Performance and Application of Real-Time Hyperspectral Imaging

Mark Dombrowski

Surface Optics Corporation
11555 Rancho Bernardo Road, San Diego, CA 92127-1441

Paul Willson, Ph. D.

U.S. Army Armament Research and Development Engineering Center
SMCAR-QAH-T, Picatinny Arsenal, New Jersey 07806-5000

Clayton LaBaw

Jet Propulsion Laboratory
4800 Oak Grove Drive, Pasadena, CA 91109-8099

ABSTRACT

Hyperspectral imaging is the latest advent in imaging technology, providing the potential to extract information about the objects in a scene that is unavailable to panchromatic imagers. This increased utility, however, comes at the cost of tremendously increased data. The ultimate utility of hyperspectral imagery is in the information that can be gleaned from the spectral dimension, rather than in the hyperspectral imagery itself. To have the broadest range of applications, extraction of this information must occur in real-time. Attempting to produce and exploit complete cubes of hyperspectral imagery at video rates, however, presents unique problems for both the imager and the proceessor, since data rates are scaled by the number of spectral planes in the cube. MIDIS, the Multi-band Identification and Discrimination Imaging Spectroradiometer, allows both real-time collection and processing of hyperspectral imagery over the range of 0.4 μm to 12 μm. Presented here are the major design innovations associated with producing high-speed, high-sensitivity hyperspectral imagers operating in the SWIR and LWIR, and of the electronics capable of handling data rates up to 160 megapixels per second, continuously. Discussion of real-time algorithms capable of exploiting the spectral dimension of the imagery is also included. Beyond design and performance issues associated with producing and processing hyperspectral imagery at such high speeds, this paper also discusses applications of real-time hyperspectral imaging technology. Example imagery includes such problems as detecting counterfeit money, inspecting surfaces, and countering CCD (camouflage, concealment, and deception).

Keywords:

Hyperspectral-imaging, real-time, visible, infrared, processing, spectral, correlation

1.0 DEMANDS OF HYPERSPECTRAL IMAGING

Real-time imaging over broad bands in the electromagnetic spectrum from the ultraviolet (UV) through the infrared (IR) has been a staple in the areas of remote sensing, surveillance, target detection and tracking, search and homing devices, spectrally tailored coating development, nondestructive inspection, and noninvasive diagnosis. Improvements are being made in these techniques all the time, with increased resolution, higher sensitivity, and greater information throughput being the benefit. The development of high speed digital processing hardware, first at the computer workstation level and later at the integrated circuit level, has spurred the development of more effective optical imaging systems, in part because it has taken up some of the burden formerly handled by analog devices.

Digital image processing paved the way for systems that not only view but analyze the scene in real time. By real time is meant "at video rates", or equivalently, at a rate such that the record of the scene evolves continuously before our eyes, with no gaps or jumps between successive snapshots or frames. Since the human visual system filters input information with a time constant down to approximately 0.1 sec, depending on light level, corresponding conservatively to a noise equivalent bandwidth of 5 Hz, then by the Nyquist criterion essentially maximal information transfer to the brain occurs at 10 Hz and above. In practice, visual perception of a "flicker effect" persists to somewhat higher frequencies, so one might set the threshold frame rate for real-time operation at approximately 20 Hz.

Such a system records the scene in terms of three of the independent variables in which it can be described: x (or azimuth), y (or elevation), and time. It may thus be termed a "3-d" system. Collecting spectrally resolved information over the broad band viewed adds a fourth dimension to our record of the scene, and provides a basis for a more discriminating analysis of the scene content. Associated with each pixel in the 3-d space defined above is a spectral breakdown of the arriving radiation for that pixel, divided into narrow, normally but not necessarily contiguous bands, several to hundreds in number. In modern parlance such a system is known as *multispectral* (few to several bands) or *hyperspectral* (many bands - which we might define here as a sufficient number of bands to record *all* the *useful* spectral variations to be found in the scene). Systems built on these principles become "4-d" imaging systems, sampling the scene in terms of wavelength as well as the other three variables.

