On-Chip Focal-Plane Image Processing

Eric R. Fossum
Imaging Systems Section, Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109 USA

Abstract

Focal-plane image processing refers to the integration of image acquisition and image processing functions in the same integrated circuit, or adjacent ICs. In this paper, the motivation for focal-plane image processing will be developed. Architectures for achieving such integration in high, medium, and low density imagers will be discussed. Future directions will be suggested.

Introduction

Focal-plane image processing refers to the integration of image acquisition and image/signal processing functions on the focal-plane of an imaging system. In this paper, on-chip focal-plane image processing will be discussed, meaning that the integration of acquisition and processing is on the same integrated circuit (IC).

There are several motivations for pursuing focal-plane image processing. These include increasing sensor imaging performance, reducing sensor system size and weight, and increasing sensor system throughput. For example, imaging performance may be improved through on-chip reduction of readout noise, exposure control, motion stabilization, and transient event suppression (e.g. cosmic rays). Parallel signal processing on the focal-plane, while constrained by available integrated circuit area, is inherently compatible with the spatially parallel nature of the imaging process. By performing image and/or signal processing functions such as smoothing, edge detection, motion detection, uniformity correction, data conversion, compression and stereo vision, in parallel on the focal-plane, system throughput can be improved. Analog processing circuits typically take less power and chip area than their digital counterparts so that performing analog processing on the focal-plane, distortion of the signal by the output amplifier and subsequent off-chip interconnects can be avoided. Finally, an additional motivation is that electronics integration has historically led to improved circuit reliability and decreased cost.

It is also important to observe that pixel size, determined in large part by imaging optics, has generally been reduced to the 10-20 μm regime and will not shrink much farther, whereas microelectronics technology continues on an exponentially increasing circuit density curve. Furthermore, microlense technology can alleviate fill-factor constraints on unit cell electronics.

Proc. of 1991 Symposium on Advanced Image-Acquisition Technology (SAIT91),pp. 48-56, T. Ando, ed., Tokyo, Japan, November 1991

Thus, there is an increasing window of opportunity to apply parallel and massively parallel electronic circuitry on the focal plane to enhance imaging system performance.

Implementation of on-chip focal-plane image processing has drawbacks. Increasing circuit density and/or chip area generally decreases IC yield. For cooled sensors, power-intensive processing can significantly increase heat load. Preferred image acquisition device structures may also differ from image processing circuitry device requirements thus increasing fabrication complexity. Clocking signals used in signal processing circuit operations may inadvertently add noise and have other detrimental effects on image acquisition. Analog processing circuit design is challenging so that off-chip "brute force" digital approaches may be easier to achieve. One must therefore carefully choose applications of on-chip focal-plane image processing.

In this paper, architectures for on-chip focal-plane image processing are discussed first. Work reported to date is used to illustrate achitectural approaches. Further directions for focal-plane image processing are suggested at the conclusion of the paper.

Architectures

Architectural approaches for performing on-chip focal plane image processing depend upon the type of processing required and the pixel pitch. If the minimum feature size is L, a low density imager is defined as having a pixel pitch exceeding 50L. A high density imager has a pitch under 10L, with medium density imagers falling in-between. In general, the higher the density and size of the imager, the more difficult it is to achieve focal plane image processing.

High Density Imagers

High density imagers have the smallest unit cell size and, in general, there is no real-estate available within the unit cell for signal processing. Thus, on-chip image processing circuitry must be located beyond the region of detection and read-out. Some exceptions to this are to provide some signal conditioning through a buffer amplifier in the unit cell or to utilize the read-out circuitry in an unconventional manner to achieve filtering [1] or binning. The preferred architectural approach is to divide the IC into an image acquisition sector, and a processing sector, as shown in Fig. 1.

One of the simplest signal processing functions envisioned for integration on the image plane is A/D conversion. For machine vision applications, a few bits of resolution is often adequate, but for video applications, 8 to 10 bits would be required. For scientific imaging, 12 to 16 bits are required. In A/D conversion, there is often a trade-off between bit resolution and throughput, or power and throughput. Since 1 to 10 MHz conversion throughput is typically required, and power consumption on the image plane must be minimized, it might be expected

Proc. of 1991 Symposium on Advanced Image-Acquisition Technology (SAIT91),pp. 48-56, T. Ando, ed., Tokyo, Japan, November 1991

that 8 to 10 bit resolution is a reasonable goal for on-chip A/D conversion. Charge domain A/D conversion is one approach [2,3]. Switched capacitor-based conversion circuits have low power but may be too slow for video applications. Semi-parallel A/D conversion can compensate for slow conversion rate but is more suited for medium density imagers [4].

