

New Massively Parallel Technique for Global Operations in Embedded Imagers

Vladimir Brajovic and Takeo Kanade

The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

1. Introduction

The need for low-latency vision systems is quickly growing. Existing vision systems, however, still do not meet this need. For example, pipelined dedicated vision hardware updates its output at standard video rate (30 frames per second) but the latency incurred is usually several frame times. Standard video cameras provide images at standard video rate of 30 frames per second. For vision application requiring less than 1/30 of a second latency, standard video rate cameras can no longer fulfill the need — it is often too late by the time the system has received an image. It is clear that an alternative approach is needed.

It has been noted in the past that spatial properties of sensed images can be exploited in VLSI computational sensors and focal plane structures. A number of designs have emerged implementing local operations on a single light sensitive VLSI chip [1] [5] [6] [9]. The local operations compute on operands within a small spatial/temporal neighborhood of data; typically represented by invariant FIR filtering, results of such operations are preprocessed images. Therefore, a large quantity of image data must be still readout and further reasoned upon before making a final decision about the scene.

Global operations, on the other hand, provide fewer entities for representation of a scene being sensed. If computed at the point of sensing, these entities could be routed from a computational sensor through a few output pins with very little latency. Often this will be sufficient for rapid decision making without requiring the image to be readout. Global operations, however, need to gather and process information over the entire set of data. This global

exchange of data among a large number of processors/sites quickly saturates device's I/O and adversely affects computing efficiency in parallel systems — parallel digital computers and computational sensors alike. It is not surprising that there are only a few computational sensors for global operations, all with modest capability and/or low resolution [4] [7].

This work proposes an efficient massively parallel technique suitable for *global computation* over large groups of fine-grained data. Typical examples of such global operations on images include histogram computation, sorting, region segmentation and labeling.

2. Method and Architecture

The proposed technique for global operations in embedded imagers comprises the following steps. The input data are fetched in parallel by focusing an image on the array of sensor-processor cells. An architecture of such an array is shown in Figure 1. A local processor in each cell performs at least one operation. The *instant* when this operation is performed (i.e. triggered) is determined by the intensity of the radiation received at the particular cell. A single global processor supervises and services the array of cells. Since local processors trigger at times determined by the magnitude of their input operands, the global processor serves only few local processors at a time.

3. Sorting Chip

By using the proposed technique we have developed a sorting computational sensor — a sensor which is able to sort all pixels of an input image by their intensities.

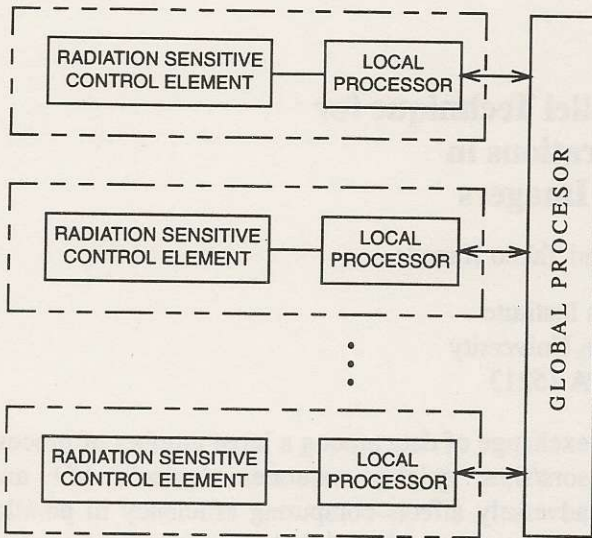


Figure 1: An architecture of computational sensor for global operations.

