

A reprint from

Optical Engineering

28(8), 865-871 (August 1989).

ISSN 0091-3286

ARCHITECTURES FOR FOCAL PLANE IMAGE PROCESSING

Eric R. Fossum

Columbia University
Department of Electrical Engineering
1312 S.W. Mudd Building
New York, New York 10027

Architectures for focal plane image processing

Eric R. Fossum, MEMBER SPIE

Columbia University
Department of Electrical Engineering
1312 S.W. Mudd Building
New York, New York 10027

Abstract. Architectures for focal plane image processing are discussed. On-chip image preprocessing for solid-state imagers using analog CCD circuits is described for low, medium, and high density detector arrays. A spatially parallel architecture for low density, high throughput applications is described. For sparse illumination or event detection, a content-addressable architecture is proposed. A new pipelined vector pixel processor architecture for medium density infrared staring focal plane arrays is described. Neighborhood reconstruction during serial readout of high density TV-quality imagers for a pixel processor is considered using delay and analog frame memory techniques. The potential of on-chip read/write analog frame memory for image transformation and frame-to-frame processing is discussed.

Subject terms: charge-coupled devices; focal plane image processing; solid-state imager sensors; space surveillance; focal plane array; machine vision.

Optical Engineering 28(8), 865-871 (August 1989).

CONTENTS

1. Introduction
2. Imager readout technology
3. Focal plane image processing issues
4. Low density arrays
5. Medium density arrays
6. High density arrays
7. Conclusions
8. Acknowledgments
9. References

1. INTRODUCTION

Solid-state imaging devices have evolved rapidly in the past five years, and this trend is expected to continue. The consumer market for home video cameras and the industrial market for machine vision and security cameras have provided a strong driving force for this evolution. Technology derived from aerospace research and development in the area of visible and infrared imagers has further pushed the state of the art.¹⁻³

Although the pace in computer electronics has also been brisk, technology for processing electronic images acquired by solid-state devices has been lagging and now represents a severe bottleneck for real-time image processing. Television-quality imagers capture images containing between 200,000 and 300,000 pixels per frame at frame rates of 30 Hz. Thus, a real-time image processor must contend with data rates of the order of 10 Mpixels/s. For simple image processing (e.g., edge detection), approximately 100 elementary operations such as addition and subtraction must be performed per pixel, resulting in a corresponding throughput requirement of 1 billion operations per second (Bops). Machine vision applications may require higher frame rates and more sophisticated processing, while higher definition imagers require more pixels to be processed, making the throughput

requirement on the image processor potentially climb to 100 Bops. Such throughput is difficult to achieve in a general-purpose digital computer without massively parallel processing.

Recently, some special-purpose digital integrated circuits have been proposed to relieve this bottleneck.⁴⁻⁶ These circuits generally consume between 500 and 1500 mW of power, and most are restricted in function. Their digital nature requires high speed analog-to-digital (A/D) conversion of the pixel data between the imager output and the digital processor input, which is often a major source of power consumption. Most digital processors operate on a subset of the image data, and additional circuitry to feed the processor pixel blocks is required. Digital storage of a single frame of imagery for frame-to-frame operations or for buffer memory requires approximately 2 to 4 Mbits, with consequent real-estate and power consumption. Thus, a complete digital image processor system may require dozens of ICs and several printed circuit boards of electronics.

Analog image processing, which occurs prior to A/D conversion, has several advantages. These include lower power consumption, lower real-estate consumption, and no A/D converter. Some approaches for implementing analog image processing with CCDs have been explored, with processing occurring either in separate ICs or on the same IC as the imager itself.⁷⁻¹⁵ In general, analog approaches to image processing are perceived to suffer from more limited accuracy, from design and fabrication complexity, and from the dynamic nature of CCDs. In practice, these perceived problems are not particularly critical, and future image processing circuits may combine analog and digital functions in a CCD/CMOS process.

In this paper, architectures for image processing circuits located on the same chip as the imager itself, or hybridized with the imager, are considered. Such focal plane image processing* has a high potential for achieving high throughput

Paper 2618 received Aug. 22, 1988; revised manuscript received March 14, 1989; accepted for publication March 14, 1989.
©1989 Society of Photo-Optical Instrumentation Engineers.

*This is actually image plane image preprocessing, but the term "focal plane array" has become an accepted part of the technical language.

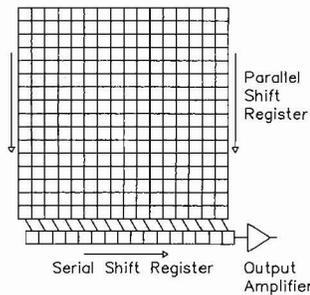


Fig. 1. Schematic illustration of a CCD imager readout structure.

with low power and chip-area consumption. The choice of architecture depends on the detector array density, the total number of pixels, the frame rate, and the processing complexity required by the application. Imager readout technology is described first, followed by discussions of architectures for low, medium, and high density detector arrays.

