Hexagonal sampling in the infrared domain: an introduction to array set addressing

Nicholas I. Rummelt, Geoffrey L. Barrows, Mark A. Massie

Nicholas I. Rummelt, Geoffrey L. Barrows, Mark A. Massie, "Hexagonal sampling in the infrared domain: an introduction to array set addressing," Proc. SPIE 8012, Infrared Technology and Applications XXXVII, 80120I (21 May 2011); doi: 10.1117/12.896888

Event: SPIE Defense, Security, and Sensing, 2011, Orlando, Florida, United States

Invited Paper

Hexagonal Sampling in the Infrared Domain: an Introduction to Array Set Addressing

Nicholas I. Rummelt[a*], Geoffrey L. Barrows[b], and Mark A. Massie[c]

[a]Air Force Research Laboratory, AFRL/RWGI, Eglin AFB, FL, 32542, U.S.A.;
[b]Centeye, Inc., Washington, DC, 20008, U.S.A.;
[c]Nova Sensors, Inc., Solvang, CA, 93463, U.S.A.

ABSTRACT

It has been known since the early 1960s that hexagonal sampling is the optimal sampling approach for isotropically band-limited images, providing a 13.4% improvement in sampling efficiency over rectangular sampling. Despite this fact and other significant advantages of hexagonal sampling, rectangular sampling is still used for virtually all modern digital image processing systems. This is arguably due to the lack of an efficient addressing system for hexagonal grids. *Array set addressing* (ASA) is a recent advance in addressing hexagonal grids that allows image processing techniques to be performed efficiently on hexagonally sampled images. This paper will describe ASA and discuss its advantages. With ASA, a renewed interest in sensors that sample hexagonally is occurring. We will describe a new visible imager that simultaneously samples both hexagonally and rectangularly. This novel research tool has the ability to provide real imagery that can be used to quantitatively compare the performance of an image processing operation on both hexagonally sampled and rectangularly sampled images. We will also describe current efforts and plans for future visible sensors that sample hexagonally. The advantages of hexagonal sampling are not limited to the visible domain and should be equally realizable in the infrared domain. This paper will discuss considerations for developing infrared sensors that sample hexagonally. On-focal plane array (FPA) processing, readout architectures, detector materials, and bump-bonding are among the topics to be discussed.

Keywords: Hexagonal grid, hexagonal imaging, hexagonal addressing, hexagonal sensors, infrared FPAs, on-FPA processing

1. INTRODUCTION

Historically, both visible and infrared focal plane array (FPA) sensors have used a rectangular arrangement of square or rectangular pixels. We presume that this grew out of the convenience of such regular spatial organizations, combined with the accepted practice of using such spatial organization for image display systems, frame grabbers, etc. While convenient, such rectangular organization does not reflect the "low entropy" hexagonal organization that is typically found in biological systems. Such hexagonal "close pack" arrangements can be found in a variety of biological examples such as insect compound eyes, the cones of the primate retina, and honeybee and wasp combs. Some of these systems are thought to have evolved so as to maximize the use of available space. However, in the case of biological vision systems, the hexagonal arrangement of sensing elements also improves performance, since it is known that hexagonal sampling is the optimal approach to sampling isotropically band-limited two-dimensional signals. Hexagonal sampling also provides other advantages over rectangular sampling such as equidistant spacing, consistent connectivity, greater angular resolution, and a higher degree of symmetry. These advantages, combined with the recent development of array set addressing, are leading to the development of novel imaging sensors and systems that can efficiently perform complex real-time image processing functions.

*Send correspondence to: nicholas.rummelt@eglin.af.mil

Infrared Technology and Applications XXXVII, edited by Bjørn F. Andresen, Gabor F. Fulop, Paul R. Norton,
Proc. of SPIE Vol. 8012, 80120I · © 2011 SPIE · CCC code: 0277-786X/11/$18 · doi: 10.1117/12.896888

2. ARRAY SET ADDRESSING (ASA)

One of the reasons rectangular sampling is used in virtually all digital imaging systems to date is that rectangular sampling leads to rectangular arrays which can be easily represented and indexed in digital hardware, whereas hexagonal sampling does not. Various addressing schemes[1–4] have been introduced to try to remedy this fact, but none of these approaches can match the efficiency of the rectangular array. Array set addressing (ASA) solves this problem by representing the hexagonal grid as a set of two rectangular arrays which can be individually indexed by integer-valued row and column indices. The two arrays are distinguished using a single binary coordinate so that a full address for any point on the hexagonal grid is uniquely represented by three coordinates

$$(a, r, c) \in \{0, 1\} \times \mathbb{Z} \times \mathbb{Z} \tag{1}$$

where the coordinates represent the array, row, and column respectively. Figure 1 shows an example of how a 4×4 hexagonal array is represented using ASA.