Image processing can involve a single pixel or multiple pixels but only a few image processing tasks involve one pixel at a time. One of these few is thresholding. In this case, a pixel signal level is compared to a reference level and either a logic 0 or 1 is produced depending on the relative magnitudes of these levels. Although this can simplify off-chip driving of the signal, thresholding is usually not used in practice until a number of other operations have taken place. Image half-toning usually involves a more sophisticated approach to thresholding to insure that the local average of the analog and binary images are equal and that objectionable artifacts are not introduced into the image [5,6]. Noise reduction by non-destructively resampling the signal can also be performed on chip [7] and involves only a single pixel at a time. Level shifting and gain adjustment (also known as non-uniformity compensation) are other examples of single pixel operations. However, the merits of performing non-uniformity compensation on-chip without performing other multi-pixel operations subsequently on-chip are unclear.

Most operations involve a local neighborhood of pixels. For example, edge detection requires comparison of a pixel value to that of its nearest neighbors. Many local neighborhood operations can be cast as a local convolution or weighting of the nearest neighbors over some window region, typically 3 x 3 or 5 x 5 in size. Hence, the values of the neighbors must be available simultaneously in order to perform the operation. Unfortunately, in conventional read-out architectures the serial output data stream retains horizontal neighbors in close proximity but vertical neighbors are separated in time by a one row delay. Circuitry to reconstruct the local neighborhood after read-out must be utilized prior to image processing.

Neighborhood reconstruction is most easily achieved through the use of a tapped on-chip delay line which provides for multi-row buffering [8]. Alternatively, in the parallel-to-serial multiplexer input process, copies of three adjacent rows may be generated if they can be non-destructively sensed [9]. In another scheme, a single row may be regenerated three times, with two of the copies being delayed by appropriate CCD circuitry. This has the advantage of avoiding the charge transfer efficiency and layout problems of the first approach, and the uniform non-destructive sensing problem of the second [10]. Neighborhood reconstruction can also be implemented by reading out the image in a non-raster scan mode so that local neighborhoods remain intact upon readout. For example, a 3x3 neighborhood read out in a single block has been used to improve image compressibility through differential encoding of nearest neighbors [11,12].

Proc. of 1991 Symposium on Advanced Image-Acquisition Technology (SAIT91),pp. 48-56, T. Ando, ed., Tokyo, Japan, November 1991

The pixel processor operates on the reconstructed neighborhood to generate the processed image. The operations can consist of linear and non-linear components. A convolver processor performs simple linear weighting of the neighborhood. More sophisticated processors would perform non-linear functions such as thresholding and conditional operations and might be digitally programmable. The high throughput requirements of high density imagers dictates that the processor be organized in a pipeline architecture.

The convolver processor weighting of the neighborhood may be realized in several ways. A transversal filter-like architecture employing split electrodes can be used to simultaneously sense, weight and sum at high speed, but with limited accuracy [13]. Alternatively, the signals may be sensed and regenerated . During regeneration, weighting can be provided through the use of a CCD fill-and-spill circuit with appropriate area ratios [14]. These two approaches provide for fixed-weight operation determined by mask layout. A programmable approach using a CCD multiplying D/A converter architecture would be of more general purpose [15,16].

In high-density imagers, where yield is already of concern, further integration of more sophisticated circuitry may be less attractive. Data transforms to improve compressibility of images for transmission might be implemented in the analog charge domain using CCD-like circuits. Frame-to-frame operations for compression, motion detection, tracking and event detection require the integration of frame memory and offer exciting directions for future research. However, implementation of these functions must be demonstrated off-chip before these operations will be accepted for on-chip integration.