2. IMAGER READOUT TECHNOLOGY

Technology currently employed for detector readout is central to a discussion of focal plane image processing architectures. Currently, there are three major readout or multiplexer technologies employed for solid-state imager readout.¹ Multiplexers deliver pixel data one row at a time, and within each row, column positions are scanned horizontally. During the readout, the pixel data are shielded from light either by use of a fast frame-transfer operation or by interline transfer, and by the use of a light shield.

In the first readout technology, the CCD approach, row data are shifted in parallel in the columnar direction by use of a slow CCD analog shift register. The bottom row is loaded, in parallel, into a serial CCD shift register that, in turn, shifts the data in the row direction to an output amplifier, as shown in Fig. 1. There is a tradeoff in the CCD approach between fill factor (percentage of the pixel unit cell used for detection) and overall chip size. Chip size is important from an IC manufacturing and packaging perspective, but imager performance is improved with increasing fill factor. Interline transfer yields a smaller chip size but a relatively poor fill factor; frame transfer improves fill factor at the expense of chip size.

The CCD multiplexer delivers the charge generated within each unit cell to the output amplifier by successive transfers. The charge transfer efficiency (CTE) is defined as the fraction of charge successfully transferred in each transfer process. A charge packet in a four-phase CCD might undergo 3000 to 4000 transfers before reaching the output amplifier, so a chip with a CTE of 99.995% would deliver only about 90% of the original charge packet to the output amplifier (and portions of other charge packets as well). This is adequate in most applications. Noise is introduced in the charge transfer process, but in a buried-channel CCD the multiplexer noise is usually smaller than the background photon shot noise.

The second readout technology is the MOS X-Y or direct readout approach, in which pixels are addressed like digital bits in a random access memory. Each pixel is individually addressed, and the analog charge packet is placed on a

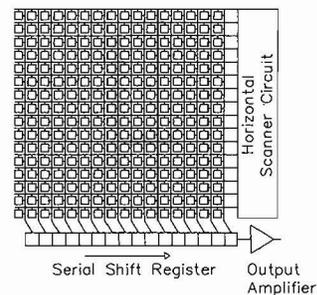


Fig. 2. Schematic illustration of a MOS-CCD imager readout structure.

global readline connected to an output amplifier. The capacitance for the readline may be large, making recovery of the pixel signal difficult. The resultant dynamic range is comparable to surface-channel CCD multiplexers but generally poorer than buried-channel CCD multiplexers. However, the MOS X-Y readout circuitry is easier to design and simpler to fabricate with high yield than are CCD circuits. The full random access capability of the direct readout architecture is generally not utilized in imaging systems, and the array is scanned in a rasterlike manner (usually by on-chip scanner circuits). The ability to focus the readout in a particular subregion of the array is interesting but so far unexploited.

The third readout architecture is the most recent and is a hybrid of the previous two approaches. The MOS-CCD approach is to use direct readout for rows but a CCD serial shift register (loaded in parallel with the row data) to perform column multiplexing. The high readline capacitance of the MOS X-Y multiplexer is reduced, and clocking is simplified. The transfer efficiency of the readline-to-CCD-bucket process is typically 95%, with overall transfer efficiency further reduced by the serial CCD shift register. This architecture is shown schematically in Fig. 2.

In general, solid-state imagers for scientific applications have a single source-follower output stage. Commercially oriented imagers may include on-chip sample and hold circuitry and drive amplifiers. Video output for images that are digitally processed are sent to off-chip scaling amplifiers and then converted to digital format using an A/D converter. On-chip A/D conversion, desirable for improvement of dynamic range and system simplification, is rarely performed, although a CCD-based A/D converter for this purpose is currently being prototyped.[†]

The choice of readout architecture for solid-state imagers depends on a number of factors ranging from fabrication capability and design expertise to system considerations. It appears that architectures utilizing analog devices and circuits compatible with digital circuit technology are the most likely candidates for success.

3. FOCAL PLANE IMAGE PROCESSING ISSUES

The degree to which an electronic representation of the photon flux is altered or manipulated on the focal plane can vary significantly. For example, simple transimpedance amplifier (TIA) circuits placed between the detector and

[†] S. E. Kemeny and E. R. Fossum, unpublished.

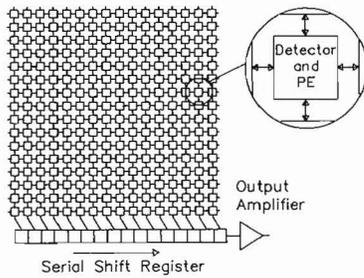


Fig. 3. Spatially parallel architecture for focal plane image processing. Each pixel is associated with a processing element.

multiplexer for buffering purposes barely constitutes focal plane image processing since they do not alter the spatial content of the image. On the other hand, circuitry for implementing on-chip discrete image transforms or image halftoning represents significant image processing rather than preprocessing. The extent of the processing possible is dictated primarily by the available chip area, which in turn depends on the detector array density and size.