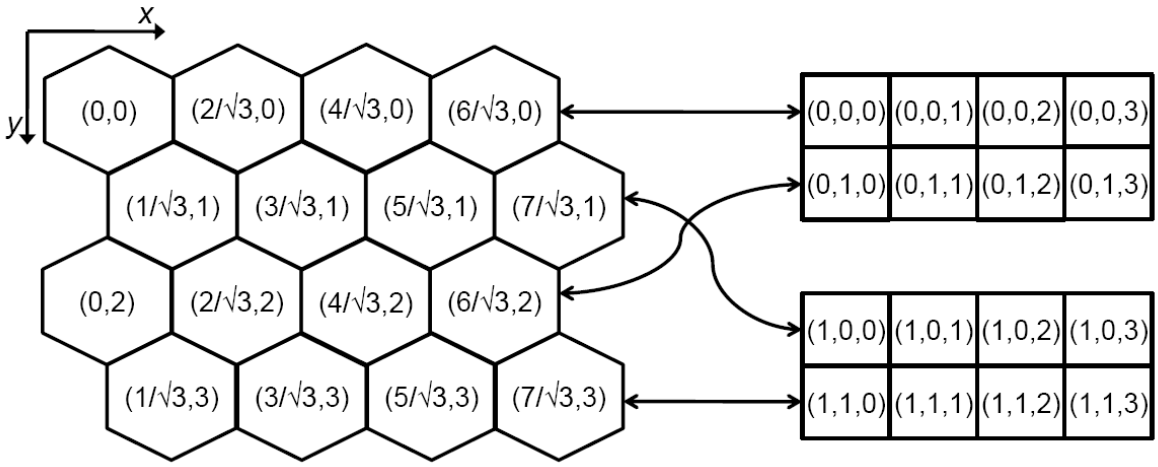

Figure 1. Example of how the hexagonal grid is separated into two arrays using ASA.

Rummelt developed ASA[5] and showed that hexagonally sampled images can be efficiently stored and processed on existing digital hardware. He developed efficient vector operations and proved that ASA satisfies the properties of a module over the ring of integers. Operations on ASA images such as coordinate conversion, Euclidean distance, city-block distance, convolution, gradient estimation, edge detection, downsampling, wavelet decomposition, and Fourier transform were developed and shown to be efficient. In fact, it turns out that the hexagonal Fourier kernel becomes separable in ASA, leading directly to a radix-2 decimation hexagonal fast Fourier transform (HFFT).

This elegant addressing scheme significantly outperforms other hexagonal addressing approaches. For example, HIP is a hexagonal addressing approach that was developed by Middleton and Sivaswamy[1] and is a variant of the generalized balanced ternary (GBT) system developed by Lucas and Gibson.[6] Some consider these "spiral" addressing approaches to be the natural way to address the hexagonal grid. However, a direct comparison between ASA and HIP was possible for several basic operations due to the Python source code for the HIP operations being provided in Middleton's book. The same operations were coded up in Python using ASA and the run times were compared. The result was that the ASA operations ran about 550 times faster than the HIP operations, on average.

The development of this truly efficient addressing scheme for hexagonally sampled images is beginning to spur a renewed interest in the development of sensors that sample hexagonally. Because of the fundamental nature of image sampling and addressing, the advantages to be gained by sampling hexagonally are independent

of the signal frequency band. Imaging applications in all frequency bands (e.g. visible, infrared, x-ray, etc.) can now exploit the advantages of hexagonal sampling without the burden of a computationally complex addressing scheme.

3. HEXAGONAL SENSORS

3.1 The Hex-Rect Imager

In order to compare the performance of hexagonal pixel arrays with rectangular pixel arrays, it is preferable to have a single image sensor that supports both modes. Placing both modes into one camera eliminates variables such as differences in silicon, readout circuitry, optics, support electronics, and other factors that may occur if hexagonal and rectangular arrays are placed in separate cameras. Below is a basic pixel geometry conceived

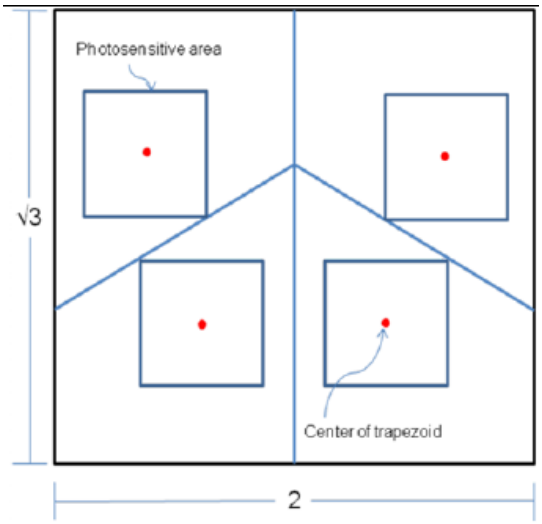

Figure 2. Trapezoidal pixel tile structure.

by author Rummelt and executed as an image sensor design by author Barrows. Figure 2 depicts a basic cell comprising four trapezoidal shaped "raw pixels" placed together to form a rectangle. Note the rectangle's aspect ratio of $\sqrt{3}/2 \approx 0.866$.

Figure 3 depicts the schematic diagram of a single pixel. There are four of these circuits in the cell shown above in Figure 2. Each pixel is a variation of a classic three transistor "active pixel sensor" in which one transistor (M1) is used for resetting a pixel value while two transistors (M2 and M3) are used for readout. Diode D1 is a photodiode between N-diffusion and the P-substrate, and is drawn so that it's centroid is located at the centroid of the trapezoid. Each pixel circuit also has one or two switching transistors (M4 and M5) that may be used to short out the pixel circuit with an adjacent pixel circuit.