Medium Density Imagers

Medium density arrays are defined to have detector pitches between 10L and 50L and a small amount of real-estate is available within the unit cell. Architectures for medium density arrays fall between those for high density imagers and those for low density imagers. Approaches described for high density imagers are applicable to medium density imagers. In addition, it may be possible to provide a degree of parallel processing in a medium density imager prior to serial multiplexing. In this case, real-estate located at the bottom of the parallel read-out multiplexer would be utilized for the processor array (see Fig. 2.) Ideally, each vertical column would have its own processor (tall and thin layout required) and processors would be able to communicate with nearest neighbor processors. The processors could be organized in a pipeline fashion within each column to maximize throughput. Thus, such an architecture is termed a pipelined vector processor[10]. Serial multiplexing of the processor array output would be readily realized. The processor throughput requirement is substantially reduced below that found in a high density imager for a given frame rate due to both the parallelism and likely reduction in the number of rows. Simplified processor design becomes possible due to the relaxation of the throughput requirements, though real-estate constraints become more critical. Fixed pattern (columnar) noise

Proc. of 1991 Symposium on Advanced Image-Acquisition Technology (SAIT91),pp. 48-56, T. Ando, ed., Tokyo, Japan, November 1991

introduced by processor-to-processor variation must be considered. For example, source-follower buffer amplifiers must utilize threshold voltage cancellation schemes, unlesss non-uniformity compensation in the processor array (necessary in any case for most IR detector applications) is applied. Semi-parallel A/D conversion (one converter per column) is an interesting candidate for this architecture [4]. A Laplacian-of-Gaussian on-chip focal-plane image processor chip recently used this architecture [17].

Low Density Imagers

Low density imagers are useful for machine vision, surveillance, and autonomous vehicles. The unit cell size is large to increase the number of photo-generated carriers through an increase in detector area. Typically, imager array sizes are 256 x 256 or smaller. As before, architectures suitable for higher density imagers are appropriate for lower density imagers. However, in the case of low density imagers, a spatially parallel architectural approach is also possible. In the spatially parallel architecture, there is one processor for each detector (see Fig. 3.) and the processors are interconnected in a way which reflects the spatial topology of the image. In the case of focal-plane image processing, spatially parallel architectures are naturally suited to the image plane, with processors communicating with their nearest neighbors.

The penalty for spatially parallel processing is the potential reduction of detector real-estate (fill factor). The use of the third dimension can alleviate this constraint. The detector region can be an amorphous silicon overlayer [18], recrystallized silicon [19] or a separate detector chip hybridized (bump-bonded) to the processor array. For applications requiring very high throughput and a large amount of processor real-estate, the "Z-plane" architecture [20] might be employed. In these cases, detector fill factor does not suffer.

The throughput of a spatially parallel architecture can be quite high due to the high degree of parallelism. For example, a 10 x 10 mm imager with a detector pitch of 150 μm could have approximately 4,000 pixels. Assuming 100 elemental operations per pixel, a serial processor operating at the rate of 1 μsec per operation would take approximately 400 msec to process the image, corresponding to a frame rate of 2.5 Hz. On the other hand, a spatially parallel architecure could process the image at a speed-up ratio of 4,000 corresponding to a frame rate of 10,000 Hz! In a pratical sense, 10,000 Hz is too high for most applications and the data read-out from the processor array would likely be a major bottleneck. A lower degree of parallelism can be traded for tighter pixel pitch with each processor serving multiple pixels.

Analog charge domain CCD-like circuits are well-suited for the circuit realization of such a discrete-time processor [21] and some recent progress has been made [22]. An array of 24 x 24 processors has been fabricated with each processor serving four pixels for a total imager size of 48

x 48. The image processor chip is approximately 1 cm in area with a 180 μm photodetector pitch. Each processor has circuitry for performing addition, subtraction, comparison, conditional differencing, short term memory, and communication with nearest neighbors, in addition to charge collection from the four photodiodes. Assuming 250 cycles/ elemental operation per pixel per frame and 0.40 μsec clock widths (25 clock cycles/elemental operation), a processed frame rate of 50 Hz can be achieved at a total power cost of under a few milliwatts. In principle, internal throughput exceeding 1000 frames per second (576 million operations/ second) can be achieved at a cost under 50 mW, but in practice the serial output multiplexer could not handle the corresponding serial data rate. Fully-parallel charge-domain processors have also been explored for performing stereo vision processing. [23] and for improving interferometric measurement [24].

Continuous time parallel processors for on-chip focal plane image processing are also receiving increased attention. Such "neural network" architectures are designed to mimic functions of the human retina [25,26]. To date, these systems have traded imager performance for massively-parallel computational throughput.