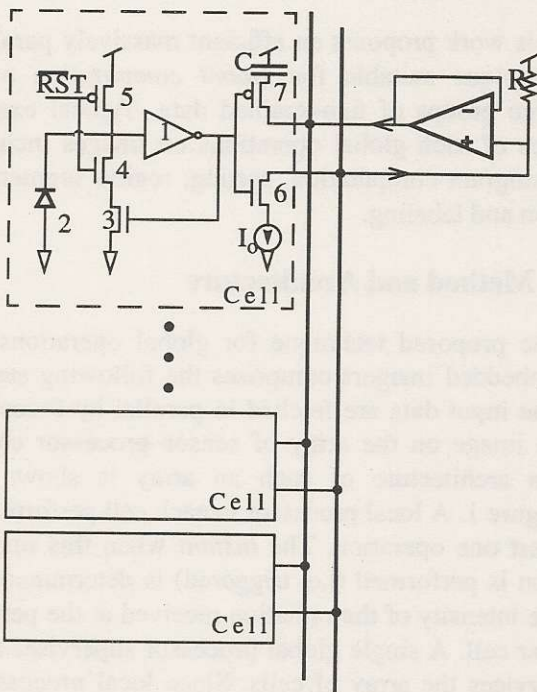


Figure 2: Schematic diagram of a sorting computational sensor

The schematic diagram of the sorting computational sensor is shown in Figure 2. Each local processor performs two function:

1. remember data supplied to them by the global processor within track-and-hold (T/H)

circuit comprised of capacitor C and switch 7; and,

2. send a current signal (I_o) to the global processor via switch 6.

The remaining portion of the cell comprises the radiation sensitive control circuit. It controls the instant when the local processor performs its function(s).

The cells share two common wires. These common wires and circuitry connected to them can be considered a global processor. The global processor sums the received currents and articulates an appropriate voltage waveform for storage in local processors.

3.1. Operation of the Circuitry

The operation of the sorting computational sensor with four cells has been simulated and is shown in Figure 3.

A photodiode (2) operating in the photon flux integrating mode [8] detects the radiation. In this mode of operation the capacitance of the diode is charged to a high potential and left to float. Since the capacitance is discharged by the photocurrent, the voltage decreases approximately linearly at a rate proportional to the amount of light impinging on the diode (Figure 3, top graph).

The diode voltage is monitored by a CMOS inverter (1). Once the diode voltage falls to the threshold of the inverter, the inverter's output changes state from low to high (Figure 3, second graph). A switch (3) is included to force rapid latching action.

The output of the inverter represents a control signal produced by the radiation sensitive control circuit. It controls the *instant* when the T/H memorizes supplied voltage. It also controls the instants when the current I_o is supplied by the cell. This is achieved by two complementary switches, (6) and (7): one turns on the internal current source (I_o), and the other disconnects the internal storage capacitor (C) from the global input wire.

Currents from all cells are summed up on a single output wire. Therefore, this wire functions as a global summer. The voltage on the resistor R (Figure 3, third graph) represents the index of a cell that is changing state and is supplied to the




Figure 3: Sorting computational sensor: simulated operation

global input wire. The capacitor within each cell follows this voltage until it is disconnected, at which point a capacitor C retains the index of the cell (Figure 3, bottom graph). The bottom graph shows that the cell with the highest intensity input has received the highest "index", the next cell one "index" lower, and so on.

The chip computes several important properties about the image focused thereon. First, the time when a cell triggers is approximately inversely proportional to the input radiation received. Second, by summing up the currents I_0 from all the local processors the global processor knows at each given time how many cells have been triggered. This time waveform is closely related to a cumulative histogram of the input image. The time derivative of this signal is related to a histogram of the input image. This is one global property of the input image that is reported by the chip with very low latency.

3.2. VLSI Realization and Experiments

A 21 x 26 cell sorting sensor has been built in 2 μ CMOS technology. The size of each cell is 76 μ by 90 μ . A micrograph of 2mm square die is

shown in Figure 4.




Figure 4: Micrograph of the sorting chip.

An image was focused onto the silicon. Several experiments were performed under varying conditions. The cumulative histogram waveform as well as the indices were digitized with 12 bits of resolution. In order to facilitate hard copy reproduction the 26x21 images obtained by the sorting chip are interpolated and magnified by the factor of 2.

