhances the Schottky barrier. These structures have demonstrated up to 0.7 A/W responsivity at 1.3 μm wavelength [24], and exhibit, for 250 μm square devices, less than 200 nA of leakage current at a 10 V bias. The measured capacitance of 150 fF is typical for a 250 mm active area I-MSMs with 2 μm finger width and 8 μm spacing.

To fabricate the I-MSMs, 250 μm mesas were etched, Ti/Pt/Au fingers were defined, and, prior to substrate removal, the photodetectors were tested on-wafer so as to integrate known good detectors. The I-MSM mesas were covered with an Apiezon W wax handling layer, and the InP substrate and etch stop layers were subsequently removed using selective etches [24]. The thin film I-MSMs were then bonded to a transparent mylar transfer diaphragm, and the handling layer was removed with TCE. The I-MSM fingers and the short contact strips on each side were facing away from the diaphragm, for bonding to the pads on the Si CMOS amplifier circuit.

Figure 5 shows the optoelectronic detector integrated onto the receiver amplifier. This chip is currently being tested.

5. Interface Circuitry (Receiver Amplifier)

The receiver amplifier provides low noise amplification of the optical detector output current. For through wafer optical data links we expect signals as low as 10 nA will arrive at the receiver input. For low error rates the signal to noise ratio must be about ten to one meaning the receiver's input referred current noise should be approximately 1 nA. The receiver must amplify the signal until a fast CMOS comparator can generate a digital output from the signal. For CMOS this means we would like about 10 – 100 mV of output signal meaning the receiver should have a gain of 1 – 10 MΩ. In addition the receiver must have a differential front end to prevent pickup of digital noise from the processor located on the same chip. Simulations indicate that gains of 1 – 10 MΩ at 10 – 100 Mbps are possible with 0.6 μm CMOS.

5.1. Experimental Results

To date we have not designed a receiver that meets all of the design goals and runs at 100 Mbps. We have designed a receiver that runs at 155 Mbps and has a gain of 1 KΩ. The receiver included in the chip shown in Figure 5 has been separately fabricated and tested before integration with the SIMPil node. The receiver die photo and measured performance are presented in Figure 7 and

Figure 6. We are beginning the design of a higher gain receiver and also working on making the input signal specification less severe.

A digital 0.8 μm CMOS transimpedance amplifier was designed to meet SONET OC-3 SR specifications, with a target bandwidth of 155 Mb/s. The design optimization yielded a multi-stage, low-gain-per-stage design. Assuming a fixed gain/power dissipation product, an open-loop gain approach provided better upper −3dB frequency cut-off than a feedback circuit. The optimized circuit consisted of 5 identical stages, each with a current gain of

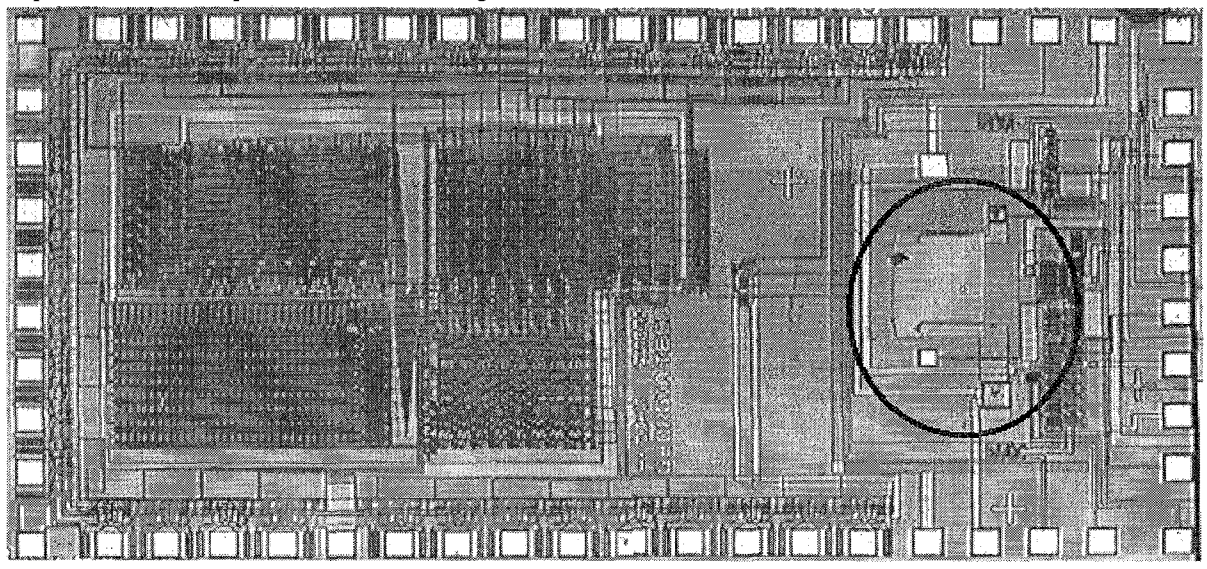

Figure 5: Photomicrograph of a SIMPil Node, Receiver Amplifier, and Integrated I-MSM Detector