This pixel circuit may be operated in a linear mode or a logarithmic mode according to how light intensity is expressed as a voltage. In the linear mode, the row reset line is pulsed high, which charges the parasitic capacitance at the top side of D1 and the gate of M2 to a positive voltage. As light strikes D1, the charge drops. The pixel circuit may then be read out by setting the row select line high to close M3, at which point transistors M2 and MC (there is one MC per column of pixels) form a source follower and read out the pixel potential. The readout voltage is a function of both the light striking D1 and the time interval between when the row reset line is pulsed and when the pixel is read out.

This pixel circuit may also be operated in a logarithmic mode by setting the row reset signal to a digital high everywhere across the chip. In this case current flows through M1 and D1 to ground in accordance with the light striking D1. This current is typically in the nanoamp or picoamp range, thus M1 will be biased in

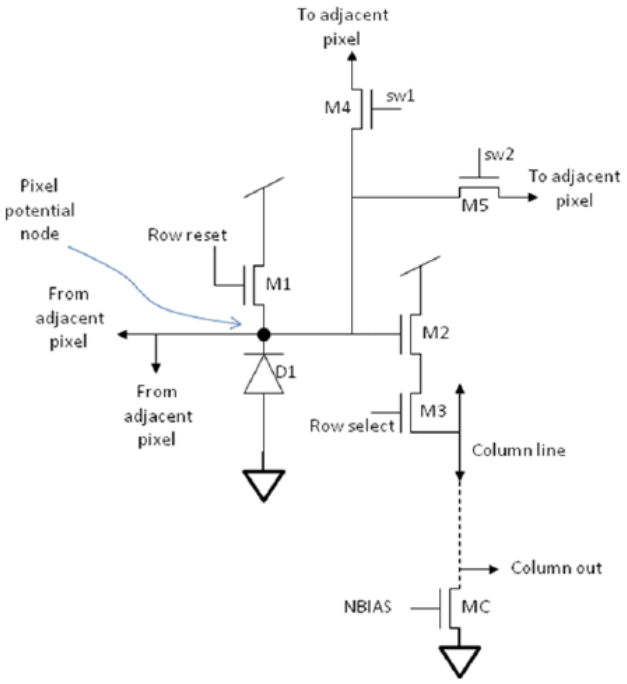

Figure 3. Conceptual schematic diagram of a single pixel.

the subthreshold region. The voltage drop across M1 is thus a logarithmic value of the light striking D1. The logarithmic mode has the advantage that many orders of magnitude of light levels may be handled with one setting, but has the disadvantage of a significantly reduced contrast sensitivity.

Figure 4 shows how the four pixel cell of Figure 2 may be combined to form a larger array, and how the trapezoidal raw pixels may be used to form either rectangular or hexagonal "super pixels". The left side of Figure 4 shows rectangular super pixels that may be formed by shorting together all four of the raw trapezoidal pixels of Figure 2 by turning on the appropriate M4 and M5 transistors. This causes the four corresponding D1 photodiodes to be shorted together and thus provide the same output. Note that in the rectangular mode, the super pixels have a $\sqrt{3} : 2$ aspect ratio rather than a perfect $1 : 1$ ratio, thus the pitch between super pixels is different for the horizontal and vertical directions. Similarly the right side of Figure 4 shows hexagonal super pixels that may be formed by shorting together the raw pixels in a different pattern, by turning on a different set of M4 and M5 transistors. The focal plane array depicted in Figure 4 may thus be operated in three array modes (which may be chosen orthogonally to the linear and logarithmic modes described above). Figure 5 shows the actual layout of the focal plane as seen by the CAD tools.

One may configure the focal plane to read out either rectangular super pixels or hexagonal super pixels. Note that only one mode may be used at a time. It is possible to alternate between the two modes every frame, however the resulting video sequences for the hexagonal and rectangular configurations will be offset in time by the time duration of one frame readout. It is also possible to turn off the M4 and M5 switches and read out the raw trapezoidal pixel array. In this case, the super pixels may be formed in software by arithmetically averaging the four corresponding raw pixel values. The averaging operation may be performed quickly by adding the four values and shifting right by two bits. The resulting hexagonal and rectangular super pixel arrays will then be perfectly synchronized in time. The disadvantage of this method is that four times as many raw pixels need to be digitized as either a rectangular or hexagonal array frame, thus any bottlenecks in the analog conversion are amplified. Table 1 shows the basic specifications and details of the Hex-Rect chip.

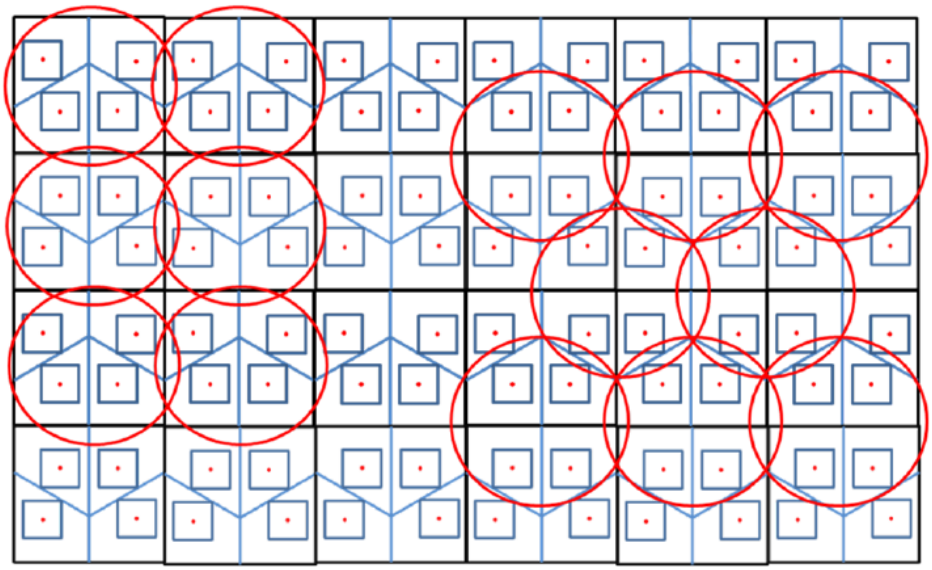

Figure 4. Forming rectangular and hexagonal super pixels from the raw trapezoidal pixel array. Red circles show the four trapezoidal pixels shorted together to form rectangular pixels (left) or hexagonal pixels (right).

