

A Reprint from the

PROCEEDINGS

Of SPIE - The International Society for Optical Engineering



Volume 1071

Optical Sensors and Electronic Photography

16-18 January 1989
Los Angeles, California

CCD architectures for focal-plane image processing

Eric R. Fossum

Department of Electrical Engineering, Columbia University
New York, New York 10027

CCD Architectures for Focal-Plane Image Processing

Eric R. Fossum

Department of Electrical Engineering
Columbia University
New York, New York 10027

ABSTRACT

Architectures for focal-plane image processing using CCD circuits are discussed. The choice of architecture depends on imager density and required throughput. High-density imagers require the reconstruction of local neighborhoods prior to image processing. Lower density imagers can utilize spatial parallelism to improve throughput. The use of three-dimensional structures can provide additional real-estate for processing circuitry.

INTRODUCTION

Solid-state imaging devices have evolved rapidly in the past few years. Television-quality imagers capture pictures containing several hundred thousand pixels at frame rates of 30 Hz, corresponding to data rates on the order of 10 megapixels per second. Applications such as robotic manufacturing, autonomously guided vehicles, and teleconferencing require the real-time processing of these images. The system throughput requirement can thus exceed 10 Bops (billion operations per second). Although digital integrated circuits may be pushed to achieve such high throughput¹⁻³, these circuits generally consume between 500 mW and 1,500 mW of power and are restricted in function. Furthermore, they require the A/D conversion of the data prior to processing, which can be another major source of power consumption.

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. The charge-coupled device (CCD) technology is well-suited for these analog circuits⁴ and some CCD image processing circuits have been demonstrated or proposed⁵⁻⁸. This paper discusses the architectural aspects of designing CCD image processing circuits for focal-plane image processing. In this approach, image acquisition and image processing circuitry are integrated on the image plane.

ISSUES

Focal-plane image processing can imply a wide degree of signal processing functions. Edge detection and image coding require a high degree of processing whereas integration of detector buffer/amplifiers is less demanding. The possible extent of processing is primarily dictated by the available chip real-estate, which in turn, depends on the detector array density and size.

One major reason for considering focal-plane image processing is that it avoids the introduction of noise and distortion during off-chip driving of the multiplexed output. A second reason is that on-chip processing, such as quantization, can reduce the bandwidth of the signal driven off chip. Third, for parallel, digital architectures, the reformatting and channeling of data from the serial imager output to the parallel processor array can represent significant overhead. Spatially parallel on-chip processing (discussed below) can avoid this bottleneck. Finally, the use of analog charge domain circuitry, as discussed above, has additional advantages.

Focal-plane image processing has disadvantages as well. The real-estate is severely constrained. Expending additional power in an imager designed for cooled detection may be undesirable, and circuit design is challenging. Off-chip "brute-force" digital approaches may be easier to achieve. However, with careful design, focal-plane image processing has the potential to improve system performance and reliability while reducing complexity and power consumption.

Imagers may be classified into high density, medium, and low density according to the pixel pitch. If the minimum device feature size is L , a low density imager has a pixel pitch exceeding $50L$. A high density imager has a pixel pitch under $10L$, with medium density imagers falling in between. The higher the density and size of the imager, the more difficult it is to achieve focal-plane image processing.

A longer version of this paper has been submitted for publication in Optical Engineering.

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. Two exceptions to this are to provide some signal conditioning through a buffer amplifier in the unit cell⁹ or to utilize the read-out circuitry in an unconventional manner¹⁰.

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 that 8 to 10 bit resolution is a reasonable goal for on-chip A/D conversion. Switched capacitor based conversion circuits have low power but may be too slow for video applications. A promising possibility is the use of CCD-based circuits for A/D conversion, but testing of experimental circuits has not yet been completed¹¹.

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. 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 line¹² which provides for multi-row buffering. Alternatively, in the parallel-to-serial multiplexer input process, copies of three adjacent rows may be generated if they can be non-destructively sensed¹³. In another scheme, a single row may be regenerated three times, with two of the copies being delayed by appropriate CCD circuitry¹⁴. 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.

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 either processor be organized as 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 accuracy¹⁵. 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 ratios¹⁶. 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^{11,17}.

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^{16,18}. 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, CCD 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 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. Ideally, each vertical column would have its own analog processor (tall and thin layout required) and processors would be able to communicate with nearest neighbor processors. The processors would be naturally organized in a pipeline fashion within each column to maximize throughput. Thus, such an architecture is termed a pipelined vector processor. 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 analog processor design becomes possible due to the relaxation of the throughput requirements, though real-estate constraints become more critical. Fixed pattern (columnar) noise introduced by processor-to-processor variations must be considered. For example, source-follower buffer amplifiers must utilize threshold voltage cancellation schemes, unless non-uniformity compensation in the processor array (necessary in any case for most IR detector applications) is applied.

