Journal of VLSI Signal Processing 41, 65–80, 2005
© 2005 Springer Science + Business Media, Inc. Manufactured in The Netherlands.

Efficiency Analysis for a Mixed-Signal Focal Plane Processing Architecture

WILLIAM H. ROBINSON* AND D. SCOTT WILLS
Portable Image Computation Architectures (PICA) Laboratory, Georgia Institute of Technology, School of Electrical and Computer Engineering, Atlanta, GA 30332-0250, USA

Received July 2003; Revised January 2004; Accepted March 2004

Abstract. Monolithic integration of photodetectors, analog-to-digital converters, data storage, and digital processing can improve both the performance and the efficiency of future portable image products. However, digitizing and processing a pixel at the detection site presents the design challenge to deliver a system with the required performance at the lowest cost, not just a system with the highest performance. This paper analyzes the area-time efficiency, the area efficiency, and the energy efficiency of a mixed-signal, SIMD focal plane processing architecture that executes front-end image applications with neighborhood processing. Implementations of the focal plane architecture achieve up to 81x higher area efficiency and up to 11x higher energy efficiency when compared to traditional TI DSP chips. Higher efficiency ratings are required to maintain portability while addressing technology limitations such as interconnect wiring density, heat extraction, and battery life. Systems can be implemented with a less expensive fabrication technology by increasing the number of pixels per processing element (PPE).

Keywords: focal plane processing, digital pixel, single-instruction multiple-data (SIMD), area efficiency, energy efficiency, area-time efficiency

1. Introduction

The demand for mobile productivity has led to incorporating embedded systems into handheld devices for real-time applications. A wide range of products, such as cellular phones, video and still cameras, and portable data assistants (PDAs), deliver multimedia processing by including inexpensive imaging chips. However, this processing creates a challenging design problem [1] and requires a change in paradigm to accommodate processing requirements [2]. Future portable imaging products will benefit from the monolithic integration of photodetectors, analog-to-digital converters, digital processing, and data storage to improve their performance, efficiency, and cost. A typical system-on-a-chip digital signal processing (DSP) architecture, shown in Fig. 1(a), assigns an entire image to a single processing core. This architecture is designed to span all stages of image processing. However, the DSP architecture uses a significant amount of area and energy resources to perform basic image enhancement and image analysis applications. For example, the digital image signal multiprocessor [3] reports 53% of its execution time for preprocessing and 47% for feature extraction. Preprocessing includes common tasks, such as noise reduction, smoothing, and segmentation, which are characterized by high spatial locality. Processing in the focal plane can more efficiently handle these basic tasks, leaving the more complex applications to the DSP architecture. A focal plane architecture, shown in Fig. 1(b), can be built by integrating analog-to-digital conversion (ADC) and digital processing at each detector site in the focal plane array.

Incorporating analog and digital components to process pixels at the detection site creates a challenging design problem [4–6]. The goal is not strictly building the system with the highest performance, but delivering

*Currently affiliated with the Department of Electrical Engineering and Computer Science at Vanderbilt University.

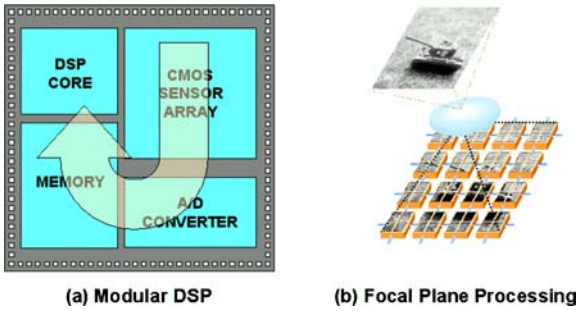

Figure 1. Architectural paradigms for image processing.

the system with the required performance at the lowest cost. The architecture must consider key performance and cost metrics, such as execution time, system throughput, chip area, power consumption, and area-time efficiency. However, in the system-on-a-chip design, two figures of merit are also important, namely area efficiency and energy efficiency [7]. Characterizing the targeted application workload leads to reasonable design choices for register file size, datapath width, and functional units to efficiently utilize silicon area and power consumption [8, 9]. In addition, higher ratings in area efficiency (system throughput per unit area) and energy efficiency (system throughput per unit power) are required because of technology limitations such as interconnect wiring density and heat extraction. The choice of both fabrication technology and pixels per processing element impact the system design.

This paper presents the area-time efficiency, the area efficiency, and the energy efficiency from analyzing a system-on-a-chip mixed-signal focal plane processing architecture. These results help quantify the system-level design issues to monolithically integrate photodetectors, analog-to-digital converters, data storage, and digital processing into focal plane architectures. Different implementations are used that vary the number of pixels per processing element (PPE). A focal plane architectural simulator provides the execution time and sustained throughput for the application suite. Component area models project the chip size for the design. A Technology Scenario Analyzer (TeSA) projects power consumption using different fabrication technologies. Because excessive die sizes would limit the realization of these systems, guidelines for affordable manufacturing are used from ITRS specifications [10, 11]. Battery life is another vital concern. Therefore, maximum power for each technology is constrained to ITRS specifications for portable battery operation, which is generally less than 3 Watts [10, 11]. Area-time efficiency, area efficiency, and energy efficiency are calculated using the appropriate projected values.

Despite a significant difference in clock frequency, implementations of the focal plane architecture perform well compared to a traditional DSP architecture. The sustained throughput achieved by the focal plane architecture implementations operating at 10 MHz exceeds the values reported in TI DSP chip specifications (Table 1). For area-time efficiency, the 16 PPE implementation is comparable to the TI TMS320C6211B chip, while the 1 PPE and the 4 PPE implementations show some improvement versus the TI TMS320C6411 chip.

Implementations of the focal plane architecture have higher energy efficiency and area efficiency when compared to traditional TI DSP chips. Energy efficiency increases by an average factor of 11x compared to the TMS320C6211B and by an average factor of 2.9x compared to the TMS320C6411. The energy efficiency demonstrates the potential for extended battery life in a portable device. The area efficiency is dramatically improved as well using the focal plane architecture (Table 2). The area efficiency for the focal plane architecture implementations includes the data acquisition and analog-to-digital conversion circuitry, which is not available with the DSP chips. Although the 1 PPE implementation provides the most efficient architecture, it generally exceeds the constraints for chip size and power consumed. A 4 PPE or 16 PPE with a less aggressive technology may provide a more cost-effective

Table 1. Increase in sustained throughput using focal plane architecture implementations.

	TMS320C6211B (150 MHz)	TMS320C6411 (300 MHz)
1 PPE	130×	65×
4 PPE	32×	16×
16 PPE	10×	5×

Table 2. Increase in area efficiency using focal plane architecture implementations.

	TMS320C6211B (180 nm technology)	TMS320C6411 (120 nm technology)
1 PPE	81×	56×
4 PPE	47×	27×
16 PPE	20×	10×

solution. In addition, detector area dominates as feature size shrinks. Therefore, denser detector technology could substantially reduce chip area. Another potential technique to reduce the analog component area is to share the analog-to-digital conversion circuitry [12].

The organization of this paper is as follows. Section 2 provides the background for architectures used by image processing systems. Section 3 describes the technique to analyze the performance of the focal plane processor array on the targeted application suite. Section 4 discusses the area models for the processing element components and presents the projected area for the focal plane processor array. Section 5 presents the projected power consumption for the focal plane processor array. Section 6 presents the analysis of resource efficiency for design implementations and provides a comparison to a traditional DSP architecture. Section 7 concludes the discussion on efficiency analysis for a mixed-signal focal plane architecture. Finally, Section 8 presents directions for future development.

2. Background

Numerous architectures have been developed for image processing systems [13, 14]. Efficient handling of the two-dimensional image data is a common issue among these designs. A natural solution to this issue is to process the spatially parallel data within the image plane [15]. Processing capability can be integrated directly into the focal plane. A logical approach for focal plane processing incorporates analog computational ability into the sensor device. Analog design methods, such as the Silicon Retina, have been presented to implement focal plane arrays [16]. This has led to the development of neuromorphic vision sensors for early image processing [17]. Another technique includes utilizing cellular neural networks for parallel processing [18]. This led to the development of the CNN Universal Chip [19]. However, the advantages of analog techniques, such as low power and small area, break down as CMOS technology scales [20]. In addition, analog architectures are typically non-programmable, requiring multiple designs to implement different functionality.