A few images of an office scene were taken. This office scene seen by a commercial CCD camera is shown in Figure 7. The illumination was standard office lightning coming from the ceiling right above the operator.

A portion of the scene was imaged by the sorting chip with operator in (Figure 5) and out (Figure 6) of the field of view. Cumulative histogram of *the scene* provided by the chip is shown in the top graph for each image. Scene 1 contains more dark regions than Scene 2, since the moderately bright wall in the background is replaced by the dark regions of the operator in the partial shadow. Therefore, the chip takes longer to compute Scene 1 than Scene 2, but the dynamic range of the output indices is maintained. The total time shown on the sample axis is about 200ms.

A histogram of the computed indices (i.e. output *image*) is shown in the bottom graph for each scene. Most of the pixels appeared to have different input intensities and, therefore, received different indi-




Image of indices




Figure 5: Scene 1 imaged by the sorting computational sensor.




Image of indices:





Figure 6: Scene 2 imaged by the sorting computational sensor.

ces. Occasionally as many as 3 pixels were assigned the same index. Overall the histogram of indices is uniform, indicating that the sorting chip




Figure 7: An office scene as seen by a commercial CCD camera.

has performed well.

There is a total of 546 pixels in this prototype, and most of them received different indices. This means that without special considerations as to the illumination conditions, low-noise circuit design and temperature and dark current control, our lab prototype readily provided images with more than 9 bits of resolution. Under the same conditions a commercial CCD camera performed poorly. (We applied Gamma of 2.5 before including the image in this document to facilitate hard copy reproduction, but generally, contrast was poor.)

4. Applications

The data that ought to be stored in the local processors or the sorting chip are provided by the global processor. The shape of the supplied data stream defines mapping from input intensities to output intensities. There are numerous operations/mappings on input images that can be performed with the same chip, making it flexible and programmable. Examples of such operations include histogram computation and equalization, arbitrary point-to-point mapping, region segmentation and adaptive dynamic range imaging. In fact, the chip provides all the information necessary to perform any mapping during the readout: knowing the shape of the cumulative histogram as well as the indices, the input (i.e. natural) scene intensities can be mapped in number of ways.

4.1. Histogram Equalization

When the voltage of the cumulative histogram (computed by the chip itself) is supplied to the

local processors, the generated image is a histogram equalized version of the input image [2]. This is the basic mode of operation for the sorting chip and is illustrated in the previous section.

4.2. Linear Imaging

When the waveform supplied to the input wire is inversely proportional to time, the values stored in the capacitors are proportional to the input intensity, implementing a linear camera. The results of such mapping for Scene 1 are shown in Figure 8.

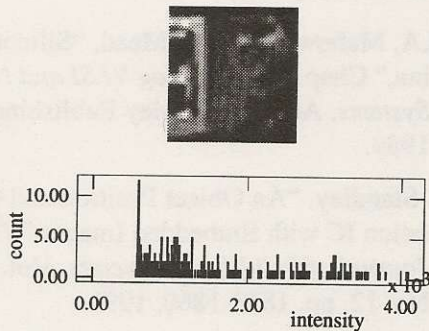


Figure 8: Sorting computational sensor as a linear imager: Image and its histogram.

As expected, the result is similar to the image obtained by the linear CCD imager (Figure 7).

4.3. Image Segmentation

The cumulative histogram of an image can be used to segment an image into regions. Pixels from a single region often have pixels of similar intensity that appear as clusters in the image histogram [2]. The values to be stored in the cells can be generated to correspond to the "label" of each such region. The global processor can perform this labeling by updating the supplied value when the transition between the clusters in the (cumulative) histogram is detected. An example of segmentation is shown in Figure 10b in which the bright background is labeled as zero.

4.4. Scene Change Detection

Analyzing the change in the histogram pattern is a basic technique to detect a scene change. The sorting computational sensor computes the cumulative histogram at real-time and can be used for low-latency scene discrimination/surveillance without requiring the image to be read out. The cumulative histograms for Scene 1 and Scene 2 are shown

together in Figure 9.