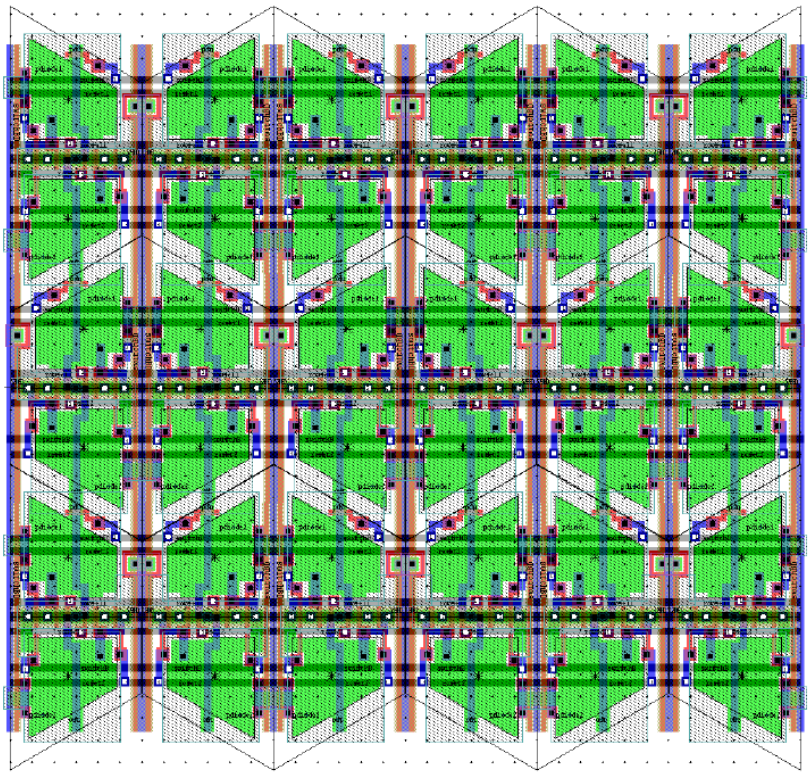

Figure 5. Layout of a section of the focal plane. The green areas are the locations of the trapezoidal shaped photodiode.

Drawn chip size	6.1mm x 11.1mm
Focal plane size	4.7mm x 9.2mm
Focal plane resolution	Raw trapezoid pixels: 304 x 512 Hexagonal array: 152 x 255 (even rows have 256 hex pixels) Rectangular array: 151 x 256
Pixel type	3-transistor active pixel, with support for both logarithmic response and linear response
Pixel pitch	18 microns wide by 15.6 microns high for raw pixels
Post-pixel circuitry	8-bit flash ADC
Interface	PIO12B parallel interface: 8 bidirectional digital, 2 digital in, 1 analog out 12-bit command bus in two 6-bit words 8-bit digital out Optional 3 input chip select Optional analog out Alternative 12 bit input / 8 bit output parallel interface
Process	ON-Semi C5N 3 metal 2 poly 0.5 micron process
Chip operating voltage	4V to 5V preferred
Digital input 0/1 threshold	About 0.95V
Voltage regulation	On-chip voltage regulator for analog circuits and bias generators

Table 1. Hex-Rect Specifications.

As an example of the types of experiments that can be conducted with the Hex-Rect imager, consider presenting a series of radial sine waves to the imager, with the frequency increasing as the experiment progresses. The reason for using a radial sine wave is that the magnitude of its Fourier transform is a circle in Fourier space with the radius of the circle being proportional to the frequency of the sine wave. This characteristic allows one to visually determine the frequency at which aliasing begins to occur (when the circles begin to overlap in Fourier space). Figure 6 shows a few of the frames from such an image sequence being presented to the Hex-Rect camera. The results clearly show that the hexagonally sampled images have a higher bandwidth than the rectangularly sampled images. The frequency at which the circles in Fourier space begin to touch is the onset of aliasing and hence the bandwidth of the sampling approach. Notice that the ratio of bandwidths is very close to the theoretical limit, $0.268/0.309 \approx 0.867 \approx \sqrt{3}/2 \approx 0.866$. Therefore, this experiment demonstrates that one can realize the full benefit of hexagonal sampling with real imagers and improve bandwidth by approximately 13.4%.