A programmable digital pixel is formed by monolithically incorporating the data acquisition device, analog-to-digital conversion (ADC) circuitry, data storage, and digital processing circuitry within a processing element (PE). A block diagram of this processing tile [20] is shown in Fig. 2. This tile element is replicated to form a focal plane array imager with integrated analog-to-digital conversion, nearest-neighbor communication, and SIMD processing. This section discusses the functionality of data acquisition, analog-to-digital conversion, data storage, and digital processing.

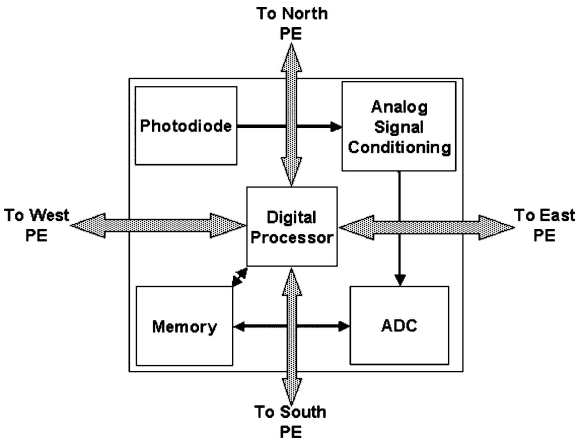

Figure 2. Programmable digital pixel.

2.1. Data Acquisition

The first stage performed by the digital pixel is the acquisition of the analog light intensity of an image. The relative response of the human eye correlates to the major color bands of the visible spectrum, which span from approximately 0.4 μm to 0.7 μm [21]. Therefore, photodetector devices used in the digital pixel should provide reasonable performance across these wavelengths. Photodetectors absorb photons with energy greater than the material bandgap to generate a current proportional to the number of optically generated electron-hole pairs [22]. The quantum efficiency of the detector is the number of electron-hole pairs generated per incident photon for a given wavelength [22]. The spectral response, or responsivity, is the quantum efficiency over a range of light wavelengths [23]. Photodetectors made from silicon can detect light wavelengths up to approximately 1.1 μm due to its bandgap energy of 1.12eV and provide reasonable performance over the visible spectrum (Fig. 3).

Implementing photodetectors in a standard CMOS process enables the integration of other processing components at the detection site. Research efforts have been made to develop and evaluate these CMOS image sensors [24–26]. A noted development is the active

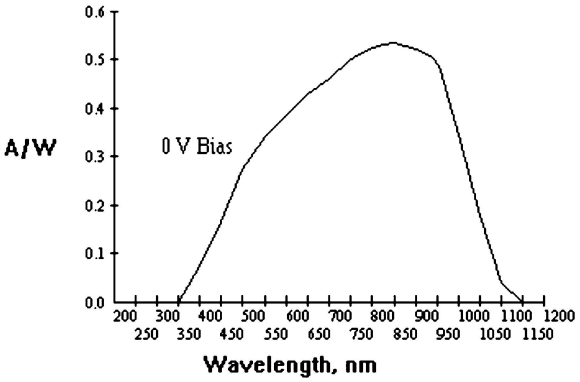

Figure 3. Spectral response of a typical silicon photodetector.

pixel sensor (APS) [27], which has been implemented into a digital camera on a chip with the functionality of analog-to-digital conversion [28].

2.2. Analog-to-Digital Conversion

The next stage is to convert the acquired image signal into a digital value. The interface between analog signals and their digital representation requires three tasks: (1) anti-alias filtering, (2) sampling, and (3) quantization [29]. These tasks can be implemented using an oversampling technique. Oversampling methods use sampling rates far above the Nyquist rate, which is the minimum rate for reconstructing a signal without aliasing [30]. Figure 4 shows a block diagram for analog-to-digital conversion [29]. The oversampled analog signal is passed through one or more integrators, represented by transfer function block H(z), before quantization. Using a feedback loop, the quantization noise is transformed (or shaped) into a high-pass response. The signal is then sent to a decimator, which combines the low-pass filtering operation and rate reduction. This removes the high-pass quantization noise and other signal information above the maximum input frequency of interest.

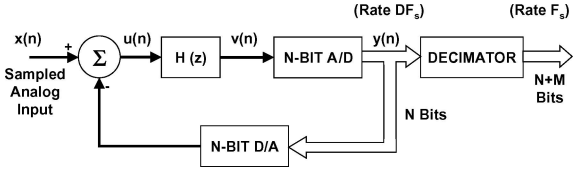

Figure 4. Generic noise-shaping feedback loop.

A candidate implementation of analog-to-digital conversion with noise shaping is the delta-sigma (or sigma-delta) converter [31]. These converters have become popular because they avoid some of the difficulties of conventional A/D conversion, such as the need for high-precision analog circuitry and vulnerability to noise or interference [31]. The performance of these circuits has been studied with standard CMOS implementation [32–34]. Sigma-delta converters have been successfully integrated with CMOS image sensors [35–37].

2.3. Data Storage

Data storage implementation, the most transistor-consuming function for a pixel-level processing element, affects most other aspects of the design [38]. The architecture must provide the storage capacity that would normally be found in a frame buffer for conventional store-and-process systems. Each application executed on the architecture has a minimum data storage threshold, typically consisting of the input image, the output image, and processing overhead. The data storage is also constrained by the minimum bit-precision required for accuracy. The architecture can utilize the six-transistor SRAM cell as the foundation of memory in the processing element. The instruction set architecture of the focal plane array imager requires a three-ported register file organization [39, 40]. In addition, models have been developed to predict silicon area usage for register file configurations [41, 42].

2.4. Digital Processing

Once the pixel information is represented and stored digitally, the programmable processing core can implement various image applications. Image processing applications are categorized into (1) point operations, (2) local operations, and (3) global operations [43]. A point operation such as thresholding occurs at the individual pixel level. A local operation such as smoothing requires knowledge of an individual pixel and its immediate neighbors. A global operation such as histogramming uses all pixel data from the image. These operations form the basis of applications found in imaging systems.

To illustrate the processing sequence required for a typical imaging system of $(N \times N)$ pixels, Fig. 5(a) describes the processing tasks, Fig. 5(b) gives example

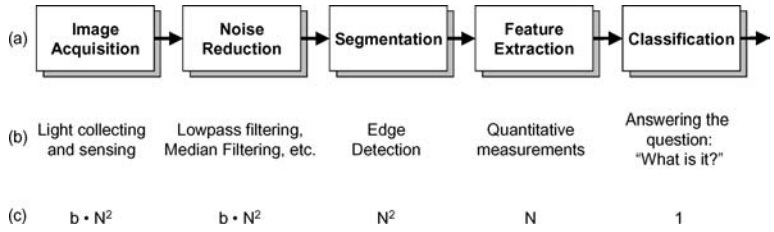

Figure 5. Typical image processing sequence (a) processing tasks (b) examples and (c) amount of data.

applications at each stage, and Fig. 5(c) quantifies the amount of data required for processing [44]. In Fig. 5(c), the variable b represents the number of bits per pixel and N represents the image resolution on side for a square image. Point operations are the first stage, and local operations include the first three stages. The fourth and fifth stages are global feature-measuring operations [45]. The first three stages represent a significant proportion of the computational workload. For example, the digital image signal multiprocessor [3] reports 53% of its execution time for preprocessing, which includes tasks from the first three stages such as noise reduction, smoothing, and segmentation. These stages are candidates for processing at the pixel level.