Figure 9: Cumulative histograms for Scene 1 and Scene 2 computed by the sorting chip.

4.5. Adaptive Dynamic Range Imaging

In many applications, for example, the illumination of the scene which is directly exposed to the sunlight is several hundreds times greater than the illumination for the surfaces in a shadow. During a single frame, both high and low intensity pixels can be mapped to the same output range. This is possible because the cells detecting high intensities "memorize" the supplied data before the low intensity cells.



a)



b)

c)

Figure 10: Imaging back lit objects with the sorting chip: a) the scene imaged with a commercial CCD camera; b) image of indices detected by the sorting sensor; c) segmented and enhanced dynamic range image obtained by the sorting sensor.

We applied this concept to back illuminated objects/scenes shown in Figure 10. As expected, the sorting sensor has adjusted its dynamic range and resolved the detail in the dark region as well as the bright (Figure 10a.) Since all 546 indices are competing to be displayed within 256 levels allowed for postscript images, one enhancement mapping for purpose of human viewing is to amplify the dark pixels. The result is shown in Figure 10b. The same method obviously can be applied to the image obtained with a standard CCD camera (Figure 10c.). However, we can only expect to have the face resolved with 2-3 bits, since the dark region of the image is quantized and digitized within a few levels. Theoretically, the sorting computational sensor allocates as many levels as there are pixels within the dark region, or the entire image for that matter.

5. Conclusion

It is essential to view a vision computational sensor as a massively-parallel computational engine which makes global or overall decisions about the sensed scene and reports such a decision on few output pins of the chip with low latency.

Presented is the technique which enables embedded imagers to perform global operations on massive amount of input image data. The power of the technique is illustrated with a VLSI implementation of sorting — an operation still challenging in computer science when performed on large groups of data [3]. This work shows that if an appropriate relationship between circuitry, algorithms and applications is maintained, VLSI vision computational sensors of appreciable spatial resolution are capable of performing such global operations.

References:

- [1] W. Bair, and C. Koch. "An Analog VLSI Chip for Finding Edges from Zero-crossings." In: *Advances in Neural Information Processing Systems*, Lippman, R., Moody, J., and Touretzky, D.S., eds., Vol. 3, pp. 399-405, Morgan Kaufmann, San Mateo, CA, 1991.B.
- [2] Ballard, D.H. and C.M. Brown, *Computer Vision*, Prentice-Hall, 1982.
- [3] Brelloch, G.E. et al. "A Comparison of Sorting Algorithms for the Connection Machine CM-2", SPAA, 1991.
- [4] DeWeerth, S.P. and C.A. Mead, "A Two-Dimensional Visual Tracking Array," *Proc. of the 5th MIT Conf. on Adv. Research in VLSI*, pp. 259-275, March 1988.
- [5] J.G. Harris, C. Koch, and J. Luo, "A two-dimensional analog VLSI circuit for detecting discontinuities", *Science*, Vol. 248, pp 1209-1211, 1990.
- [6] M.A. Mahowald and C. Mead, "Silicon Retina," Chap. 15 of *Analog VLSI and Neural Systems*, Addison-Wesley Publishing Co., 1989.
- [7] D. Standley, "An Object Position and Orientation IC with Embedded Imager," *IEEE Journal of Solid-State Circuits*, Vol. 26, No. 12, pp. 1853-1860, 1991.
- [8] Weckler, G.P. "Operation of p-n Junction Photodetectors in a Photon Flux Integrating Mode," *IEEE Jour. of Solid-State Circuits*, pp. 65-73, Vol. sc-2, No. 3, September, 1967
- [9] P.C. Yu, S.J. Decker, H.S. Lee, C.G. Sodini and J.L. Wyatt, Jr., "CMOS Resistive Fuses for Image Smoothing and Segmentation," *IEEE Journal of Solid-State Circuits*, Vol. 27, No. 4, pp. 545-553, April 1992.