As an outgrowth of the explosion of imaging technology mentioned at the outset, a number of diverse multispectral and hyperspectral imaging systems have been developed and fielded in recent years. The amount of information generated by these 4-d systems becomes so voluminous that their application to many practical tasks depends on the ability to process the data and make decisions on it in real time, that is at rates keeping pace with the real-time acquisition of images. The area of remote threat sensing and reconnaissance includes many such problems. At this point we clarify our definition of real time by asserting that one spatio-temporal frame, complete with all of its spectral points, should be collected at 20 Hz or more. For example, a system operating at 20 spatio-temporal frames per second, collecting 50 spectral bands for each of these, for a total of 1000 *images* per second, may be said to be real-time.

Since 1990, SOC has been engaged in the development of a real-time hyperspectral imaging system, known as MIDIS (multi-band identification and discrimination imaging spectroradiometer). With a combination of Government and in-house resources, a highly capable system has emerged and has been demonstrated in various application areas. A unique aspect of MIDIS is its ability to process the images *spectrally* in *real time*, from initial gain and offset compensation and spectroradiometric calibration through decision oriented processing including spectral correlation using varied algorithms and sensor emulation by spectral integration of the response of other sensors. The uniqueness stems from SOC not only being the first to develop a system collecting and processing spectral images at a 1,000 Hz rate, using custom developed processing boards, but also from the fact that the system, specifically its real-time spectral acquisition and processing capability, is covered by two U.S. patents (5,424,543 and 5,602,394). Figure 1-1 presents a representation of a visible-band hyperspectral image cube of a set of spectrally diverse objects, showing the differing contrasts at a few of the image planes, and illustrating the vast amount of data contained in a single hyperspectral image cube.