Several architectures, such as the Near Sensor Image Processing (NSIP) [46], the Programmable and Versatile Large Size Artificial Retina (PVLSAR) [38], and the Simple and Smart Sensory Processing Elements (S^3PE) [47], follow the pixel-level model to perform early image processing in the focal plane. However, these architectures utilize bit-serial processing techniques with limited memory, which can either restrict processing to binary images or require multiple cycles to perform a single instruction on data words. This may prevent the implementation of some early image processing algorithms. However, the integration of additional functional units at the pixel level enables the versatility to execute a broader set of applications.

3. Application Suite Analysis

This section describes the technique to characterize the performance of the focal plane architecture on an image application suite. A focal plane architectural simulator is used to implement the various algorithms. A description is given for each application. The performance is calculated and hardware constraints are determined for the application suite.

3.1. Focal Plane Architectural Simulator

Applications for the focal plane architecture can be programmed using the SIMD Pixel Processor (SIMPil) Simulator [48]. This software tool is a windows-based instruction level simulator, running on a PC platform. The SIMPil Simulator allows editing, assembling, executing, and debugging parallel image applications in a single integrated workbench. The current version of the SIMPil Simulator is available on the download page [48].

An image processing application is executed on the architectural simulator to provide a dynamic workload. Each application is implemented using an instruction set architecture (ISA) that corresponds to the available functional units within the digital pixel. The number of pixels per processing element is assigned in the application. The simulator is instrumented to measure the operand resolution for each instruction, the distribution for instructions issued in parallel, the concurrency level of the processing elements, and the number of processor cycles required for execution.

3.2. Description of Applications

The grayscale (8-bit) front-end applications, shown in Table 3, provide workload characteristics for image enhancement and image analysis. The enhancement and analysis applications of median filtering, convolution, and morphological processing represent the typical early image processing sequence of: (1) noise removal, (2) smoothing, and (3) segmentation. The workload characteristics for these applications are used to make efficient architectural choices for processing in the focal plane. This section briefly discusses each application.

3.2.1. Median Filtering. Median filtering (MED) is useful to remove impulse noise from an image while

Table 3. Selected early image processing applications.

Applications	Description
Image enhancement	
Median filtering	• Removes impulse noise from an image while preserving spatial resolution.
Convolution	• Performs different filtering operations, such as shadowing, smoothing, and edge-detection.
Image analysis	
Morphological processing	• Performs feature extraction and segmentation of binary or grayscale images such as inside edge detection.

preserving spatial resolution. The algorithm is a rank-order filter [49] that replaces each pixel in the image with the median value in the window. Generally, a window size is selected to generate a rank-order filter with odd length. A larger window size increases the severity of the median filtering effect [50]. This implementation performs a 2-D nonseparable rank and selects the median value for a 3×3 window.

3.2.2. Convolution. Convolution-based filtering (CONV) has been implemented to perform different filtering operations, such as shadowing, smoothing, and edge detection [50]. The filter mask elements are broadcast one at a time to every PE. All calculations requiring the mask element are performed before the next element is broadcast. Each PE multiplies the mask element by the corresponding pixel value from the original image and accumulates the result. Values are accumulated in a spiral pattern that places the final result in the center pixel of the filter mask. A 3×3 filter mask is used for this implementation.

3.2.3. Morphological Processing. Morphological image processing refers to the study of the topology or structure of objects from their 2D spatial representation [51]. Binary or grayscale images are morphologically transformed by passing a structuring element over the image in a process similar to convolution. At each pixel position, a specified logical function is performed between the structuring element and the underlying image. For grayscale images, erosion is the minimum pixel value in the structuring element, and dilation is the maximum pixel value in the structuring element [50]. Depending upon the size and content of the structuring element, different effects such as inside edge detection (IED) can be produced from erosion and dilation operations. A 3×3 structuring element has been used to implement the morphological operations.

3.3. Application Performance

The execution time was calculated for application suite described in Section 3.2. The execution time is determined with reference to a 10 MHz target platform This clock frequency addresses both the speed of analog components in each processing element and power density limitations for the high utilization factor of PEs in the array. Simulations were run with a system resolution equal to Quad-CIF (176×144 pixels). This resolution is one specification of the H.261 and H.263 video codec standards of the International Telecommunications Union (ITU) [52]. However, the execution time of the focal plane architecture is independent of the image dimension due to the parallel data processing.

In Table 4, the total execution times of the early image processing sequence generally follow the increase of the PPE factor with a baseline of the 1 PPE case. Median filtering is the dominant component of the processing sequence. The median filtering application for 16 PPE is not as efficient as the versions for 1 PPE and 4 PPE; therefore it does not follow the trend of the PPE factor. The total execution time for each PPE is within real-time constraints of 30 frames/sec (33.3 ms per frame).

Using the execution time, the sustained throughput Throughput_{sust}, measured in billion operations per second, is determined for the focal plane architecture as

Table 4. Total execution time for Quad-CIF focal plane architecture implementations using a 10 MHz clock.

	Application execution cycles				
PPE	MED	CONV	IED	Total	Total execution time (μs)
1	385	56	84	525	52.5
4	1531	183	238	1952	195.2
16	13531	826	979	15336	1533.6

follows:

$$\text{Throughput}_{sust} = \frac{IC \cdot U \cdot N_{\text{PE}}}{t_{\text{exec}}} \quad \left[\frac{\text{Gops}}{s}\right] \tag{1}$$

where IC is number of parallel instructions issued to the PE array during the application. The system utilization U is calculated as the average number of active processing elements determined from the simulator's concurrency meter. N_{PE} is the number of processing elements in the PE array and is determined from the following formula:

$$N_{\text{PE}} = \frac{\text{system_resolution}}{\text{PPE}} \tag{2}$$

where PPE is the number of pixels in each processing element (1, 4, or 16). For $PPE > 1$, pixels are arranged in a square (i.e. 2×2 or 4×4). The target system resolution is QCIF (176×144 pixels). The values for N_{PE} are shown in Table 5.

The sustained throughput has been calculated for focal plane architecture with 1 PPE, 4 PPE, and 16 PPE implementations (Table 6). The maximum number of PE instructions executed equals $(IC \cdot N_{\text{PE}})$. The utilization is derived from the weighted average of executing each application. The sustained throughput achieved by the focal plane architecture implementations exceeds the values achieved by a traditional DSP architecture despite a significant difference in clock frequency. For example, the throughput specification of the TI TMS320C6211B DSP chip running at 150 MHz is 1.2 Gops/s, and the throughput specification of the TI TMS320C6411 DSP chip running at 300 MHz is 2.4 Gops/s [53]. Although operating at 10 MHz, the focal plane architecture implementations for 1 PPE, 4 PPE, and 16 PPE increase sustained throughput by factors of 65x, 16x, and 5x respectively compared to the TI TMS320C6411 chip. These factors are doubled when comparing against the TI TMS320C6211B chip.

Table 5. Focal plane processor array characteristics for QCIF resolution.

	PE array dimensions		
PPE	X	Y	Total number of PEs
1	176	144	25344
4	88	72	6336
16	44	36	1584

Table 6. Application characterization for focal plane architecture implementations using a 10 MHz clock with Quad-CIF resolution.

PPE	Execution time (μs)	Maximum PE instructions executed	Utilization	Sustained throughput (Gops/s)
1	52.5	9,529,344	85.56%	155
4	195.2	8,838,720	84.95%	38
16	1533.6	19,209,168	87.91%	11

3.4. Application Hardware Constraints

The hardware must satisfy both the data storage and the data precision constraints to execute the algorithms. Table 7 shows the architectural design parameters for each processing element determined by the application suite. Register requirements are derived from the code for the implemented algorithm and include twice the PPE (to store the input and output images) plus overhead for intermediate calculations. Datapath requirements are determined from the measured operand resolutions for each instruction during simulation. With an overhead of 2 registers, convolution required the fewest total registers, $[(2 \cdot \text{PPE}) + 2]$, but the widest datapath (12 bits) to accommodate the successive multiply-accumulates of pixel values with the corresponding 3×3 mask values. The other applications only required the original grayscale resolution of 8 bits. Median filtering required $[(2 \cdot \text{PPE}) + 7]$ registers to store and sort the pixels in the 3×3 window. A sufficient architecture capable of executing all three applications would require $[(2 \cdot \text{PPE}) + 7]$ registers with a 12-bit datapath. These values are used to select the sizes of both the ALU and the register file in each processing element.