There are strong motivations for performing image processing on the focal plane of the imager. Although digital circuitry is often easier to design and fabricate, analog signal processing uses little power and real estate and avoids the need for prior (and higher resolution) A/D conversion. The low power aspect of analog signal processing is especially important in aerospace applications, in which the focal plane array may be cooled or in which total system power is limited (e.g., satellites or missile systems). Thus, most of the ensuing discussion is made with analog circuitry in mind.

A major reason for considering focal plane image processing is that it avoids the introduction of noise and distortion through off-chip driving of the multiplexed output. Output amplifiers are often a major source of noise in readout circuits, and pickup on output lines is also a difficulty. Distortion introduced by the output amplifier can require nonlinear gain compensation.

Additionally, focal plane image processing has the potential to reduce the bandwidth of the signal driven off-chip. For example, for thresholded images (which become binary in nature) the requirements on off-chip drive circuitry and A/D converter resolution are significantly reduced. Alternatively, for video compression, frame-to-frame comparison of pixels can be performed to transmit only those pixels that change.

Interconnection of processing elements in a parallel processor adds a substantial hardware and power burden to conventional digital computing systems. On-chip spatially parallel processing elements (described in Sec. 4) operating in the analog domain have particularly simple interconnect structures, which adds to the desirability of performing focal plane image processing.

Finally, on-chip processing has the potential to alleviate bottlenecks in massive detector arrays that are sparsely illuminated or have sparse event occurrences. For example, the detection of a sudden bright spot in an otherwise deactivated array might trigger full readout of the array. Surveillance of large fields of view is another example in which full array readout could be avoided until motion is detected by on-chip processing circuitry.

There is an unfortunate relationship between array size and processor complexity that exists for all image processing systems and is particularly acute for focal plane image processing. In the latter, for a given chip size, as the detector array size becomes larger the throughput requirements of the processor become more stringent as the chip area available for image processing is reduced. Three-dimensional stacked or hybridized structures (e.g., flip-chip or z-plane topologies) can be used to retain the advantages of focal plane image processing while extending the real estate available for the processor.^{16,17}

4. LOW DENSITY ARRAYS

Low density detector arrays, in which chip real estate is readily available, offer the largest opportunity for focal plane image processing. A low density array is defined as one having a detector pitch greater than approximately $50L$, where L is a typical feature size. For example, a low density detector array with a feature size of $L = 3 \mu\text{m}$ would have a detector pitch greater than approximately $150 \mu\text{m}$. Low density arrays are used in low resolution applications (e.g., event or motion detection) or in low carrier generation (e.g., low light or poor quantum efficiency) applications. In the former case, detectors may be monolithically integrated within the unit cell, and in the latter case, hybrid flip-chip configurations or amorphous silicon overlayers provide high detector fill factor without sacrificing readout chip real estate.

Circuitry for performing image processing functions can be placed within the available real estate in the unit cell. The circuitry, or processing element (PE), may provide only buffer/amplifier functions for the unit cell or more sophisticated functions. For example, circuitry simulating neuron behavior to perform motion or edge detection has been prototyped for low density arrays.¹⁸ Implemented as a switched capacitor CMOS IC, the unit cell size is $164 \mu\text{m} \times 143 \mu\text{m}$ and the array size is 48×48 .

Such an architecture is termed spatially parallel¹⁹ since the physical interconnect relationship between processing elements corresponds to the spatial connectivity of the image, as shown in Fig. 3. Spatially parallel general-purpose charge-domain analog computing circuitry can also be located in the unit cell to provide more sophisticated kernel functions.^{9,10,12,14} Implemented in a double-polysilicon CCD process, the detector pitch is typically 150 to $200 \mu\text{m}$. Processing elements are designed to communicate with their nearest neighbors and can perform functions such as smoothing, signal averaging, edge detection, and A/D conversion. Such a focal plane image processor is a single-instruction, multiple-data (SIMD) architecture. Unlike the neural network circuit, it is digitally programmable by applying various clocking sequences.

Spatially parallel architectures can also be implemented with the z-plane technology.¹⁷ In this technology, a laminated stack of perhaps 128 chips is mated perpendicularly to a detector array using flip-chip technology. The detector array pitch is typically $100 \mu\text{m}$, and the edge of each chip becomes mated to one detector array column, providing unit cell real estate in the z-direction. The z-plane hybridized approach has the advantage of providing more real estate per unit cell for pixel processing, such as nonuniformity correction or multiple frame buffering,²⁰ but

it is significantly more difficult to manufacture. Furthermore, PE communication in the columnar direction is easily achieved, but communication in the row direction (perpendicular to the lamination plane) must be performed through a backplane connection. It is expected that as this technology matures and becomes applicable to medium and perhaps high density detector arrays, it will provide the ultrahigh throughput required in future real-time image processing systems.