3.2 Plans for Future Hexagonal Imagers

In the past, there have been a few attempts to build imagers that sample hexagonally such as Carver Mead's Silicon Retina,[7] and the prototype described in Hauschild's paper,[8] but there are no commercially available hexagonally sampling imagers. Gaber et al.[9] produced a design for a hexagonal thermal infrared (IR) imager that exhibited very good performance characteristics, but it seems the design has not yet been fabricated. Although the Hex-Rect image sensor chip is suitable for studying the performance of hexagonally sampled images vs. rectangularly sampled images, the additional circuitry to support both modes adds overhead to the focal plane. Centeye is currently working on two pure hexagonal image sensors for mass production, both for a silicon process and both for usage in visible and near infrared light. The first hexagonal image sensor chip will be a low resolution version suitable for peripheral sensing and may be operated by a microcontroller or a DSP. The second hexagonal image sensor chip will be a higher resolution version containing hardware support for smoothing and downsampling to enable variable acuity image sensing, and will thus be suitable for foveating or zooming in on areas of interest. Both image sensor chips are expected to be available by late 2011.

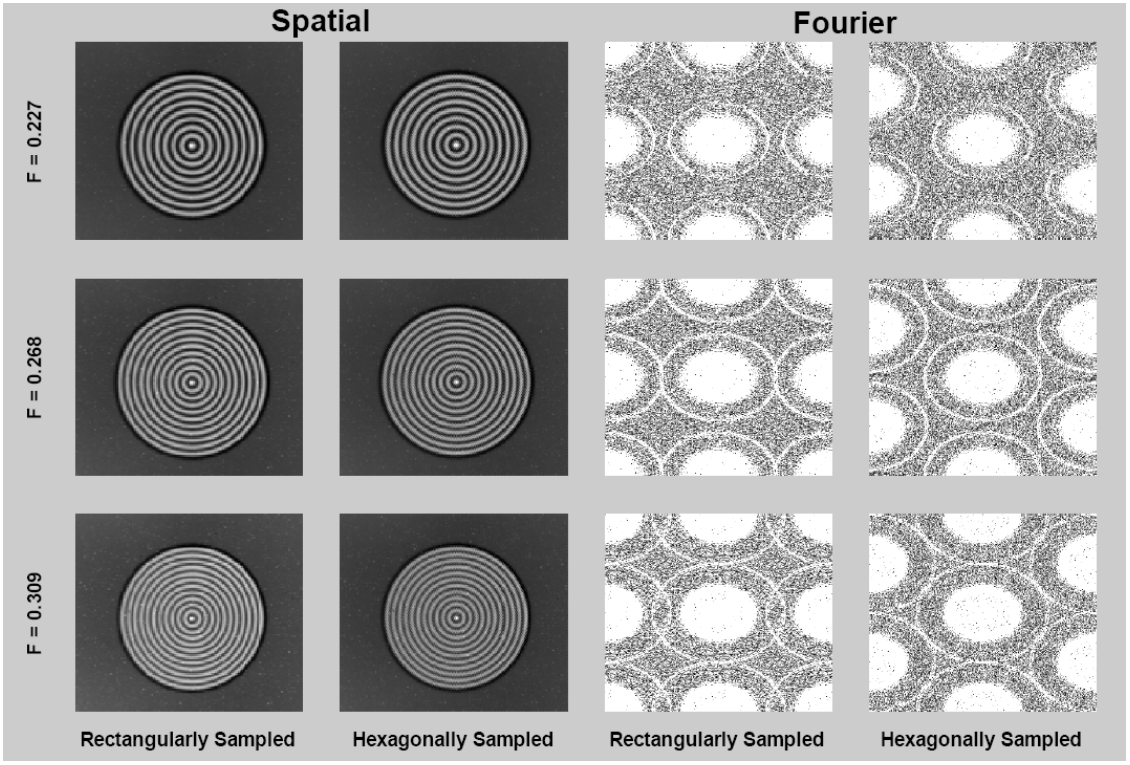

Figure 6. Experimental results from the Hex-Rect imager. Frequency (F) is given in radians per pixel.

3.3 Infrared Hexagonal Imagers

Mead's silicon retina includes a "Horizontal Resistive Layer" that is integrated into the hexagonal pixel architecture and patterned after the horizontal cells in the retina of vertebrate animals. Each photoreceptor in the network is linked to its six neighbors with resistive elements, to form the hexagonal array as shown in Figure 7. Each node of the array has a single bias circuit to control the strength of the six associated resistive connections. The photoreceptors act as voltage inputs that drive the horizontal network through conductances. As Figure 7 shows, the horizontal network of hexagonal resistors is a flat mesh of effective resistances (actually implemented in CMOS through the use of switched capacitors). This method provides a nearly ideal means to derive a reference with which local signals can be compared. As Mead indicates, this hexagonal network is particularly attractive because it has the highest symmetry and connectivity of any regular, two-dimensional structure. This network makes use of resistive connections R between neighboring nodes and a conductance G (not shown) connected from each node to ground. At each node, the pixel circuitry injects an amount of current that is proportional to the in-band illumination present at the pixel.

In 1991, Massie, et al.[10] implemented such a horizontal resistive layer using a rectangular pixel geometry and a network of switched capacitors in the development of the "Neuromorphic Infrared Sensor" (NIFS) and produced the first working IR system that offered logarithmic sensitivity, local contrast enhancement and tolerance to a very wide range of absolute illumination levels. Much like the silicon retina, this device electronically blurred the image and used this information in a subtraction operation with the unblurred image to implement a "Difference of Gaussians" (DoG) algorithm on the real-time data. Although not optimal, the NIFS(shown in Figure8) used a rectangular arrangement of pixels for this analog computation. The log conversion of intensity at the photoreceptors increases incoming dynamic range and low level sensitivity. High level sensitivity may be sacrificed, but this is not as relevant since the DoG filter highlights local contrast differences and not absolute intensities. Thus the available biochemical dynamic range for vertebrate retinas is reserved for the more pertinent information contained in the local contrast differences. The inhibition of information redundancy by certain retinal cells (amacrine) provides for more efficient data information packing.