4. Projected Area

Because of limited chip resources in a focal plane architecture, silicon area usage within an integrated digital

Table 7. Application constraints for focal plane processor implementation.

Application	Registers required	Datapath required
Median filtering	$(2 \cdot \text{PPE}) + 7$	8 bits
Convolution	$(2 \cdot \text{PPE}) + 2$	12 bits
Inside edge detection	$(2 \cdot \text{PPE}) + 3$	8 bits
Full application suite	$(2 \cdot \text{PPE}) + 7$	12 bits

pixel is a critical design factor. A pixel design tool provides a common context for component area models within an integrated pixel-processing array [54]. Area models based upon implementations of detector array circuitry and CMOS functional units are used to project silicon area for the architecture using different fabrication technologies.

4.1. Component Area Models

Area models are developed from fabricated analog and digital components. Area projections of the photodiode and analog-to-digital conversion are based upon the CMOS focal plane array [5]. This 8×8 array of Si CMOS detectors, fabricated in 0.8 μm technology, incorporates a current input first-order sigma-delta analog-to-digital converter at each pixel. The transistor circuitry is scaled to feature sizes ranging from 250 nm to 100 nm. Area projections for memory and digital circuitry are based upon the SIMPil16 focal plane architecture [55] that maps 16 pixels to each processing element. SIMPil16 is a 16-bit implementation fabricated on MOSIS HP 0.8 μm (1.0 μm drawn) CMOS technology. Selected functional units of a SIMPil16 PE used by the pixel-level processing architecture include: (1) ALU, (2) register file, (3) decoder, (4) bus driver, (5) sleep unit, and (6) communication unit. The original areas for these functional units are scaled to feature sizes ranging from 250 nm to 100 nm for area projections of the focal plane processor. Bit slicing is used to adjust both the ALU area and the register file area for a reduced datapath width. In addition, the register file area is adjusted to correspond to the number of words in the design.

4.2. Processing Element Area

Silicon area allocation is a significant issue because in single-level VLSI, the photodiode, the analog signal conditioning, the analog-to-digital converter, the memory, and the digital processing core compete for silicon area in a small replicated processing element (Fig. 2). A P-N photodiode is used in the CMOS focal plane array [5]. As a first-order approximation, the area of the sampling capacitor represents the analog signal conditioning because it dominates that component. The ADC is a first-order sigma-delta circuit. Memory is a register file with the number or registers based upon the formula required to execute the full application suite (Table 7). The digital processor contains the functional units mentioned in the previous section.

Processing element area has been projected using various fabrication technologies for 1 PPE (Table 8), 4 PPE (Table 9), and 16 PPE (Table 10). By observing the component area trends for 1 PPE (Fig. 6), 4 PPE (Fig. 7), and 16 PPE (Fig. 8), the digital components (digital processor, memory, ADC) benefit from technology scaling. However, the analog components (photodiode and analog signal conditioning) do not see the same benefit [25, 56]. The photodiode has a fixed area requirement to ensure acquisition of light. The analog signal conditioning (sampling capacitor) also has a fixed area cost. As the PPE increases, a larger percentage of the silicon area within the processing element is used for data acquisition and data storage.

4.3. Focal Plane Processor Array Area

Portable devices with an embedded focal plane architecture must reduce their silicon area cost for a compact design. Excessive die sizes would limit the

Table 8. Area of a single processing element using 1 PPE.

	Components					
Technology	Photodiode (mm^2)	Analog signal conditioning (mm^2)	ADC (mm^2)	Memory (mm^2)	Digital processor (mm^2)	Total area (mm^2)
250 nm	0.0046	0.0024	0.0008	0.0287	0.0453	0.082
180 nm	0.0046	0.0024	0.0004	0.0149	0.0235	0.046
150 nm	0.0046	0.0024	0.0003	0.0103	0.0163	0.034
130 nm	0.0046	0.0024	0.0002	0.0078	0.0122	0.027
120 nm	0.0046	0.0024	0.0002	0.0066	0.0104	0.024
100 nm	0.0046	0.0024	0.0001	0.0046	0.0073	0.019

Table 9. Area of a single processing element using 4 PPE.

	Components					
Technology	Photodiode (mm^2)	Analog signal conditioning (mm^2)	ADC (mm^2)	Memory (mm^2)	Digital processor (mm^2)	Total area (mm^2)
250 nm	0.0185	0.0096	0.0030	0.0478	0.0453	0.124
180 nm	0.0185	0.0096	0.0016	0.0248	0.0235	0.078
150 nm	0.0185	0.0096	0.0011	0.0172	0.0163	0.063
130 nm	0.0185	0.0096	0.0008	0.0129	0.0122	0.054
120 nm	0.0185	0.0096	0.0007	0.0110	0.0104	0.050
100 nm	0.0185	0.0096	0.0005	0.0077	0.0072	0.043

Table 10. Area of a single processing element using 16 PPE.

	Components					
Technology	Photodiode (mm^2)	Analog signal conditioning (mm^2)	ADC (mm^2)	Memory (mm^2)	Digital processor (mm^2)	Total area (mm^2)
250 nm	0.0739	0.0384	0.0121	0.1244	0.0453	0.294
180 nm	0.0739	0.0384	0.0063	0.0645	0.0235	0.207
150 nm	0.0739	0.0384	0.0044	0.0448	0.0163	0.178
130 nm	0.0739	0.0384	0.0033	0.0336	0.0122	0.161
120 nm	0.0739	0.0384	0.0028	0.0287	0.0104	0.154
100 nm	0.0739	0.0384	0.0019	0.0199	0.0072	0.141

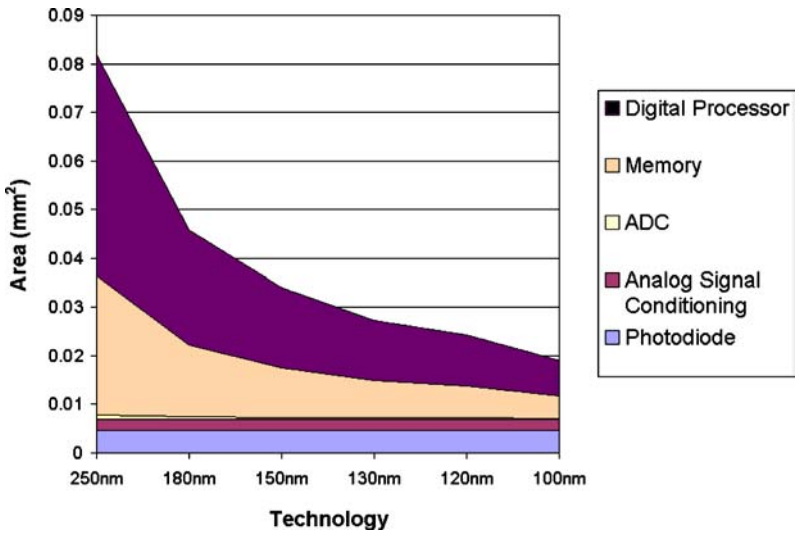

Figure 6. Component area trend for 1 PPE with decreasing feature size.

realization of these systems. The chip area is tightly coupled to the system resolution because of the integrated functionality within each processing element. For this analysis, the desired image resolution for processing is grayscale (8-bit) Quad-CIF (176 × 144 pixels). The Quad-CIF resolution is one specification of the H.261 and H.263 video codec standards of the International Telecommunications Union (ITU) [52]. The