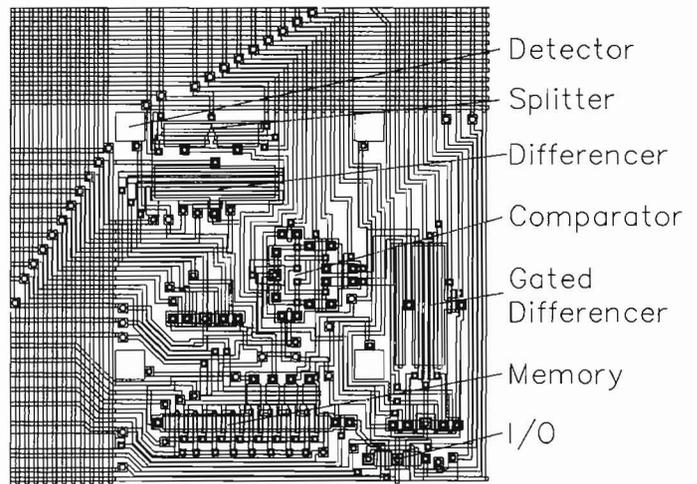
Three-dimensional integrated circuits utilizing laser recrystallized silicon for spatially parallel image processing have also been reported.¹⁶ Digital CMOS technology is used after unit cell A/D conversion (2 bits) to perform some logic functions. Since it is a digital PE requiring approximately 7000 transistors per unit cell, the unit cell size is nearly $1000 \mu\text{m} \times 1000 \mu\text{m}$.

The design of the PE depends strongly on the application, although some general-purpose operations can be anticipated. A serious difficulty in PE design is the reduction of the number of control lines required to operate the PE. These control lines can consume a significant portion of the available chip area. Decoding of control signals within the unit cell can reduce the number of control lines, but the decoding circuitry also requires significant unit cell real estate.

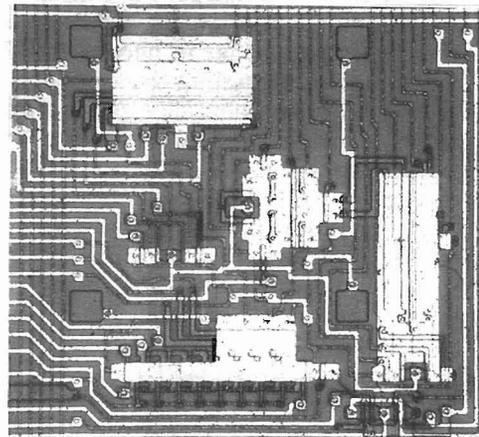
The throughput of a spatially parallel architecture can be very high. For an array size of $N \times M$, the throughput increases simply as NM . For example, a $10 \text{ mm} \times 10 \text{ mm}$ imager with a detector pitch of $150 \mu\text{m}$ could have approximately 4000 pixels. Assuming 100 elemental operations per pixel, a serial processor operating at a rate of $1 \mu\text{s}/\text{operation}$ would take 400 ms to process an acquired image, corresponding to a frame rate of 2.5 Hz. However, a spatially parallel architecture could process the image at a frame rate of 10,000 Hz! Thus, the image processing throughput is taken from a realm that is barely real time to one acceptable for ultrahigh velocity intercept applications. Alternatively, the PE design can be simplified by employing slower but more efficient circuits. The spatially parallel architecture is limited not by image processing functional throughput but by readout (or I/O) rate.

A lower degree of parallelism in a spatially parallel architecture can also be utilized to conserve real estate. For example, a single PE could serve p pixel detectors by time-domain multiplexing, reducing the parallelism by the factor p . The penalty for a lower degree of parallelism in a SIMD machine is increased software complexity for shuffling the data. This may be more cumbersome in an analog-circuit-based PE.

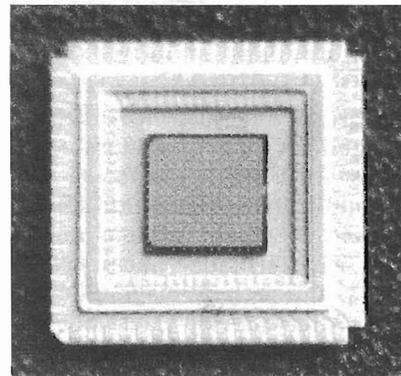
In a CCD-based spatially parallel image processor currently being fabricated,¹⁰ an array of 48×48 detectors ($180 \mu\text{m}$ pitch) and 24×24 charge-coupled computer PEs ($p=4$) are monolithically integrated on a $9.4 \text{ mm} \times 9.4 \text{ mm}$ chip, as shown in Fig. 4. Each PE is designed to perform nearest-neighbor I/O, magnitude comparison, differencing, and halving in the analog charge domain. A bidirectional stack is used for local memory. This processor, which is digitally programmable, can perform algorithms for smoothing, thresholding, edge detection, and A/D conversion. The chip is projected to provide a maximum throughput of 0.5 Bops at a power cost of 12 mW. Since the CCD technology is capacitive in nature, lower clocking rates would scale the power down with throughput.