Scanned Imagers

Scanned imagers such as TDI (time delay and integration) multiplexers can also be considered for focal-plane image processing. The architectures described for medium density and high density imagers are valid for TDI imagers, depending on the detector pitch. In general, more real-estate is available for the processor since the array size (and chip size) is usually smaller, resulting in higher yield. However, for wide scanners, neighborhood reconstruction through the use of delay line could become much more susceptible to transfer efficiency problems. Other methods of neighborhood reconstruction might be more useful.

Future Directions

Biological models such as the human eye indicate that evolution favors focal-plane image processing for lower level tasks, with cognitive function reserved for larger, centralized processing. As microelectronics technology advances cause high density imagers to become reclassified as medium and low density imagers through the shrinking minimum feature size L, opportunitites for focal-plane image processing will dramatically increase. Commercial IC process advances, such as BiCMOS, coupled with X-Y readout unit cell architectures will also speed commercial application of focal-plane image processing for machine vision, surveillance and possibly home video. Scientific imaging systems will also benefit from these advances. It is inevitable that focal-plane image processing will be adopted for many future imaging systems.

Proc. of 1991 Symposium on Advanced Image-Acquisition Technology (SAIT91),pp. 48-56, T. Ando, ed., Tokyo, Japan, November 1991

Acknowledgments

The author gratefully acknowledges years of discussion with his graduate students, especially S. Kemeny, and E-S. Eid, and his colleagues at Columbia University and elsewhere. This paper was written at the Jet Propulsion Laboratory, California Institute of Technology, under a contract from the National Aeronauctics and Space Administration, Office of Aeronautics Exploration and Technology.

References

1. J.P. Sage, "A high speed analog two-dimensional gaussian image convolver", Opt. Soc. America Topical Meeting on Machine Vision, FD5/1-4, Incline Village, NV, March 1985.
2. E.R. Fossum, "Charge domain analog signal processing for detector arrays," Nuclear Instruments and Methods, vol. A275, pp. 530-535 (1989).
3. S.E. Kemeny, CCD focal-plane reorganization processors, Ph.D. Thesis, Columbia University 1991.
4. S. Mendis, R. Nixon and E.R. Fossum, unpublished
5. J. Jarvis, C. Judice, and W. Ninke, "A survey of techniques for display of continuous-tone pictures on bi-level displays," Computer Graphics and Image Processing 5, pp. 12-40 (1976).
6. E.-S. Eid and E.R. Fossum, "Design of a CCD focal-plane image half-toner," in Charge-Coupled Devices and Solid-State Sensors, M. Blouke, ed. Proc. SPIE 1242, pp. 126-132 (1990)
7. C.E. Chandler, R. A. Bredthauer, J.R. Janesuck, J. A. Westphal, and J.E. Gunn, "Sub-electron read noise charge-coupled devices," Proc. SPIE vol. 1242, pp. 238-251 (1990).
8. J.E. Hall, J.F. Breitzmann, M.M. Blouke, and J.T. Carlo, "A multiple output CCD imager for imag eprocessing applications," Proc. 1978 International Electron Devices Mtg., pp. 415-418, San Francisco, CA, Dec. 1978.
9. J.P. Vermeiren and G.J. Declerck, "A novel CCD output structure for picture processing applications," Proc. 1985 IEEE Custom Integrated Circuits Conf., pp. 128-131, New York, NY, May 1985.
10. E.R. Fossum, "Architectures for focal-plane image processing, Optical Engineering, vol. 28(8) pp. 865-871 (1989).
11. S.E. Kemeny, H. Meadows, and E.R. Fossum, "Design of a CCD focal-plane codec preprocessor for lossless image compression," in Charge-Coupled Devices and Solid-State Optical Sensors, Proc. SPIE vol. 1242, pp. 118-125 (1990).
12. S.E. Kemeny, H. Torbey, H.E. Meadows, R. Bredthauer, M. LaShell, and E.R. Fossum, "CCD focal-plane image reorganization processors for lossless image compression," to appear in IEEE J. Solid-State Circuits, March 1992.