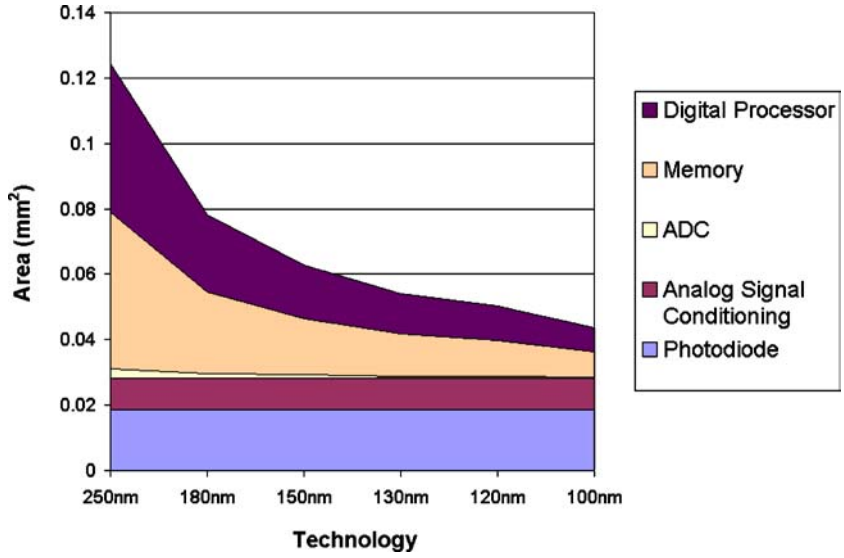

Figure 7. Component area trend for 4 PPE with decreasing feature size.

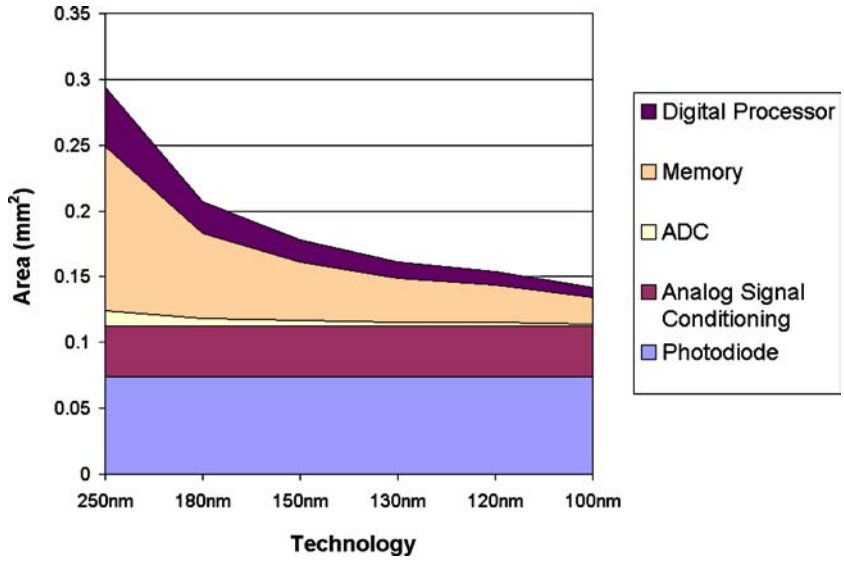

Figure 8. Component area trend for 16 PPE with decreasing feature size.

ITRS provides guidelines for affordable chip sizes for each technology [10, 11].

Using the array dimensions for different PPE values (Table 5) and the processing element areas (Table 8, 9, 10), the focal plane processor array area has been projected for various technologies (Fig. 9). Using the ITRS guideline for manufacturing affordability, the 1 PPE implementation does not project to the die size target. The 4 PPE implementation exceeds the guideline by approximately 20% starting with the 150 nm technology and reaches the ITRS guideline using 100 nm technology. The 16 PPE implementation meets the guideline for all the selected technologies except 250 nm, allowing it to be implemented using a less aggressive and less expensive fabrication technology. Because detector area dominates as feature size shrinks, denser detector technology could substantially reduce chip area, making the chip more affordable. Another

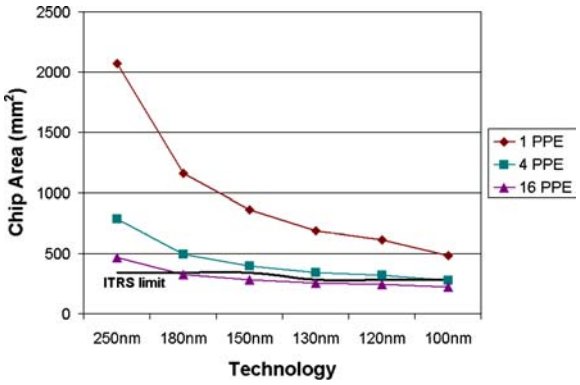

Figure 9. Projected area for a QCIF processor array implementation.

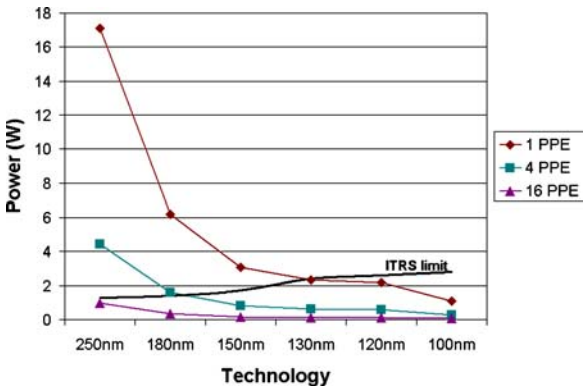

Figure 10. Projected power for a QCIF processor array implementation using a 10 MHz clock.

potential technique to reduce the analog component area is to share the analog-to-digital conversion circuitry [12].

5. Projected Power

Portable devices with an embedded focal plane architecture must provide a meaningful battery life. Setting the minimum time between battery charges (*MTBC*) to a desired value translates into a limit for maximum power consumed during operation for a fixed battery energy. Typical double barrel, NiCd AA sized batteries (3.6 V at 700 mA·hours) have an energy capacity of about 10 Watt·hours. Therefore for $MTBC \geq 3$ hours, the power should not exceed 3 Watts.

A Technology Scenario Analyzer (TeSA) was developed to project the power consumption of the digital components within the focal plane architecture using different fabrication technologies [57]. The power consumption has been projected for fabrication technologies ranging from 250 nm to 100 nm using a 10 MHz clock (Fig. 10). The maximum power limit for each technology is used from ITRS specifications for portable battery operation [10, 11]. The 1 PPE implementation consumes the most power because of the large number of processing elements required to provide QCIF resolution ($176 \times 144 = 25344$). It requires technology generations of 130nm or smaller to satisfy the ITRS power constraint. However, it may be too close to the limit to account for the analog components. The 4 PPE implementation satisfies the constraint for all technologies except 250 nm and 180 nm by reducing the required number of processing elements ($88 \times 72 = 6336$). The power consumed is reduced by a factor of 3.9 (i.e. 74%) when compared to the 1 PPE implementation. The 16 PPE implementation, which uses the least number of processing elements ($44 \times 36 = 1584$), consumes the least power and satisfies the constraint for all the chosen technologies. The overall reduction factor is 18 (i.e. 94%) when compared to the 1 PPE implementation.

6. Resource Efficiency Analysis

Higher ratings in area-time efficiency, area efficiency, and energy efficiency metrics are desired for image processing systems because of technology limitations such as interconnect wiring density and heat extraction [58]. This section examines the area-time efficiency, area efficiency, and energy efficiency of the focal plane architecture implementations. Comparisons for TI DSP chips are based upon application benchmarks and chip specifications [53].

6.1. Area-Time Efficiency

The area-time efficiency is defined as:

$$\eta_{A-T} = \frac{1}{\text{Area} \times t_{\text{exec}}} \quad \left[\frac{1}{s \cdot mm^2}\right] \tag{3}$$

where the reciprocal is taken for the product of the area for each system and the execution time for the imaging processing sequence (Table 4). Because a system should execute in the shortest time using the smallest area, the optimal solution is determined by maximizing the reciprocal product.