3. The total transimpedance gain is about 12 KΩ with a 50 Ω load.

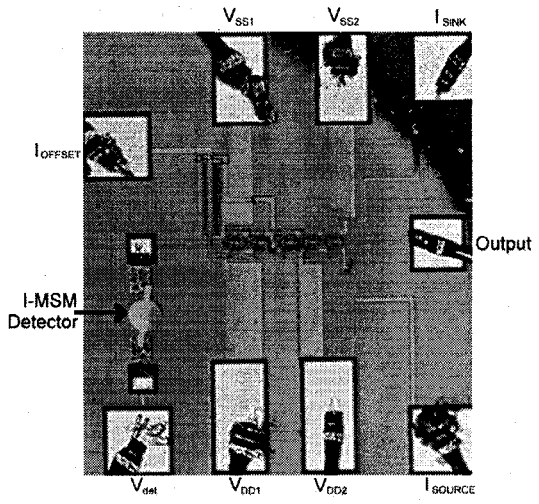

Figure 7: A 0.8 μm CMOS Receiver Fabricated in a Digital CMOS Process

This photoreceiver has been fully characterized with an integrated 50 μm square thin film InP- based I-MSM, and exhibited excellent characteristics, matching the SONET OC-3 SR specifications. In order to test the receiver, a commercial 1 mW CW 1300 nm laser diode was directly modulated by a current source and the pseudorandom data was fed through a bias tee to test the integrated circuit.

Figure 6 (left) shows a 155 Mb/s wide open eye diagram measured at −19 dBm. This circuit demonstrates operation at 200 Mb/s with 2^7-1 NRZ optical pseudorandom data at −18 dBm, as shown in Figure 6 (right).

6. Applications

The SIMPil architecture targets a high-throughput, low-memory class of applications. These problems requires short parallel tasks to be performed on massive array streams of data. An argument for the trade-off of less local memory and higher data throughput is made in [25]. Examples of this application class include image processing (e.g., filtering, edge detection, convolution), object recognition, and data compression. To evaluate the SIMPil architecture, several image processing algorithms were programmed and simulated using an instruction-level simulator. These algorithms are also being used to test SIMPil implementation prototypes.

6.1. Morphological Edge Detection

A common image processing application is edge detection. In one implementation, the morphological edge detection algorithm has been programmed on the SIMPil simulator to determine edges on binary images. Gray scale images can also be processed with minor modification to the algorithm. With morphological operations [20], edge detection can be accomplished in one of two ways: dilate the image and subtract the original image from it, or erode the image and subtract it from the original image. The former technique was used where dilation is defined as $(X \oplus B)(m, n) = MAX\{X(m + i, n + j) \mid \forall (i, j) \in B\}$. B(m,n) is the structuring element and X(m,n) is the set of pixel values. Of the different possible structuring element arrangements, the 3×3 neighborhood structuring element was used. With each SIMPil processor mapped to a 4×4 detector array, more communication to neighboring processors is required than with systems having a one