(a)



(b)



(c)

Fig. 4. Prototype spatially parallel focal plane image processor chip (IRET) currently under test.¹⁰ (a) Layout of unit cell. (b) Photograph of unit cell (size is $360 \mu\text{m} \times 360 \mu\text{m}$). (c) Photograph of complete IRET chip. Chip size is $9.4 \text{ mm} \times 9.4 \text{ mm}$.

Ideally, PEs would be addressed externally in a direct readout method since the output would be buffered or digital and thus immune to the parasitic effects described previously. Other readout techniques may be more applicable in later generations of spatially parallel focal plane image processors. For example, a proposed content-addressable array readout^{15,21} for searching for particular

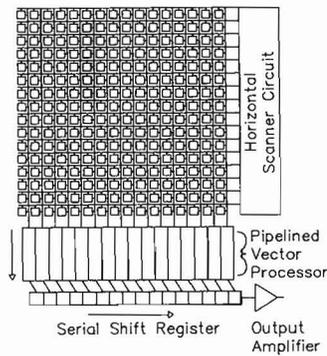


Fig. 5. Schematic illustration of proposed pipelined vector processor architecture for column-parallel pixel processing.

pixel values or subpatterns would be useful for sparse illumination or motion detection applications. In this case, unit cells that are illuminated above a particular threshold or in which a frame-to-frame change in photon flux has been detected set a system flag requesting readout. The unit cell then places its address on a global bus. Arbitration of bus contention can be achieved in a number of ways. For example, priority encoding based on location, asynchronous enabling of unit cells, and multiple bus lines are some ways to ensure that only one address is on the bus in a given cycle. Thus, unlike normal architectures in which an address is given and the data are returned, here the data are prescribed and the address of the pixel with the prescribed data is returned.

The readout of contours generated by edge detection is a second example for which nonconventional readout might be more useful. In this case it would be desirable to have adjacent pixels on the same edge read out sequentially. Such "stitched" readout could be implemented on the focal plane by appropriate PE design.

5. MEDIUM DENSITY ARRAYS

Medium density arrays have detector pitches between approximately $10L$ and $50L$. The small unit cell size prohibits all but the simplest PE circuits from being implemented. A unit cell PE would permit signal modification prior to readout to improve the overall dynamic range of the sensor array. Possible PEs include linear and logarithmic amplifiers, buffers, and perhaps magnitude comparison for A/D conversion. Alternatively, externally adjustable unit cell amplifier gain would provide for adaptive imaging or perhaps fixed-pattern noise removal.

For low and medium density arrays, a new approach would be to have image processing circuitry located at the bottom of each column, as illustrated in Fig. 5. For an $N \times M$ array, the degree of parallelism becomes just M , reduced by the factor N over spatially parallel architectures. The throughput requirement of each PE is increased over that of a spatially parallel architecture, but throughput requirements for real-time processing can most likely be met.

This new architecture, even for low density arrays, has the advantage of making more real estate available for each PE. A pipelined PE design (in the columnar direction) can be readily achieved. The resulting architecture is referred to as a pipelined vector processor. The architecture would reduce

the total fixed-pattern noise introduced by the analog PE circuits (since there are fewer PEs) and is compatible with time delay and integration (TDI) imaging and MOS-CCD readout structures. However, the possibility of pixel data distortion between the detector and the PE is increased.

The pipelined vector processor architecture has the highest potential for monolithic focal plane image processing for infrared and other non-TV-quality images.

6. HIGH DENSITY ARRAYS

High density arrays are defined as arrays having a detector pitch less than approximately $10L$, thereby prohibiting PE circuitry. The imager chip area is used for detection and readout. For some dedicated applications and harsh environments it may be possible to fan out an otherwise dense array using electronic or optical means (e.g., optical fibers) so that the detector density is reduced, but in this discussion, it is assumed that the on-chip image processing circuitry must be located beyond the region of detection and readout. One exception to this is the use of the readout circuitry in a nonconventional manner to perform some limited image processing functions such as Gaussian convolution.¹³ In this case, buckets are clocked to effect charge mixing and simulate a diffusion process.

The advantages of on-chip focal plane image processing given a serial readout of the detector array are more limited since only the off-chip transmission of the serial data would be avoided by on-chip processing. If the on-chip processor utilizes significant chip area or I/O pads or introduces noise through clock feedthrough, the advantages of on-chip processing could be negated.