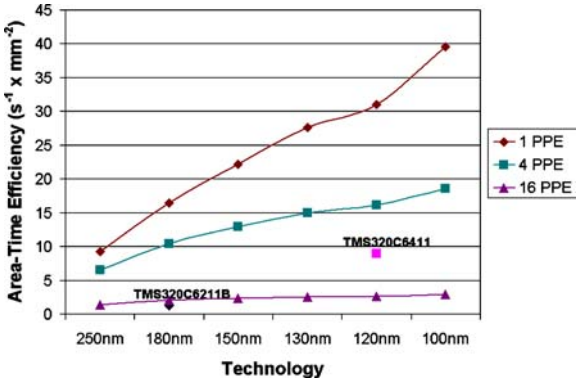

Figure 11. Comparison of $(\text{area} \cdot \text{time})^{-1}$ efficiency for image processing architectures.

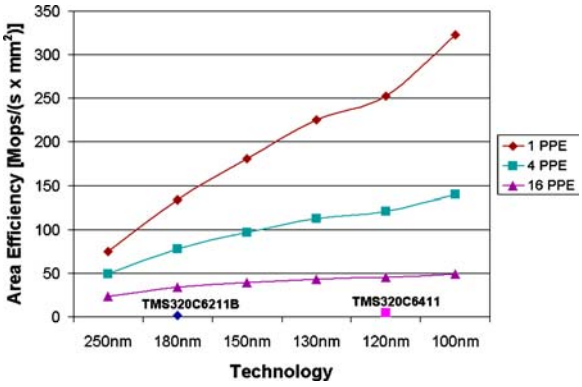

Figure 12. Comparison of area efficiency for image processing architectures.

The $(A \cdot T)^{-1}$ efficiency of focal plane architecture implementations using 1 PPE, 4 PPE, and 16 PPE, has been calculated for different fabrication technologies (Fig. 11). The area includes the data acquisition and analog-to-digital conversion circuitry. Two TI DSP chips, the TMS320C6211B and TMS320C6411, are shown for comparison. The 16 PPE implementation is comparable to the TMS320C6211B operating 150 MHz. The 1PPE and the 4 PPE implementations provide a 3.5× and 1.8× factor increase respectively when compared to the TMS320C6411 operating at 300 MHz. This demonstrates that a focal plane architecture can be utilized as an embedded component to deliver efficient processing for common image applications. These applications can consume over half of the processing for an image.

6.2. *Area Efficiency*

Area efficiency is defined as the number of operations executed per second per unit area:

$$\eta_A = \frac{\text{Throughput}_{sust}}{\text{Area}} \quad \left[\frac{\text{Mops}}{s \cdot mm^2}\right] \tag{4}$$

Area efficiency has been established as an important figure of merit for system-on-a-chip architecture design [7]. Because of limited chip resources in a focal plane architecture, silicon area usage within an integrated digital pixel is a critical design factor. Higher area efficiency implies better component utilization within the architecture.

The area efficiency of focal plane architecture for implementations using 1 PPE, 4 PPE, and 16 PPE, has been calculated using different fabrication technologies (Fig. 12). Generally, an increase in PPE decreases the area efficiency of the focal plane architecture. However, there is a significant gain when compared to a DSP architecture in the same fabrication technology. For the TMS320C6211B, which uses 180nm technology, the area efficiency increases by factors of 81×, 47×, and 20× for 1 PPE, 4 PPE, and 16 PPE respectively. For the TMS320C6411, which uses 120 nm technology, the increase factors are 56×, 27×, and 10× for 1 PPE, 4 PPE, and 16 PPE respectively. This result is significant because the gain in efficiency is for common image enhancement and image analysis tasks that can represent a significant portion of the total operations within an image processing sequence. In addition, the area efficiency for implementations of the focal plane architecture includes the data acquisition and analog-to-digital conversion circuitry, which is not available with the DSP architecture. This does not suggest that the architecture replaces a DSP but can handle certain common image enhancement and image analysis tasks more efficiently through specialization.

6.3. *Energy Efficiency*

Energy efficiency is defined as the number of operations executed per unit energy:

$$\eta_E = \frac{\text{Throughput}_{sust}}{\text{Power}} \quad \left[\frac{\text{Mops}}{\text{Joule}}\right] \tag{5}$$

Previous work [59, 60] has illustrated the validity of energy efficiency for fixed throughput computation in uniprocessor systems. The validity is extended to

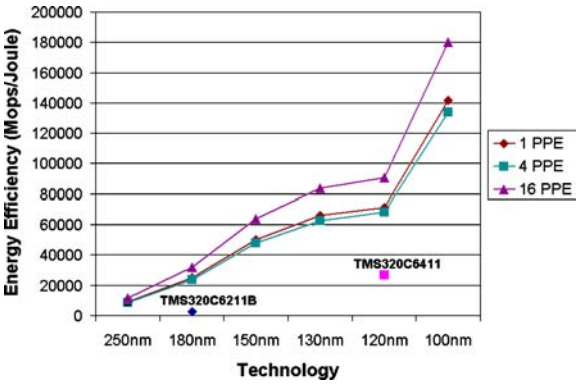

Figure 13. Comparison of energy efficiency for image processing architectures.

massively parallel embedded-focal plane architectures by introducing system utilization in (1) for calculating the throughput. Increasing energy efficiency implies enhancing the sustainable battery life in portable devices. Minimizing power dissipation translates into minimizing energy per operation.

The energy efficiency of a focal plane architecture using 1 PPE, 4 PPE, and 16 PPE, has been calculated for the digital components using different fabrication technologies (Fig. 13). The 16 PPE implementation delivers the highest energy efficiency. The 4 PPE implementation is slightly lower than the 1 PPE implementation. This may indicate that the application suite optimizes slightly better for using a single pixel in a processing element. Two TI DSP chips, the TMS320C6211B and TMS320C6411, are also shown for comparison. For the TMS320C6211B with 180 nm technology, the energy efficiency increases by an average factor of 11× when using implementations of the focal plane architecture. For the TMS320C6411 with 120 nm technology, the average increase is 2.9×. This demonstrates the potential to extend battery life in portable devices.

7. Conclusion

Recently, a new dynamic has emerged with users desiring mobile productivity, where image processing functionality is encapsulated within portable devices containing embedded hardware. Portable imaging products will benefit from the monolithic integration of photodetectors, analog-to-digital converters, digital processing, and data storage to perform image acquisition and computation. Processing on the focal plane utilizes data-parallelism naturally found in image applications. Technological advances in device fabrication and integration are enabling the development of these monolithic focal plane architectures with the potential for improved performance, efficiency, and cost versus traditional imaging architectures.

This paper presents the analysis of a mixed-signal focal plane processing architecture. Key performance and cost metrics for focal plane architectures include execution time, system throughput, chip area, power consumption, area-time efficiency, area efficiency, and energy efficiency. Although the area-time efficiency has been traditionally used, it does not provide guidance for addressing energy consumption. Area efficiency and energy efficiency are important figures of merit when evaluating a focal plane architecture. The choice of both fabrication technology and pixels per processing element impact the system design, particularly for portable devices. Die size for manufacturing cost and power consumption for battery life should follow ITRS guidelines. In addition, higher ratings in area efficiency and energy efficiency are required because of technology limitations such as interconnect wiring density and heat extraction.

The performance, area, and power were projected for focal plane architecture implementations using 1 PPE, 4 PPE, and 16 PPE. This architecture is more area and energy efficient when compared to traditional TI DSP chips. However, the design problem is not strictly building the system with the highest performance, but delivering the required performance with the lowest cost. Therefore, a 4 PPE or 16 PPE with a less aggressive technology may provide a better implementation solution.

8. Future Research Directions

Future work includes analyzing the design of the CMOS image sensor circuitry to quantify the performance versus both the area and energy consumed. Although the analog components do not receive dramatic benefits from technology scaling, improvements in optoelectronic fabrication can reduce the size requirements of the photodiode and the sampling capacitor. In addition, the analog components may require a significant amount of power when compared to the digital components. The components of the digital pixel will be implemented in VLSI layout to extract circuit parameters for analysis. The layout may reveal additional opportunities for shared components, such as power or signal lines. This layout would enable a prototype of the architecture to be fabricated and tested, as well as

provide feedback to improve the area models for the photodetectors, analog-to-digital converters, data storage, and digital processing circuitry.