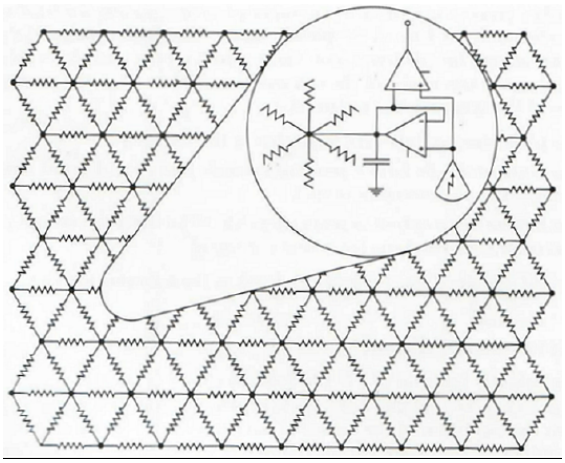

Figure 7. The "Horizontal Resistive Layer" developed by Mead et al. at the California Institute of Technology.

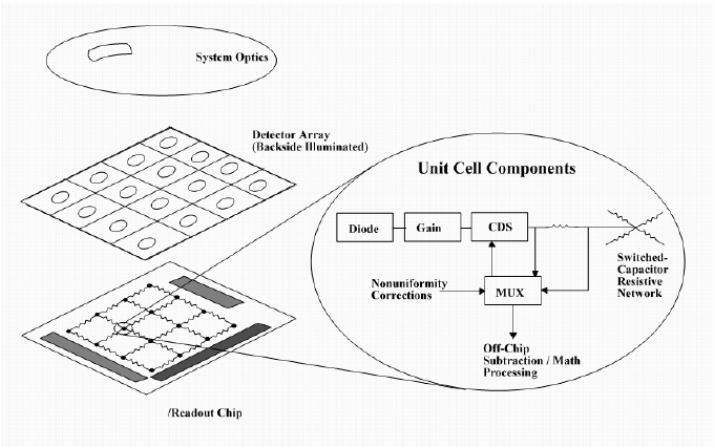

Figure 8. The Neuromorphic Infrared Sensor (NIFS).

3.3.1 Applications of the DoG in Infrared Imaging

The authors have made use of the DoG, either directly on-FPA or in post-processing hardware, to the benefit of numerous infrared systems. Hexagonal imaging architectures making use of such DoG techniques are expected to produce even better results than those achieved with rectangular pixel arrangements. Listed here is a brief summary of some implemented applications.

Local Contrast Enhancement, Zero-Crossings and Motion Detection. A visual representation of local contrast enhancement produced by the DoG algorithm is shown in Figure 9. A frame of infrared imagery of a hummingbird in flight was processed using the DoG technique as described above to indicate that low-contrast image variations present in the original image are clearly visible with higher contrast in the DoG-processed image. Not only are the bright (hot) portions of the image retained near the deltoid musculature and near the eyes, but the lower-contrast variations are enhanced to show individual feathers and texture present in the body of the hummingbird. Radiometric accuracy is lost through this image transformation, but the essential information represented by texture variations and heat centers are preserved. A variety of researchers have identified the

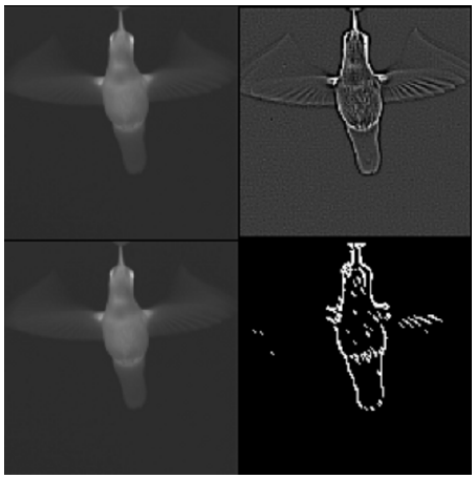

Figure 9. The DoG response of a hummingbird shows enhanced edge contrast as shown in the upper-right panel (input scene is in the upper-left). The zero-crossing image as shown in the lower-right image preserves the essential geometries in the scene, but conveys the information in a 1-bit representation. (Infrared imagery courtesy of the SE-IR Corporation, Goleta, CA)

capability of biological systems to detect when a particular sensor-produced signal crosses from the "positive" to the "negative" (or vice-versa) signal domain. Such a "zero-crossing" may be detected and used to identify the precise edge of an object within the field of view. The zero-crossing operation is naturally applied after a DoG, because the DoG produces zero-mean image data (i.e., approximately half of the DoG image distribution is in the positive domain, half is in the negative domain).

To demonstrate how the zero-crossing image preserves the spatial integrity of the image, Figure 9 compares the original input scene (left-most images) with the DoG image (upper right) and the zero-crossing image (lower right). Many algorithms that follow depend upon the reduced size of the zero-crossing image dataset. In his chapter on "Zero-Crossings and the Primal Sketch", Marr[11] lays the foundation for a description of how zero-crossings are used in biological vision systems. Following this, his chapter on "Directional Selectivity" makes use of such zero-crossings in order to perform "directional selectivity" in order to separate surfaces which are moving independently of one another.