Many image processing tasks involve the convolution of the image data with a 3×3 or 5×5 kernel; i.e., a processed pixel is a weighted sum of its surrounding neighborhood. The pixel may be further processed by applying nonlinear operators on it and its neighbors. These two steps may be repeated several times before the image processing task is complete. Rarely does the convolution kernel need to exceed a size of 3×3 , especially if multiple convolutions can be performed. Thus, the reconstruction of a small local neighborhood is required, after serial readout, in order to process a given pixel.

The image processing architecture can be divided into two parts: neighborhood reconstruction and pixel processing. These two parts are intimately interrelated. For example, if the pixel processor destructively senses neighborhood data, those data must be regenerated prior to processing since a 3×3 neighborhood implies that each pixel must be utilized at least nine times. The necessity of regeneration in turn influences the manner in which the neighborhood is reconstructed.

Methods for neighborhood reconstruction using delay lines have been demonstrated both on-chip²² and off-chip.⁷ In the on-chip case, serial output data from an $N \times M$ detector array were loaded into a CCD delay line $2M$ stages long. The delay line pixel data were sensed nondestructively using a floating-gate technique at the beginning, middle, and end of the delay line. Three adjacent pixels were sensed at each location, the locations corresponding to the same columns on adjacent rows. Thus, nine analog outputs corresponding to the local neighborhood were simultaneously provided to an off-chip image processor. The off-chip method used two

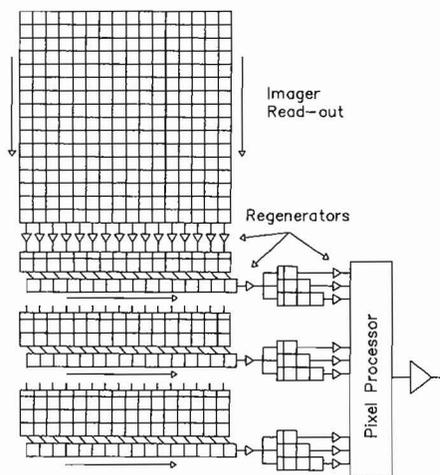


Fig. 6. Schematic illustration of proposed architecture for neighborhood reconstruction and pixel processing based on on-chip regeneration and delay.

separate delay lines, each M stages long, with data regeneration required between the two. A drawback of both approaches is that pixels delayed by $2M$ stages can suffer from CTE effects. A technique utilizing a floating-diffusion direct-readout approach has also been reported, but it suffers from fixed pattern noise and low speed.²³

An improved version of the delay approach to neighborhood reconstruction is proposed in Fig. 6. Pixel regenerators located at the bottom of each column fan out the pixel data into three row reconstruction registers. The first register is immediately loaded into a serial horizontal shift register and shifted toward the pixel processor. The second and third registers are delayed by one and two rows, respectively, prior to loading into their serial registers. At the end of each serial register, the data are regenerated a second time for fan-out into three column reconstruction registers. These add a further delay to the pixel data. The three column reconstruction registers for each of the three row reconstruction registers provide a total of nine simultaneous inputs to the 3×3 neighborhood pixel processor. To maximize throughput, a pipelined approach to the on-chip pixel processor might be employed.

Neighborhood reconstruction might also be performed utilizing a buffer analog frame memory. An analog frame memory would require approximately the same amount of chip area as the imager, but if the frame store region of a frame-transfer mode imager is used, the total chip area remains approximately the same. The data in the frame memory are stored in a spatially parallel format. Ideally, the analog frame memory would have nondestructive random-access readout capability to avoid the pixel delay circuitry described above. Without complex unit cell circuitry in the frame memory, neighborhood access time would increase since there is only one readline per column and multiple read cycles might be needed to obtain the 3×3 neighborhood. However, with buffer memory added to the pixel processor, the access time could be reduced since actually only three new pixel data are added to the neighborhood as the array is scanned, and these could be sensed in parallel.

The full potential of a frame buffer memory could be realized if the unit cell included read *and* write capability.

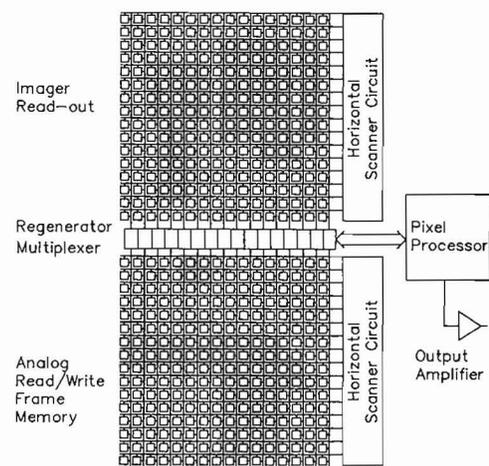


Fig. 7. Schematic illustration of proposed architecture using on-chip read/write analog frame memory for focal plane image processing.