Improved architectural developments must also address data storage and data communication issues found in focal plane systems. Data storage and data communication are tightly coupled within the focal plane system. Because data storage is the most transistor-consuming function for a pixel-processing element, implementation choices affect most other aspects of the design [38]. Data communication among processing elements follows directly from the method of storage. Pixel level processing could benefit from a distributed or reconfigurable processing architecture that allows data storage and data transfer to merge into a communication network.

Acknowledgment

The first author would like to acknowledge the generous support of this research from the Ford Foundation Dissertation Fellowship for Minorities and the Facilitating Academic Careers in Engineering and Science (FACES) Fellowship.

References

1. R.B. Lee and M.D. Smith, "Media Processing: A New Design Target," *IEEE Micro*, vol. 16, 1996, pp. 6–9.
2. K. Diefendorff and P.K. Dubey, "How Multimedia Workloads will Change Processor Design," *Computer*, vol. 30, 1997, pp. 43–45.
3. M. Maruyama, H. Nakahira, T. Araki, S. Sakiyama, Y. Kitao, K. Aono, and H. Yamada, "An Image Signal Multiprocessor on a Single Chip," *IEEE Journal of Solid-State Circuits*, vol. 25, 1990, pp. 1476–1483.
4. E.R. Fossum, "CMOS Image Sensors: Electronic Camera-on-a-Chip," *IEEE Transactions on Electron Devices*, vol. 44, 1997, pp. 1689–1698.
5. Y. Joo, J. Park, M. Thomas, K.S. Chung, M.A. Brooke, N.M. Jokerst, and D.S. Wills, "Smart CMOS Focal Plane Arrays: A Si CMOS Detector Array and Sigma-Delta Analog-to-Digital Converter Imaging System," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 5, 1999, pp. 296–305.
6. T. Morris, E. Fletcher, C. Afghahi, S. Issa, K. Connolly, and J.-C. Korta, "A Column-Based Processing Array for High-Seed Digital Image Processing," in *Conference on Advanced Research in VLSI*. Atlanta, GA, USA: IEEE Computer Society, 1999, pp. 42–56.
7. R.W. Brodersen, "Why We Need a Custom Chip-in-a-Day Design Methodology," http://www.mseconference.org/Talks/mse01_brodersen.pdf, International Conference on Microelectronic Systems Education, Mixed Signal & Fabrication Services Keynote Address, 2001.
8. M.C. Herbordt and C.C. Weems, "An Empirical Study of Datapath, Memory Hierarchy, and Network in SIMD Array Architectures," in *International Conference on Computer Design: VLSI in Computers and Processors*. Austin, TX, USA: IEEE Computer Society Press, 1995, pp. 546–551.
9. M.C. Herbordt, A. Anand, O. Kidwai, R. Sam, and C.C. Weems, "Processor/memory/array Size Tradeoffs in the Design of SIMD Aays for a Spatially Mapped Workload," in *Fourth IEEE International Workshop on Computer Architecture for Machine Perception*. Cambridge, MA, USA: IEEE Computer Society Press, 1997, pp. 12–21.
10. Semiconductor Industry Association, *The International Technology Roadmap for Semiconductors 1999 Edition*, 1999.
11. Semiconductor Industry Association, *The International Technology Roadmap for Semiconductors 2002 Update*, 2002.
12. D. Yang, H. Tian, B. Fowler, X. Liu, and A. El Gamal, "Characterization of CMOS Image Sensors with Nyquist Rate Pixel Level ADC," in *Sensors, Cameras, and Applications for Digital Photography*, vol. 3650. San Jose, CA, USA: SPIE—The International Society for Optical Engineering, 1999, pp. 52–62.
13. P. Pirsch and H.-J. Stolberg, "VLSI Implementations of Image and Video Multimedia Processing Systems," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 8, 1998, pp. 878–891.
14. A. Moini, *Vision Chips*, Boston: Kluwer Academic, 2000.
15. E.-S. Eid and E.R. Fossum, "Real-Time Focal-Plane Array Image Processor," in *Automated Inspection and High-Speed Vision Architectures III*, vol. 1197. Philadelphia, PA, USA: SPIE—The International Society for Optical Engineering, 1989, pp. 2–12.
16. C. Mead, *Analog VLSI and Neural Systems*, Reading, Mass.: Addison-Wesley, 1989.
17. J. Kramer and G. Indiveri, "Neuromorphic Vision Sensors and Preprocessors in System Applications," in *Advanced Focal Plane Arrays and Electronic Cameras II*, vol. 3410. Zurich, Switzerland: SPIE—The International Society for Optical Engineering, 1998, pp. 134–146.
18. L.O. Chua and L. Yang, "Cellular Neural Networks: Theory and Applications," *IEEE Transactions on Circuits and Systems*, vol. 35, 1988, pp. 1257–1290.
19. S. Espejo, R. Dominguez-Castro, G. Linan, and A. Rodriguez-Vazquez, "A 64*64 CNN Universal Chip with Analog and Digital I/O," in *IEEE International Conference on Electronics, Circuits and Systems*, vol. 1, 1998, pp. 203–206.
20. A. El Gamal, D. Yang, and B. Fowler, "Pixel Level Processing—Why, What, and How?" in *Sensors, Cameras, and Applications for Digital Photography*, vol. 3650. San Jose, CA, USA: SPIE—The International Society for Optical Engineering, 1999, pp. 2–13.
21. R.F. Pierret, *Semiconductor Device Fundamentals*, Reading, Mass.: Addison-Wesley, 1996.
22. M.L. Riaziat, *Introduction to High-Speed Electronics and Optoelectronics*, New York: Wiley, 1996.
23. P. Bhattacharya, *Semiconductor Optoelectronic Devices*, Upper Saddle River, NJ: Prentice Hall, 1997.
24. A. Dupret, E. Belhaire, and J.-C. Rodier, "A High Current Large Bandwidth Photosensor on Standard CMOS Processes," *Advanced Focal Plane Arrays and Electronic Cameras*, vol. 2950, 1996, pp. 36–44.