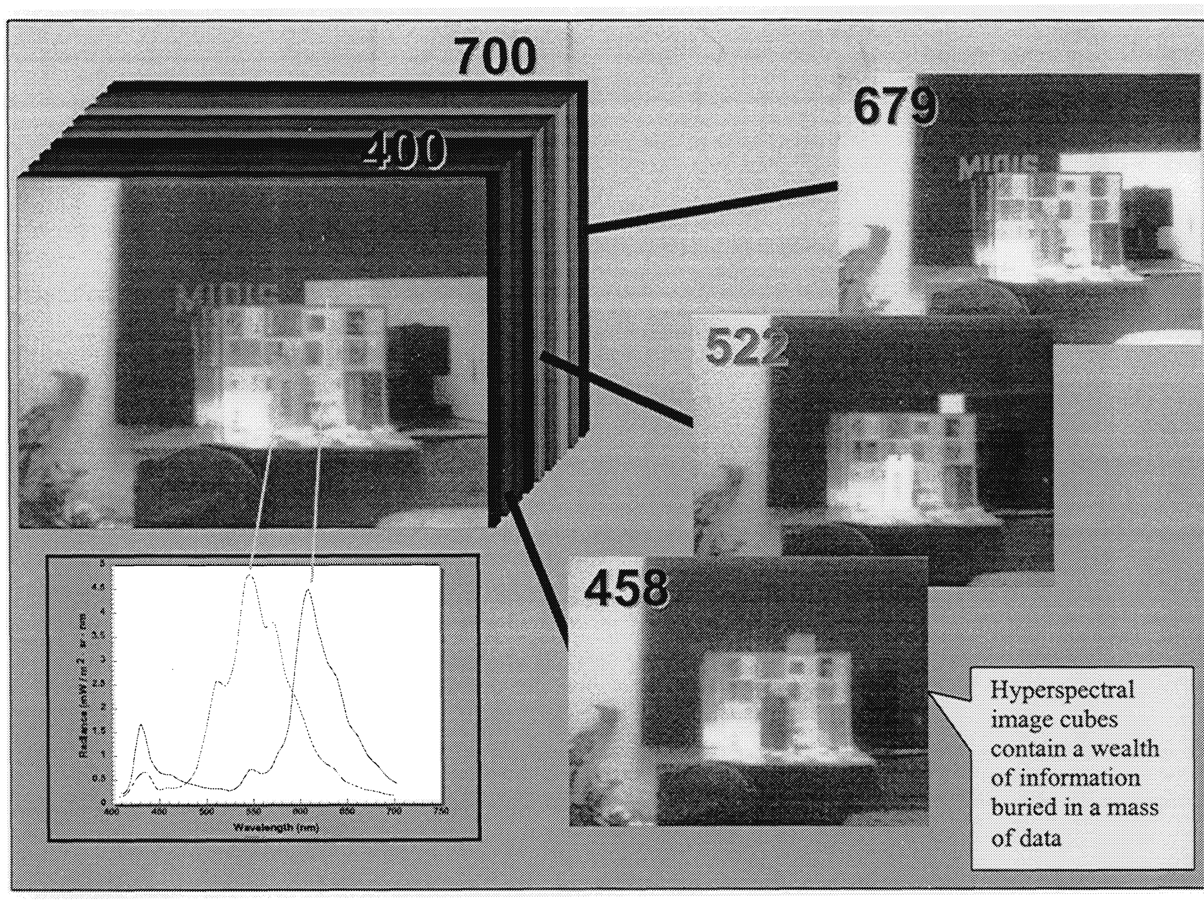

Figure 1-1. The hyperspectral image cube. With complete spectra at every pixel, a single cube of hyperspectral imagery can contain from a few megabytes to hundreds of megabytes. Timely extraction of the information content of the cube is the greatest challenge of hyperspectral imaging

The goal of hyperspectral imaging is to increase the amount of information available about a scene. Typically this means identifying objects or areas in a scene as a given material, or at least differentiating among objects based on material classes, without absolute identification of the material. Ability to perform such identification or differentiation stems from the fact that all materials posses relatively unique complex indices of refraction, giving rise to similarly unique reflectance, transmittance, and emittance. These optical properties, of course, are dependent both on chemical composition and physical structure, on both a microscopic and macroscopic scale. For instance, transmittance characteristics of water vapor, water, and ice are markedly different due both to the packing of the molecules and the increased effective mass loading on the O-H bonds as water moves from vapor to liquid to solid – a microscopic effect. Reflectance characteristics of snow and sheet ice, both the same material in the solid state, are also significantly different due to multiple scattering effects in the snow – a macroscopic effect. Figures 1-2 and 1-3 show the power of spectral correlation techniques to reveal otherwise hidden objects based on material differences.

Figure 1-2. Camouflaged Model Tank (left) viewed by VIS-NIR Instrument, Revealed by Real-time Correlation (right).

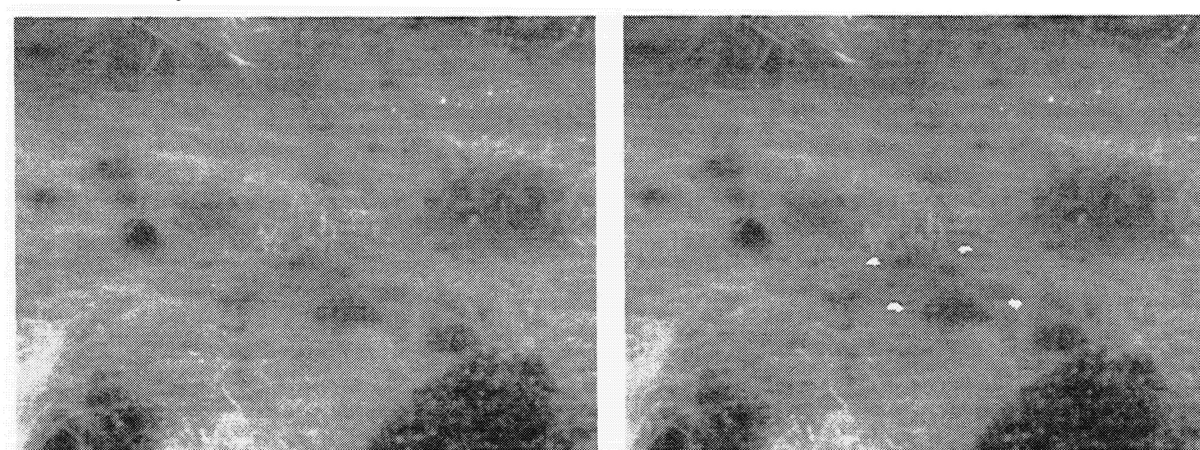

Figure 1-3. Cluttered scene (left) makes detection of embedded mines nearly impossible. Spectral Matched Filtering (right) reveals their locations.