(In a sense, frame buffer memory in a frame-transfer imager already does have read and write capability.) The unit cell complexity would increase, as would the real estate requirement. However, if a hybrid flip-chip or silicon-on-insulator 3-D IC approach is used, the overall chip size would not increase significantly. The proposed analog frame memory architecture is illustrated in Fig. 7.

With a read and write frame memory, algorithms having slow convergence to the final processed state could be executed. The overall frame rate of the system might be diminished, but some applications do not require high frame rates. A good example of this is image halftoning,^{24,25} in which neural-network-like algorithms might be used, with convergence occurring within a few hundred cycles. Since the resultant image is binary, off-chip transmission bandwidth requirements drop considerably. A 1 Hz frame rate would be acceptable for facsimile transmission and other document scanning systems.

A second application that can be implemented with a frame buffer memory architecture of Fig. 7 with nondestructive readout is image transformation. For example, implementation of the discrete cosine transform (DCT) using CCDs¹¹ is a good candidate for this architecture. Other transforms to enable image data compression²⁶ are currently under investigation.

Further increasing the chip complexity but also enhancing processing capability would be a second analog frame memory. This memory could be used for frame-to-frame operations such as motion or event detection. It might also be used for the temporary storage of intermediate results. The real-estate penalty for a second frame memory is significant, and off-chip processing could become an attractive alternative. The off-chip processing of two frames of analog memory using CCD-like circuits continues to provide power and real-estate advantages. Cooled CCD circuits can provide charge storage times measured in hours, and proximity to a cooled detector array (i.e., in the same dewar) would help reduce noise and pickup. The z -plane architecture might also be employed for frame memory applications.²⁰

Design of the pixel processor for either on-chip or off-chip processing is dependent on the application and choice

of technology. Fixed algorithm architectures are easier to design and require fewer clock control lines but do not offer flexibility. General-purpose processors provide flexibility for multiple applications but may be less efficient. Pipelined architectures require more chip area, but the throughput scales as the number of stages in the pipeline with a small signal delay as the penalty.

7. CONCLUSIONS

Focal plane image processing, particularly in the analog domain, shows promise for reducing the severe throughput, power, and real-estate problems associated with current digital technology. Several new architectures have been proposed and discussed. These included a spatially parallel architecture for low density arrays, a pipelined vector processor architecture for medium density arrays, and a read/write frame memory architecture for high density arrays. The choice of architecture depends on the application. The higher the degree of parallelism, the higher the throughput, power, and chip-area consumption.

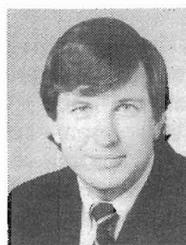
It can be anticipated that focal plane image processing in solid-state imaging systems will develop rapidly in the next few years. Since system input is expected to remain analog in nature, and high level processing digital, such systems will need to combine the best features of analog and digital processing circuitry.

8. ACKNOWLEDGMENTS

The author gratefully acknowledges thoughtful discussions with R. Colbeth, S. Kemeny, E. Eid, C. Zukowski, G. Hart, D. Anastassiou, and H. Meadows of Columbia University. The encouragement provided by A. Chiang, W. Kosonocky, C. Mangelsdorf, G. Anderson, R. Sahai, S. Othmer, M. Blouke, R. Bredthauer, M. White, and R. Barker is greatly appreciated. Portions of this work were supported by a National Science Foundation Presidential Young Investigator Award, the NSF Center for Telecommunications Research at Columbia University, an IBM Faculty Development Award, and a grant from Analog Devices, Inc.