25. T. Lule, S. Benthien, H. Keller, F. Mutze, P. Rieve, K. Seibel, and M. Sommer, "Sensitivity of CMOS Based Imagers and Scaling Perspectives," *IEEE Transactions on Electron Devices*, vol. 47, 2000, pp. 2110–2122.
26. M.L. Simpson, M.N. Ericson, G.E. Jellison, Jr., W.B. Dress, A.L. Wintenberg, and M. Bobrek, "Application Specific Spectral Response with CMOS Compatible Photodiodes," *IEEE Transactions on Electron Devices*, vol. 46, 1999, pp. 905–913.
27. S.K. Mendis, S.E. Kemeny, R.C. Gee, B. Pain, C.O. Staller, Q. Kim, and E.R. Fossum, "CMOS Active Pixel Image Sensors for Highly Integrated Imaging Systems," *IEEE Journal of Solid-State Circuits*, vol. 32, 1997, pp. 187–197.
28. E.R. Fossum, "Digital Camera System on a Chip," *IEEE Micro*, vol. 18, 1998, pp. 8–15.
29. M.W. Hauser, "Principles of Oversampling A/D Conversion," *Journal of the Audio Engineering Society*, vol. 39, 1991, pp. 3–26.
30. J.G. Proakis, *Digital communications*, New York: McGraw-Hill, 1995.
31. J.C. Candy and G.C. Temes, "Institute of Electrical and Electronics Engineers, and IEEE Circuits and Systems Society," *Oversampling delta-sigma data converters: Theory, Design, and Simulation*, Piscataway, NJ: IEEE Press, 1992.
32. M.W. Hauser and R.W. Brodersen, "Circuit and Technology Considerations for MOS Delta-sigma A/D Converters," in *International Symposium on Circuits and Systems*, vol. 3. San Jose, CA, USA: IEEE, 1986, pp. 1310–1315.
33. M.W. Hauser and R.W. Brodersen, "Monolithic Decimation Filtering for Custom Delta-Sigma A/D Converters," in *International Conference on Acoustics, Speech, and Signal Processing*, New York, NY, USA: ICASSP, vol. 4, 1988, pp. 2005–2008.
34. M.W. Hauser, "Technology Scaling and Performance Limitations in Delta-Sigma Analog-Digital Converters," *International Symposium on Circuits and Systems*, vol. 1, 1990, pp. 356–359.
35. J. Nakamura, B. Pain, T. Nomoto, T. Nakamura, and E.R. Fossum, "On-Focal-Plane Signal Processing for Current-Mode Active Pixel Sensors," *IEEE Transactions on Electron Devices*, vol. 44, 1997, pp. 1747–1758.
36. L.G. McIlrath, "A Low-Power Low-Noise Ultrawide-Dynamic-Range CMOS Imager with Pixel-Parallel A/D Conversion," *IEEE Journal of Solid-State Circuits*, vol. 36, 2001, pp. 846–853.
37. Y. Perelman and R. Ginosar, "A Low-Light-Level Sensor for Medical Diagnostic Applications," *IEEE Journal of Solid-State Circuits*, vol. 36, 2001, pp. 1553–1558.
38. F. Paillet, D. Mercier, and T.M. Bernard, "Making the Most of 15 k Lambda^2 Silicon Area for a Digital Retina PE," in *Advanced Focal Plane Arrays and Electronic Cameras II*, Zurich, Switzerland: SPIE—The International Society for Optical Engineering, vol. 3410, 1998, pp. 158–167.
39. R.L. Franch, J. Ji, and C.L. Chen, "A 640-ps, 0.25- mu m CMOS, 16* 64-b Three-Port Register File," *IEEE Journal of Solid-State Circuits*, vol. 32, 1997, pp. 1288–1292.
40. M. Nomura, M. Yamashina, K. Suzuki, M. Izumikawa, H. Igura, H. Abiko, K. Okabe, A. Ono, T. Nakayama, and H. Yamada, "A 500-MHz, 0.4- mu m CMOS, 32-word by 32-bit 3-port Register File," *IEEE 1995 Custom Integrated Circuits Conference*. Santa Clara, CA, USA: IEEE, 1995, pp. 151–154.
41. J.M. Mulder, N.T. Quach, and M.J. Flynn, "An Area Model for On-Chip Memories and its Application," *IEEE Journal of Solid-State Circuits*, vol. 26, 1991, pp. 98–106.
42. S. Rixner, W.J. Dally, B. Khailany, P. Mattson, U.J. Kapasi, and J.D. Owens, "Register Organization for Media Processing," in *Sixth International Symposium on High-Performance Computer Architecture*, Touluse, France, 2000, pp. 375–386.
43. R.S. Bajwa, R.M. Owens, and M.J. Irwin, "Image Processing with the MGAP: A Cost Effective Solution," in *International Parallel Processing Symposium*. Newport, CA, USA, 1993, pp. 439–443.
44. R. Forchheimer and A. Astrom, "Near-Sensor Image Processing: A New Paradigm," *IEEE Transactions on Image Processing*, vol. 3, 1994, pp. 736–746.
45. A. Astrom, J.-E. Eklund, and R. Forchheimer, "Near-Sensor Image Processing—Theory and Practice," in *Advanced Focal Plane Arrays and Electronic Cameras*, vol. 2950. Berlin, Germany: SPIE—The International Society for Optical Engineering, 1996, pp. 242–253.
46. J.-E. Eklund, C. Svensson, and A. Astrom, "VLSI Implementation of a Focal Plane Image Processor—A Realization of the Near-Sensor Image Processing Concept," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 4, 1996, pp. 322–335.
47. T. Komuro, I. Ishii, and M. Ishikawa, "General-Purpose Vision Chip Architecture for Real-Time Machine Vision," *Advanced Robotics*, vol. 12, 1999, pp. 619–627.
48. SIMPil Programming Environment, http://www.ece.gatech.edu/research/pica/curr_projects/civp/index.html.
49. G. Heygster, "Rank Filters in Digital Image Processing," *Computer Graphics and Image Processing*, vol. 19, 1982, pp. 148–164.
50. M.J.T. Smith and A. Docef, *A Study Guide for Digital Image Processing*, Riverdale, GA: Scientific Publishers, 1997.
51. R.M. Haralick, S.R. Sternberg, and X. Zhuang, "Image Analysis Using Mathematical Morphology," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 9, 1987, pp. 532–550.
52. International Telecommunications Union, http://www.itu.int/home/index.html.
53. Texas Instruments Incorporated, http://dspvillage.ti.com/.
54. W.H. Robinson, G.E. Triplett, and D.S. Wills, "Component Modeling for an Integrated Digital Pixel," in *15th Annual Meeting IEEE Lasers and Electro-Optics Society (LEOS)*, Glasgow, Scotland: IEEE, vol. 1, 2002, pp. 37–38.
55. A. Gentile, A. López-Lagunas, S.M. Chai, H.H. Cat, K.S. Chung, L. Codrescu, M. Deb, S.J. Ryu, M.F. Wolff, J.C. Eble, W.H. Robinson, T. Taha, D.S. Wills, M. Viera-Vera, W.E. Lugo-Beauchamp, L. Bustelo, R. Olivieri, J. Figueroa, and J.L. Cruz-Rivera, "Portable Multimedia Supercomputing," *Submitted to IEEE Transactions on Computers*, vol. TC 114531, July 2001, pp. 43.
56. H.-S. Wong, "Technology and Device Scaling Considerations for CMOS Imagers," *IEEE Transactions on Electron Devices*, vol. 43, 1996, pp. 2131–2142.
57. S.M. Chai, A. Gentile, and D.S. Wills, "Impact Of Power Density Limitation In Gigascale Integration For The SIMD Pixel Processor," in *20th Anniversary Conference on Advanced Research in VLSI*. Atlanta, GA, USA: IEEE Computer Society, 1999, pp. 57–71.

58. A. Gentile, "Portable Multimedia Supercomputers: System Architecture Design and Evaluation," in *Electrical and Computer Engineering*, Atlanta, GA: Georgia Institute of Technology, 2000, pp. 170.
59. A.P. Chandrakasan, S. Sheng, and R.W. Brodersen, "Low-power CMOS Digital Design," *IEEE Journal of Solid-State Circuits*, vol. 27, 1992, pp. 473–484.
60. T.D. Burd and R.W. Brodersen, "Energy Efficient CMOS Microprocessor Design," *Twenty-Eighth Annual Hawaii International Conference on System Sciences*, vol. 1, 1995, pp. 288–297.

William H. Robinson is an Assistant Professor in the Department of Electrical Engineering and Computer Science at Vanderbilt University. He received his B.S. in electrical engineering from Florida Agricultural and Mechanical University in 1996 and his M.S. in electrical engineering from the Georgia Institute of Technology (Georgia Tech) in 1998. He received his Ph.D. in electrical and computer engineering from Georgia Tech in 2003. His research explores the system-level integration of computer architecture to understand the impact of technology on architecture design. Topics of interest include computer architecture design, VLSI design, image processing, and mixed-signal integration with applications to portable imaging devices, integrated sensor technology, and system-on-a-chip multimedia processing. He is a member of the IEEE and participates in the Computer Society, the Education Society, and the Lasers and Electro-Optics Society.
William.h.robinson@vanderbilt.edu

D. Scott Wills is a Professor of Electrical and Computer Engineering at the Georgia Institute of Technology. He received his B.S. in Physics from Georgia Tech in 1983, and his S.M., E.E., and Sc.D. in Electrical Engineering and Computer Science from M.I.T. in 1985, 1987, and 1990, respectively. His research interests include short wire VLSI architectures, high throughput portable processing systems, architectural modeling for gigascale (GSI) technology, and high efficiency image processors. He is a senior member of the IEEE and the Computer Society and he is an associate editor of *IEEE Transactions on Computers*.