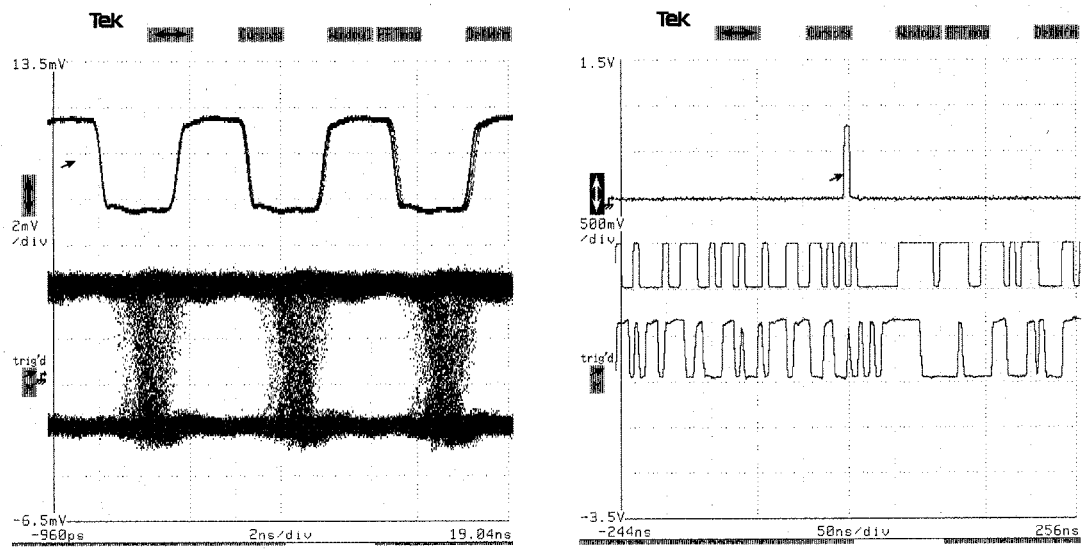

Figure 6: Measured Performance of Receiver at 155 Mbps

detector per processor organization.

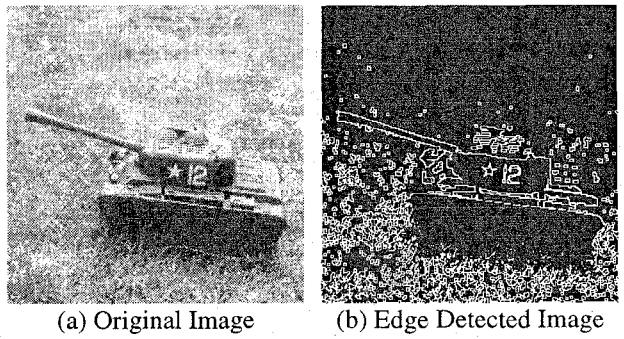

(a) Original Image (b) Edge Detected Image

Figure 8: Input and Output Images for Edge Detection

Figure 8a shows the input image (256×256 pixels) for the morphological edge detection algorithm. This image was first passed through a thresholding filter. Then, the morphological dilation and subtraction were performed to output the image in Figure 8b. The simulated system is 64×64 (4K) SIMPil nodes with each node having a 4×4 (16) array of detectors to provide a total resolution of 256×256 pixels. Figure 9 illustrates the average concurrency for the morphological edge detection algorithm running on the SIMPil simulator. The concurrency plot shows the three major sections of the algorithm indicated by the breaks in concurrency. The image is sampled and loaded to local memory in the first 100 cycles. Then, the following 300 cycles (with the little dips in Figure 9) represent the section where the morphological operations were performed in rows and columns. Finally, the rows and columns are combined to complete the morphological operation and image subtraction (shown by the 16 large dips in Figure 9). SIMD execution offers high processor utilization in these types of applications.

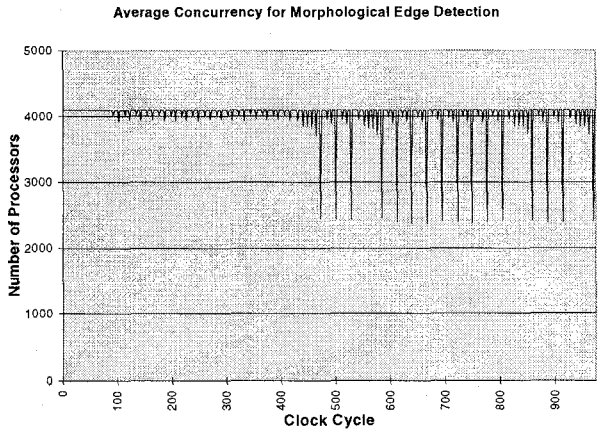

Figure 9: Concurrency Plot

6.2. Convolution

Another filtering application programmed for the SIMPil architecture is convolution. Convolution, defined in equation (1), is a basic image transform used to accomplish many image filtering operations.

$$X(u) = \sum_{k=0}^{N-1} x(k)w(u-k) \qquad (1)$$

In SIMPil's implementation, 3×3 Sobel masks are used as gradient operators for edge detection. From equation (1), the summation of the partial products is simplified by using the multiply-accumulate instruction rather than a multiply and then an addition. This reduces the 9 multiplies and 8 additions to only 9 multiply-accumulates in the general case.