9. REFERENCES

- W. Kosonocky, ed., Special Issue on Solid-State Image Sensors, IEEE Trans. Electron Devices ED-32(8), (1985).
- E. Dereniak, ed., Special Issue on Infrared Focal Plane Arrays, Opt. Eng. 26(3), (1987).
- J. Janesick, ed., Special Issues on Charge-Coupled Devices, Opt. Eng. 26(8,9,10), (1987).
- P. A. Ruetz and R. W. Brodersen, "Architectures and design techniques for real-time image processing ICs," IEEE J. Solid-State Circ. SC-22(2), 233-250 (1987).
- H. Yamada, K. Hasegawa, T. Mori, H. Sakai, and K. Aono, "A microprogrammable real-time image processor," IEEE J. Solid-State Circ. SC-23(1), 216-223 (1988).
- T. Denayer, E. Vanzieleghem, and P. G. A. Jespers, "A class of multiprocessors for real-time image and multidimensional signal processing," IEEE J. Solid-State Circ. SC-23(3), 630-638 (1988).
- G. R. Nudd, P. A. Nygaard, G. D. Thurmond, and S. D. Fouse, "A charge-coupled device image processor for smart sensor applications," in *Image Understanding Systems and Industrial Applications*, R. Nevatia, ed., Proc. SPIE 155, 15-22 (1978).
- E. R. Fossum, "Charge-coupled analog computing elements and their application to smart image sensors," Ph.D. thesis, Yale Univ. (1984).
- J. D. Joseph, P. C. T. Roberts, J. A. Hoschette, B. R. Hanzal, and J. C. Schwanebeck, "A CCD-based parallel analog processor," in *State-of-the-Art Imaging Arrays and Their Applications*, K. N. Prettyjohns, ed., Proc. SPIE 501, 238-241 (1984).
- E.-S. Eid and E. R. Fossum, "CCD focal plane array analog image processor," in *Real Time Signal Processing XI*, L. P. Letellier, ed., Proc. SPIE 977, 298-306 (1989).
- A. M. Chiang, "A video-rate CCD two-dimensional cosine transform processor," in *Visual Communications and Image Processing II*, T. R. Hsing, ed., Proc. SPIE 845, 2-5 (1987).
- P. R. Beaudet, "CCD focal plane convolver (smart eyeball)," OSA Topical Meeting on Machine Vision, FD6/1-4 (Incline Village, Nev., March 1985).
- J. P. Sage, "A high speed analog two dimensional gaussian image convolver," OSA Topical Meeting on Machine Vision, FD5/1-4 (Incline Village, Nev., March 1985).
- E. R. Fossum, "Charge-coupled computing for focal plane image preprocessing," Opt. Eng. 26(9), 916-922 (1987).
- E. R. Fossum, "Charge domain analog signal processing for detector arrays," Nucl. Instrum. Methods A275, 530-535 (1989).
- T. Nishimura, Y. Inoue, K. Sugahara, S. Kusunoki, T. Kumamoto, S. Nakagawa, M. Nakaya, Y. Horiba, and Y. Akasaka, "Three-dimensional IC for high performance image signal processor," in *Tech. Dig. 1987 IEEE Int. Electron Devices Meeting*, pp. 111-114 (1987).
- J. C. Carson, "Applications of advanced 'Z' technology focal plane architectures," in *Infrared Detectors and Arrays*, E. L. Dereniak, ed., Proc. SPIE 930, 164-182 (1988).
- M. A. Sivilotti, M. A. Mahowald, and C. A. Mead, "Real-time visual computations using analog CMOS processing arrays," in *Proc. 1987 Stanford Conf. on Advanced Research in VLSI*, MIT Press, Cambridge, Mass. (1987).
- D. H. Schaefer and P. Strong, "Tse computers," Proc. IEEE 65(1), 129-138 (1977).
- S. E. Kemeny, E.-S. Eid, and E. R. Fossum, "Novel CCD image processor for Z-plane architecture," in *Materials, Devices, Techniques, and Applications for Z-Plane Focal Plane Array Technology*, J. C. Carson, ed., Proc. SPIE 1097, in press (1989).
- O. Wing, "A content-addressable systolic array for sparse matrix computation," J. Parallel Distrib. Comput. 2, 170-181 (1985).
- J. E. Hall, J. F. Breitzmann, M. M. Blouke, and J. T. Carlo, "A multiple output CCD imager for image processing applications," in *Proc. 1978 IEEE Int. Electron Devices Meeting*, pp. 415-418 (1978).
- J. P. Vermeiren and G. J. Declerck, "A novel CCD output structure for picture processing applications," in *Proc. 1985 IEEE Custom Integrated Circuits Conf.*, pp. 128-131 (1985).
- T. Bernard, P. Garda, A. Reichart, B. Zavidovique, and F. Devos, "Design of a half-toning integrated circuit based on analog quadratic minimization by non-linear multistage switched capacitor network," in *Proc. 1988 IEEE Symp. on Circuits and Systems*, pp. 1217-1220 (1988).
- D. Anastassiou, "Neural net based digital halftoning of images," in *Proc. 1988 IEEE Symp. on Circuits and Systems*, pp. 507-510 (1988).
- H. Torbey and H. Meadows, "System and method for lossless hierarchical image coding/decoding," CTR Tech. Rept. 094/88-22, Columbia Univ. (1988).



Eric R. Fossum received the Ph.D. degree from Yale University in 1984 and joined the Department of Electrical Engineering at Columbia University, where he is now an associate professor. Prof. Fossum and his graduate students are primarily working in the area of charge-coupled devices. Architectures for focal-plane image processing, including A/D conversion, are being investigated in silicon, while very high speed (GHz) charge-transfer devices are under investigation in GaAs. Novel silicon and GaAs device structures for the direct connection of optical fibers to integrated circuits are also being studied. Prof. Fossum has published more than 30 technical papers and organized the 1986 IEEE Workshop on Charge-Coupled Devices. He has received the IBM Faculty Development Award and the Analog Devices Career Development Award. In 1986, he received the National Science Foundation Presidential Young Investigator Award.