13. S.D. Fouse, G.R. Nudd, and P.A. Nygaard, "Implementation of image preprocessing functions using CCD LSI circuits," Proc. SPIE vol. 225, pp. 118-130 (1980).
14. A.M. Chiang, "A video-rate CCD two-dimensional cosine transform processor," Proc. SPIE vol. 845, pp. 2-5 (1987).
15. A.M. Chiang, and B.E. Burke, "A high speed digitally programmable CCD transversal filter," IEEE J. Solid-State Circuits, vol. SC-18(6) pp. 746-753 (1983).
16. E.R. Fossum, S.E. Kemeny, R. Bredthauer, and M. LaShell, "Digitally programmable gain control circuit for charge-domain signal processing," IEEE J. Solid-State Circuits, vol. SC-26(4) pp. 683-686 (1991).
17. W. Yang and A.Chiang, "A full fill-factor CCD imager with integrated signal processors," Proc. 1990 IEEE International Solid-State Circuits Conference, pp. 218-219 (1990).
18. R. Miyagawa, et al., Proc. 1988 IEEE International Electron Devices Meeting, pp. 74-77, San Francisco, CA, Dec. 1988.
19. 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).
20. J.C. Carson, "Applications of advanced 'Z' technology focal plane architectures," Proc. SPIE vol. 930, pp. 164-182 (1988).
21. E.R. Fossum, "Charge-coupled computing for focal-plane image preprocessing," Optical Engineering, vol. 26(9) pp. 916-922 (1987)
22. E-S. Eid and E.R. Fossum, "Real-time focal-plane array image processor," in Automated Inspection and High Speed Vision Architectures III, Proc. SPIE vol 1197, paper 1 (1989)
23. J. Hakkarainen,J. Little, H-S. Lee, and J. Wyatt, "Interaction of algorithm and implementation for analog VLSI stero vision," Proc. SPIE vol. 1473, paper 20 (1991).
24. H.W. Tomlinson, M.P. Weir, G.J. Michon, G.E. Possin, J. Chovan, "Focal-plane processing algorithm and architecture for laser speckle interferometry," Proc. SPIE vol. 1541 paper 22 (1991).
25. C. Mead, M. Ismail, Analog VLSI and Neural Systems, Addison Wesley, Reading MA (1989).
26. B. Mathur, S-C. Liu, and H. Wang, "Analog neural networks for focal-plane image processing," Proc. SPIE vol. 1242 pp. 141-151 (1990).

Dr. Eric R. Fossum
MS 300-315
Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, CA 91109 USA
Voice: (818) 354-3128
Fax : (818) 393-0045
email: efossum@nasamail.jpl.nasa.gov

Proc. of 1991 Symposium on Advanced Image-Acquisition Technology (SAIT91),pp. 48-56, T. Ando, ed., Tokyo, Japan, November 1991

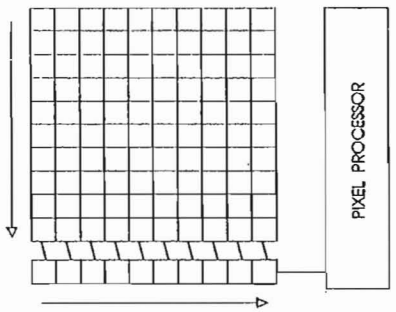

Fig. 1. Architecture for on-chip focal-plane image processor for high density imager. Pixel processor is located after serial multiplexer. Architecture applies to both CCD and switched FET X-Y readout approaches.

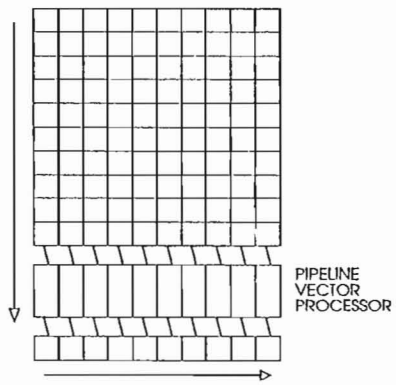

Fig. 2. Semi-parallel architecture for medium density imager. Each column has its own processor.

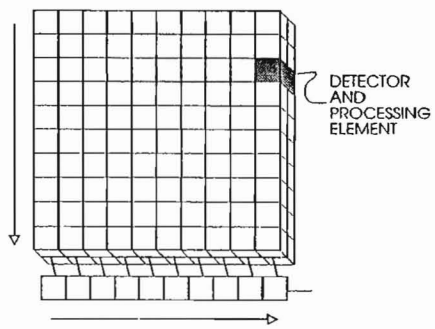

Fig. 3. Spatially-parallel architecture for low density imager. Each pixel (or sub-group of pixels) has its own processing element. Processing elements are interconnected and perform single-instruction, multiple-data (SIMD) operations.

Proc. of 1991 Symposium on Advanced Image-Acquisition Technology (SAIT91),pp. 48-56, T. Ando, ed., Tokyo, Japan, November 1991