Both morphological edge detection and convolution are elementary image processing algorithms. However, they help to illustrate examples of high-throughput low-memory algorithms. At 50 MHz, these filtering operations can easily be performed at 60 frames per second. The monolithic design and SIMD operation node allow this rate to be sustained at variable image size. The bandwidth bottleneck between detector array and parallel processors does not exist even as image size is increased.

6.3. Full-Search Vector Quantization

With efficient image compression algorithms, the communication bandwidth required for image transmission and the required memory capacity for storage is minimized. Moderate image sizes of 512×512 pixels per frame at a rate of 30 frames per second will require a transmission rate of almost 60 Mbits per second for an 8-bit per pixel (bpp) intensity. HDTV standards or greater data range (bpp) will quickly increase the transmission rate requirements. Even with satellite images where the frame rates are lower, the data explosion is evident and transmission rate is limited. Among compression algorithms, vector quantization (VQ) has become a popular technique [11][14][19]. We have programmed a parallel implementation of a full-search VQ encoding algorithm for the SIMPil architecture. Compression is achieved by subdividing the image into small blocks (e.g. 4×4 pixels) and finding the best match for each block in a provided codebook. Once the best fit is found, the index to the codebook is the compressed value for that block. SIMPil's implementation details are explained elsewhere [9][10]. Figure 10 shows the original image of Lena requiring 8 bits/pixel and its compressed image using vector quantization requiring only 0.5 bits/pixel. The lossy effects of vector quantization are acceptable for the achieved compression rate.


(a) Lena, 8 bits per pixel


(b) Lena, 0.5 bits per pixel

Figure 10: Original and Compressed Image from VQ

It is interesting to analyze performance comparisons between SIMPil and different hardware platforms. Table 2 compares a 256 node (16×16) SIMPil system with a 2048 node Maspar 1 and 32 T800 nodes. The performance measure uses 4×4 blocks as a unit to compensate for the larger codebook in the Maspar implementation. Although SIMPil's performance is simulated, the performance numbers suggest that a potential improvement exists. Vector quantization is a significant application demonstrating that high-throughput computation can be supported with low memory.

7. Conclusions and Future Work

We have presented the integration of a thin film detector to a programmable SIMD processor. The optoelectronic detector, receiver amplifier, and SIMPil processing node were tested successfully. An instruction level simulator is available to execute the applications presented. While these applications are representative of high-throughput, low-memory applications, we will continue to implement other applications for focal plane image processing.

Future work includes continuing to characterize the behavior of the prototype through further tests. In the next experiment, a modulated light source (controlled by the LV500) will be focused on the I-MSM detector for emulation of the detector plane image stream. Programs will then be executed on the prototype SIMPil node using image data input from the detector. A detailed examination of digital switching noise interaction with the analog circuitry will be made, and general performance measurements will be taken.

A new 16 bit version of the SIMPil system is currently being designed for fabrication in Fall 1996.

8. Acknowledgments

This work is being supported by AFOSR, NSF (ECS-9422452), NSF Packaging Research Center (EEC-9402723), and the DARPA Low Power Electronics Program.

9. References

[1] G. Baccarini *et al.*, "Generalized scaling theory," IEEE Transactions on Electron Devices, ED-31, Apr. 1984, pp. 452-462.

[2] K. E. Batcher, "Design of a Massively Parallel Processor," IEEE Transactions on Computers C-29, 9, Sept. 1980, pp. 836-840.

[3] T. Blank, "MasPar MP-1 Architecture," Proceedings of COMPCON Spring '90 - The Thirty-Fifth IEEE Computer Society International Conference, San Francisco, CA, 1990, pp. 20-24.

[4] CPU Info Center, http://infopad.eecs.berkely.edu/CIC/summary/index.html.

[5] H. H. Cat., J. C. Eble, D. S. Wills, V. K. De, M. Brooke, N. M. Jokerst, Low Power Opportunities for a SIMD VLSI Architecture Incorporating Integrated Optoelectronic Devices, *GOMAC'96 Digest of Papers*, pages 59-62, Orlando, FL, March 1996.

[6] M. Duff, "CLIP4: A Large Scale Integrated Circuit Array Parallel Processor," Proceedings Third IJCPR, Coronado, CA, 1976, pp. 728-733.

[7] W-C Fang, T. Shaw, J. Yu, "VLSI Focal-Plane Array Processor for Morphological Image Processing," Proceedings of Fifth Annual IEEE International ASIC Conference and Exhibit, Rochester, NY, Sept. 1992, pp. 423-426.

[8] A. L. Fisher, "Scan Line Array Processors," Annual Symposium on Computer Architecture, 1986, pp. 338-345.