As illustrated above, hyperspectral image data is only useful for the information that it contains. Extraction of the information, however, given the tremendous volume of data that must be processed is indeed a challenge. Information extraction has typically been a job that has taken weeks or months of processing on a workstation, and this for only a few image cubes. Such extended processing times are adequate to analyze data taken from static scenes. But to have the greatest range of applications, e.g., for this technology to find applications on a production line, in the operating room, or on a military vehicle, data collection and exploitation must take place in real-time.

This difference in approach is much akin to the difference between snapshot pictures and motion pictures. Whether collected by exposing film or using CCDs, the image generation technique is the same for snapshot and motion pictures. Only the speed at which the imagery is collected and displayed is different. The information content of motion pictures is significantly greater than that of snapshots, and along with this greater information content comes significantly greater numbers of applications. So it is with hyperspectral imaging. Most systems today are indeed snapshot hyperspectral imagers, capturing a view of a scene at a given instant in time, with no ability to continuously resolve temporal scene changes. Only MIDIS can truly be termed a hyperspectral video camera, capable of continuous real-time collection, exploitation, and presentation of hyperspectral imagery.

The demands on such a system are tremendous. Take for example the aforementioned cube of data comprising 50 images generated 20 times per second, or 1,000 images per second. If each image comprises 256 x 256 spatial samples of the scene, then the pixel generation rate is over 65 million pixels per second. If each pixel, after calibration, is resolved to 16 bits, then the total data rate is over 130 megabytes per second. This sustained data rate is beyond the ability of simple processors to handle, especially when considering that tens to hundreds of operations must be performed on each pixel in the cube to effectively extract the information. Even highly parallelled multi-processor digital signal processors require an excessive number of processors, translating to large size and high power consumption, to accommodate such a rate.

This is the problem that Surface Optics Corporation has sought to solve over its last seven years of involvement in the field of hyperspectral imaging. The resultant MIDIS system is unique not only in its ability to generate high-speed spectral imagery, but also in its ability to extract information from the spectral domain in real-time. In fact, the throughput rate of the processor, the MIDIS Millenium, is a full 320 megabytes per second, nearly a third of a gigabyte per second. With the memory and calibration section of the system occupying a single full-sized PC card, the pre-processor a second card, and the main processor a third, this massive computational power is extremely compact and low-power. The following sections describe this system, its performance, and its applications in further detail.

2.0 MIDIS FUNCTIONAL DESCRIPTION

Figure 2-1 presents a block functional diagram of a MIDIS system with complete coverage from 0.4 μm to 12 μm, delineating both the currently available portions of the system, and those under development. VIS/NIR, SWIR/MIWR, and LWIR spectral imagery is generated at up to 1000 frames per second (per head) by a CVF-based spectral imager. Raw imagery data is fed to a memory and calibration board, which uses two ping-pong image memory banks to allow continuous real-time collection and processing of hyperspectral imagery. The memory and calibration board also contains a calibrator circuit with associated calibration memories to translate the raw imagery to true spectroradiometric form. Calibrated imagery is then sent from two or more heads to the input multiplexer and re-mapping board, which selects which, and in which order, bands are sent to the pre-processor. Using a bilinear, random input address re-mapper, this board also corrects for any