Marr makes use of the concept that if a zero-crossing is produced, and if the vision system producing the zero-crossing representation has the ability to store at least one previous frame of information, then the system has the ability to infer the direction of motion of the zero-crossing edge. "Selectivity" may be utilized such that the imaging system simply does not report edges which are moving in certain user-selected directions. This may be shown to be a very powerful technique if, for example, detections due to background motion in the scene must be eliminated. If the direction of the background motion relative to the stationary sensor is known, such a technique could completely eliminate the background's motion from the resulting imagery. Directional selectivity should be more accurate in a hexagonal sampling domain due to the higher angular resolution of the hexagonal grid.

Edge Extraction. DoG filtering followed by a simple positive and negative number range threshold operation around the zero gray level results in a very effective edge extraction technique. An example of this from Marr is shown in Figure 10. In this example, the intensity of the lines has been made to vary with the slope of the zero-crossing (i.e., the magnitude of the contrast difference across particular edges within the scene). Implemented with a hexagonal sensor architecture, this approach would offer increased spatial fidelity over rectangular-based imagery. Rummelt[5] showed that edge detection performance improved in the hexagonal domain when using a Canny[12] edge detector.

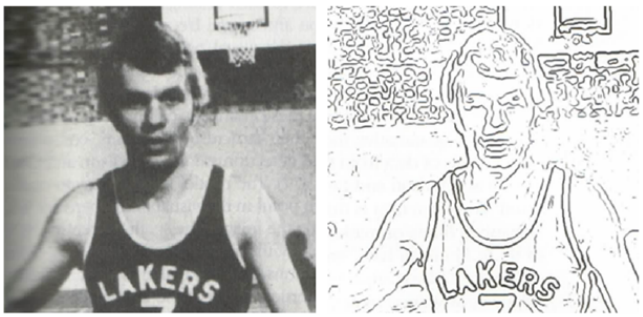

Figure 10. Example of zero-crossings and edge extraction; here, the intensity of the lines have been made to vary with the slope of the zero-crossing, so that it is easier to see which lines correspond to the greater contrast.

Anisotropic Filtering. We have used the DoG as a means to identify regions within imagery containing "salt and pepper noise" that have high relative contrast. These "edge regions" define those portions of the image over which two-dimensional pixel averaging is restricted. Lower-contrast regions (in this case containing the salt/pepper noise) become spatially averaged to a user-defined radius but the averaging does not cross the high contrast edges. Referred to as "anisotropic filtering", this technique may be implemented directly on the readout integrated circuit (ROIC). An example of this is shown in Figure 11.

3.3.2 Infrared Readout Architectures

The most typical infrared readout integrated circuits (ROICs) make use of horizontal and vertical output multiplexing architectures that presuppose the existence of a rectangular pixel organization. These typical devices read out a complete frame of information at a fixed rate and the resulting rectangular images are subsequently stored in memory and/or displayed to the user. Modifications from this typical approach have been developed over time to incorporate single "windowing" ROIC architectures[13] as well as multiple windowing "foveating" architectures.[14] In all of these cases, rectangular readout techniques have been considered. However, straightforward modifications to the typical architecture should enable the pixel data to be read out directly into the ASA data structure, enabling efficient processing. In addition, the recent advent of on-ROIC analog to digital converters may also be applied to the digitization and output multiplexing of the hexagonal arrays. Having

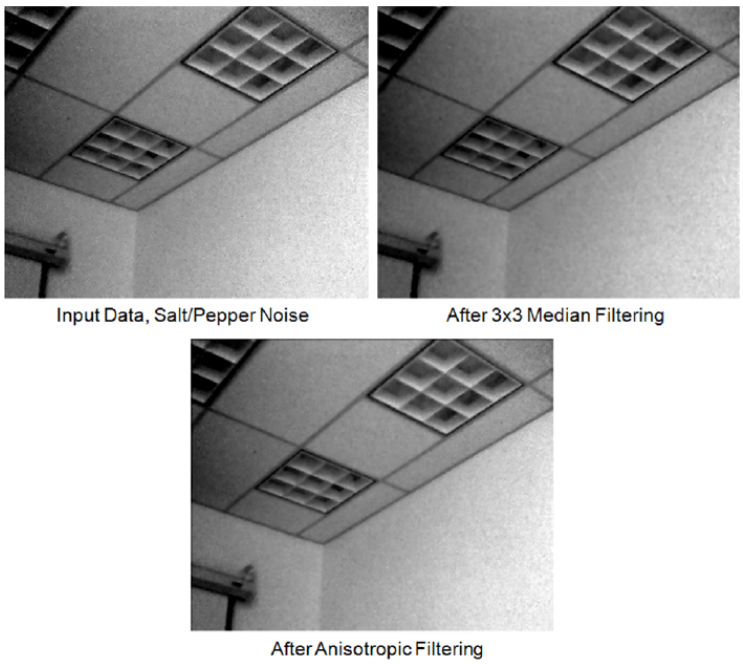

Figure 11. DoG is used to compute an internal edge representation, followed by anisotropic filtering that preserves edge resolution while averaging lower contrast image data. Salt/pepper noise, when filtered with the DoG-based anisotropic filter, produces better results than a conventional 3x3 median filter that tends to blur across edges and produce more mottled uniform grey imagery.

digital data present directly on the ROIC will enable future devices to be developed that implement many of the pre-filtering and pre-processing functions that we have described in this paper in the digital domain without the potential of further contaminating pixel signals with noise.