[9] A. Gentile *et. al.*, "Implementation of VQ-based Image Compression Algorithm on SIMPil," Fine-Grained Laboratory Technical Report 96-4, Georgia Institute of Technology, May 1996.

Table 2: VQ Performance Measurements

Distortion Measure	Platform	Clock Rate	Frame Size	Codebook Size	Encoding Time (Sec)	Performance (blocks/sec)
Euclidean	2048 node Maspar 1	12 MHz	128×128	512	**1.79**	**1,144**
Euclidean	32 T800 nodes	20 MHz	512×512	256	**4.8**	**3,413**
Absolute	256 node SIMPil	50 MHz	256×256	256	**0.0294**	**139,000**

[10] A. Gentile, H. H. Cat, F. Kossentim, F. Sorbello, and D. S. Wills, "Real-Time Implementation of Full-Search Vector Quantization on a Low Memory SIMD Architecture," IEEE Data Compression Conference, Snowbird, Utah, Apr. 1996, page 438.

[11] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*, Kluwer Academic Publisher, 1992.

[12] T. Gross, H. T. Kung, M. Lam, J. Webb, "Warp as a Machine for Low-Level Vision," Proceedings of the 1985 IEEE International Conference on Robotics and Automation, Mar. 1985, pp. 790-800.

[13] M. W. Haney, M. P. Christensen, "Smart Pixel Based Viterbi Decoder," Optical Computing, 1995 Technical Digest Series, vol. 10, 1995, pp. 99-101.

[14] Henderson, and K. S. Thyagarajan, "Implementation of VQ Algorithms on a Reconfigurable Array Processor," Professional Paper AD-A236 483/4/XAB, Naval Ocean Systems, San Diego, CA, May 1991.

[15] F. Hieronymi, E. H. Bottcher, E. Droge, D. Kuhl, D. Bimberg "High performance large area InGaAs metal-semiconductor-metal photodetectors," IEEE Phot. Tech. Letter, vol. 5, no. 8, 1993, pp. 910-913.

[16] M. Klingenstein, J. Kuhl, J. Rosenzweig, C. Moglestue, A. Hulsman, Jo. Schneider and K. Kohler, "Photocurrent gain mechanisms in metal-semiconductor-metal photodetectors," Solid State Elect. vol. 37, no. 2, 1994, pp. 333-340.

[17] C. B. Kuznia, A. A. Sawchuk, and L. Cheng, "FET-SEED Smart Pixels for Free-Space Digital Optics Systems," Optical Computing, 1995 Technical Digest Series, vol. 10, 1995, pp. 108-110.

[18] R. P. Leavitt and J. L. Bradshaw, "Molecular-Beam-Epitaxial growth and characterization of high quality alloys and multiple quantum wells on InP using a post-evaporation-heated arsenic source," J. Appl. Phys. 76, 1994, pp. 3429-3442.

[19] Lee, L. C. Liu, A. K. Chan, and C. K. Chui, "A Parallel Vector Quantization Algorithm for Single-Instruction-Multiple-Data (SIMD) Multiprocessor Systems," Proceedings of the 5th Data Compression Conference, Snowbird, UT, Mar. 1995, page 479.

[20] P. Maragos and R. W. Schafer, "Morphological Systems for Multidimensional Signal Processing," Proceedings of the IEEE, vol. 78, no. 4, Apr. 1990, pp. 690-710.

[21] C. Mead, *Analog VLSI and Neural Systems*, Addison-Wesley, Reading, Massachusetts, 1989.

[22] J. B. D. Soole, H. Schumacher, R. Bhat, and M. A. Koza, "High-speed performance of OMCVD grown InAlAs/InGaAs MSM photodetectors at 1.5 μm and 1.3 μm wavelengths," IEEE Phot. Tech. Letter, vol. 1, no. 8, 1989, pp.250-252.

[23] Connection Machine Model CM-2 Technical Summary, Thinking Machines Corporation, Version 51., May 1989.

[24] O. Vendier, N. M. Jokerst and R. P. Leavitt, "Thin film inverted MSM photodetectors," IEEE Phot. Tech. Letter, vol. 8, no. 2, 1996, pp. 266-268.

[25] D. S. Wills *et. al.*, "High Throughput, Low Memory Applications on the Pica Architecture," to appear in IEEE Transactions on Parallel and Distributed Systems.

[26] W. F. Wong and K. T. Lua, "A Preliminary Evaluation of a Massively Parallel Processor: GAPP," Microprocessing and Microprogramming, vol. 29, no. 1, July 1990, pp. 53-62.