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. Typical imager array sizes are 32 x 32 or 64 x 64. 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. 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. The use of the third dimension can alleviate this constraint. The detector region can be an amorphous silicon overlayer¹⁹ 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 might be employed²⁰. 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 mm² imager with a detector pitch of 150 um could have approximately 4,000 pixels. Assuming 100 elemental operations per pixel, a serial processor operating at the rate of 1 usec 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 architecture could process the image at a speed-up ratio of 4,000 corresponding to a frame rate of 10,000 Hz! In a practical 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 processor⁴ and some recent progress has been made⁷. 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 cm² in area with a 180 um 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 elemental operations per pixel per frame and 0.40 usec 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. Testing of this chip is currently underway.

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 lines could become much more susceptible to transfer efficiency problems. Other methods of neighborhood reconstruction might be more useful.

CONCLUSIONS

Focal plane image processing, particularly in the analog charge domain, shows promise for reducing severe throughput, power, and real-estate problems associated with current digital off-chip technology. It can be anticipated that focal-plane image processing in solid-state imaging systems will develop rapidly in the next few years. A marriage of analog CCD and digital CMOS technology will provide interesting possibilities. Since system input is analog in nature and high level processing digital, systems will need to combine the best features of analog and digital processing circuitry.

ACKNOWLEDGMENTS

The author gratefully acknowledges thoughtful discussions with his graduate research assistants and faculty colleagues at Columbia University. Portions of this work were supported by an NSF Presidential Young Investigator Award, the NSF Center for Telecommunications Research at Columbia, an IBM Faculty Development Award and a grant from Analog Devices, Inc.

REFERENCES

1. P.A. Ruetz and R.W. Brodersen, "Architectures and Design Techniques for Real-Time Image Processing IC's", *J. Solid-State Circuits*, SC-22(2), 233-250 (1987).
2. H. Yamada, K. Hasegawa, T. Mori, H. Sakai, and K. Aono, "A Microprogrammable Real-Time Image Processor", *J. Solid-State Circuits*, SC-23(1), 216-223 (1988).
3. T. Denayer, E. Vanzieleghem, and P.G.A. Jespers, "A Class of Multiprocessors for Real-Time Image and Multidimensional Signal Processing", *J. Solid-State Circuits*, SC-23(3), 630-638 (1988).
4. E.R. Fossum, "Charge-Coupled Computing for Focal Plane Image Preprocessing", *Optical Eng.*, 26(9), 916-922 (1987).
5. G.R. Nudd, P.A. Nygaard, G.D. Thurmond, and S.D. Fouse, "A CCD Image Processor for Smart Sensor Applications", *Proc. SPIE* 155, 15-22 (1978).
6. J.D. Joseph, P.C.T. Roberts, J.A. Hoschette, B.R. Hanzal, and J.C. Schwanebeck, "A CCD-based parallel analog processor", *Proc. SPIE* 501, 238-241 (1984).
7. E-S. Eid and E.R. Fossum, "CCD Focal Plane Array Analog Image Processor", *Proc. SPIE* 977, paper 32, (1988).
8. P.R. Beaudet, "CCD Focal Plane Convolver (Smart Eyeball)", *Opt. Soc. America Topical Meeting on Machine Vision*, FD6/1-4, Incline Village, NV, March 1985.
9. H. Yamashita, et al., *Proc. 1988 IEEE International Electron Devices Meeting*, 78-81, San Francisco, CA, Dec. 1988.
10. 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.
11. S.E. Kemeny and E.R. Fossum, "CCD A/D and D/A Converters", in progress.
12. J.E. Hall, J.F. Breitzmann, M.M. Blouke, and J.T. Carlo, "A Multiple Output CCD Imager for Image Processing Applications", *Proc. 1978 IEEE International Electron Devices Meeting*, 415-418, San Francisco, CA, Dec. 1978.
13. J.P. Vermeiren and G.J. Declerck, "A Novel CCD Output Structure for Picture Processing Applications", *Proc. 1985 IEEE Custom Integrated Circuits Conf.*, 128-131, New York, NY, May 1985.
14. E.R. Fossum, unpublished.
15. S.D. Fouse, G.R. Nudd, and P.A. Nygaard, "Implementation of image preprocessing functions using CCD LSI circuits", *Proc. SPIE* 225, 118-130 (1980).
16. A.M. Chiang, "A Video-Rate CCD Two Dimensional Cosine Transform Processor", *Proc. SPIE* 845, 2-5, (1987)
17. A.M. Chiang and B.E. Burke, "A High Speed Digitally Programmable CCD Transversal Filter", *IEEE J. Solid-State Circuits*, SC-18(6), 746-753 (1983).
18. H. Torbey and H. Meadows, "System and Method for Lossless Hierarchical Image Coding/Decoding", *CTR Tech. Rep. 094/88-22*, Columbia University (1988).
19. R. Miyagawa, et al., *Proc. 1988 IEEE International Electron Devices Meeting*, 74-77, San Francisco, CA, Dec. 1988.
20. J.C. Carson, "Applications of advanced "Z" technology focal plane architectures", *Proc. SPIE* 930, paper 18 (1987).