3.3.3 Infrared Detector Materials and Hybrid FPA Bump Bonding Considerations

The hybrid FPA structure shown in Figure 8 offers great flexibility in terms of the specific spectral band over which an FPA (both conventional and those using hexagonal pixel architectures) may operate. The underlying CMOS ROIC supports the operation of the overlaying photovoltaic photodetector array, supplying it with necessary bias voltage and collecting resulting photocurrent.

Typically, near infrared (NIR, 0.4 to 1.6μm) operation is supported by indium gallium arsenide (InGaAs) photovoltaic detector material. It may operate at room temperature but its dark current performance is improved substantially by cooling it by 10 to 20°C. Midwave infrared (MWIR, 3.0 to 5.0μm) operation is typically supported by both indium antimonide (InSb) and/or mercury cadmium telluride (HgCdTe). These materials have historically been required to operate at 77K to manage their temperature-induced noise current (Johnson noise) but recent advances in HgCdTe technology has produced materials capable of operating at elevated temperatures (approx. 120K and higher) while maintaining signal and noise performance. Longwave infrared (LWIR, typically 8.0 to 12.0μm operation) is supported today by HgCdTe, quantum well infrared photodiode (QWIP) structures as well as strained layer superlattice (SLS) materials. Ongoing research is underway to further optimize the performance of these latter materials that offer the promise of lower cost, higher manufacturability and potentially higher performance operation as compared with HgCdTe.

All of these materials may be considered for use with specially-designed ROICs that support hexagonal pixel architectures for future applications. The process of bonding the detector array to the underlying CMOS ROIC is typically accomplished by using an indium "bump bonding" approach in which mating indium bumps on both structures are pressed together, forming a reliable cold-weld that can survive extreme thermal cycling. The

smallest indium bumps that have been reported in recent meetings have been on the order of 2μm in diameter but typical indium bump diameters have historically been on the order of 5 to 6μm in diameter. The indium bump bonding technique will be able to accommodate a hexagonal arrangement of pixels.

4. CONCLUSION

This paper has introduced the array set addressing (ASA) method for indexing the pixels of hexagonally sampled imagery. This new technique is an enabling technology for the development of image processing systems that exploit the advantages of hexagonal sampling. Since previously known hexagonal addressing schemes are inefficient, there has been little motivation to develop sensors that sample hexagonally or displays that display hexagonally sampled imagery. We have described current work being done in the development of hexagonal visible sensors and have discussed some considerations for hexagonal infrared sensors. The authors hope that the current work leads to more excitement about hexagonal sampling and that the push to move beyond sub-optimal rectangular sampling continues.

REFERENCES

[1] Middleton, L. and Sivaswamy, J., [*Hexagonal Image Processing: A Practical Approach*], Springer, London (2005).

[2] Snyder, W. E., Qi, H., and Sander, W., "A coordinate system for hexagonal pixels," in [*Proc. SPIE Medical Imaging: Image Processing*], **3661**, 716–727 (1999).

[3] Her, I., "A symmetrical coordinate frame on the hexagonal grid for computer graphics and vision," *J. Mech. Des.* **115**(3), 447–449 (1993).

[4] Sheridan, P., *Spiral Architecture for Machine Vision*, PhD thesis, Univ. of Technology Sydney (1996).

[5] Rummelt, N. I., *Array Set Addressing: Enabling Efficient Hexagonally Sampled Image Processing*, PhD thesis, University of Florida (2010).

[6] Lucas, D. and Gibson, L., "A system for hierarchical addressing in Euclidean space," Interactive Systems Corporation (Jan. 1980).

[7] Mead, C., [*Analog VLSI and Neural Systems*], Addison-Wesley, Reading, MA (1989).

[8] Hauschild, R., Hosticka, B. J., Müller, S., and Schwarz, M., "A CMOS optical sensor system performing image sampling on a hexagonal grid," in [*Proc. 22nd European Solid-State Circuits Conf.*], 304–307 (1996).

[9] Gaber, H., El-Semary, H., and Ali, H., "Hexagonal shaped microcantilever based IR imaging sensor," in [*2009 International Conference on Microelectronics (ICM)*], 370–373 (Dec. 2009).

[10] Massie, M. A., Woolaway, J. T., Curzan, J. P., and McCarley, P. L., "Neuromorphic infrared focal plane performs sensor fusion on-plane local-contrast-enhancement spatial and temporal filtering," in [*SPIE Visual Information Processing II*], Huck, F. O. and Juday, R. D., eds., 160–174 (1993).

[11] Marr, D., [*Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*], W.H. Freeman (1982).

[12] Canny, J. F., "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-8**, 679–697 (1986).

[13] McCarley, P. L., Massie, M. A., Baxter, C. R., and Huynh, B. L., "Neuroseek dual-color image processing infrared focal plane array," in [*Proc. SPIE Infrared Readout Electronics IV*], Pain, B. and Lomheim, T. S., eds., **3360**, 13–27 (1998).

[14] McCarley, P. L., Massie, M. A., and Curzan, J. P., "Foveating infrared image sensors," in [*Proc. SPIE Infrared Systems and Photoelectronic Technology II*], Longshore, R. E., Sood, A. K., Dereniak, E. L., and Hartke, J. P., eds., **6600** (2007).