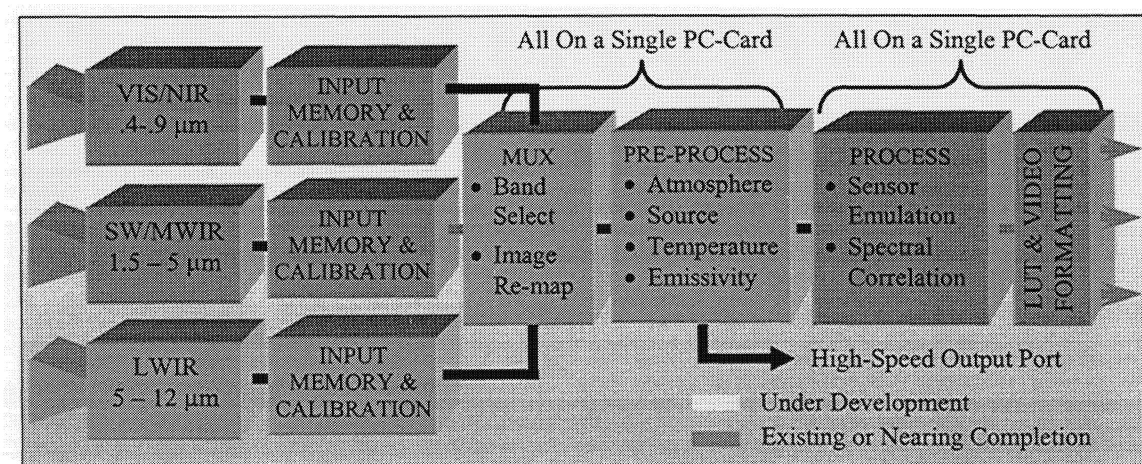

Figure 2-1. The MIDIS system. Hyperspectral imagery is generated by up to three CVF-based hyperspectral imagers. A high-speed processor accepts, calibrates, corrects, and exploits the 320 megabyte per second stream of data.

distortions or mis-registration from head-to-head. Re-mapped calibrated imagery is then sent to the pre-processor, part of the same PC board, which corrects the calibrated apparent scene radiances for atmospheric effects, including absorption, scattering, and path radiance, to produce inherent scene radiances. Using JPL, MIT Lincoln Lab, and SOC algorithms, this board also translates inherent scene radiances to scene reflectance or temperature and emittance (depending on the band). Corrected spectral imagery then flows to the main processor, which comprises three independent sensor emulation channels and a spectral correlation channel. Final imagery is output as standard video data, both to the host computer interface and to an external monitor, including a viewer in the imaging head.

At any point along the processing chain, data can be sent to the host computer for storage or alternate processing. The high-speed output port from the pre-processor allows data at any level of calibration/correction to be sent to a high-speed mass storage device, such as a RAID device. With such flexibility in data exploitation, MIDIS provides functionality both as a scientific instrument, and as an instrument suited for duty in a tank, an aircraft, in the hospital, or in the factory

Utilizing a circular variable filter (CVF) to perform spectral imaging, MIDIS collects a sequence of two-dimensional spatial images, each at a different wavelength. Figure 2-2 illustrates the spectral filtering technique, particularly the synchronization of array read-out with the filter rotation, which generates spectrally pure images. Imagery is converted to 12 bits resolution, ensuring wide intra-scene dynamic range. Variable integration time from one spectral image to the next guarantees maximum sensitivity, even with the wide inter-scene dynamic range seen by a HS imager. This type of full-frame hyperspectral imager is platform independent. That is, it can operate equally well on an aircraft, on a vehicle, held by a person, or sitting on a table. No relative motion between scene and imager is required, as is the case for push-broom imagers. This technique produces a "forward looking" hyperspectral video camera, as opposed to a "down-looking" hyperspectral snapshot camera, analogous to the difference between a FLIR and down-looking IR imager.

2.1 Sensor Performance

The SWIR/MWIR and LWIR spectral imagers utilize arrays designed specifically for MIDIS, exhibiting read-noise levels of only 120 e^-, maintaining high SNR even in the photon-starved 2-4 μm region. Figure 2-3 shows the high SNR achieved by MIDIS viewing a sunlit grass field at 300°K,

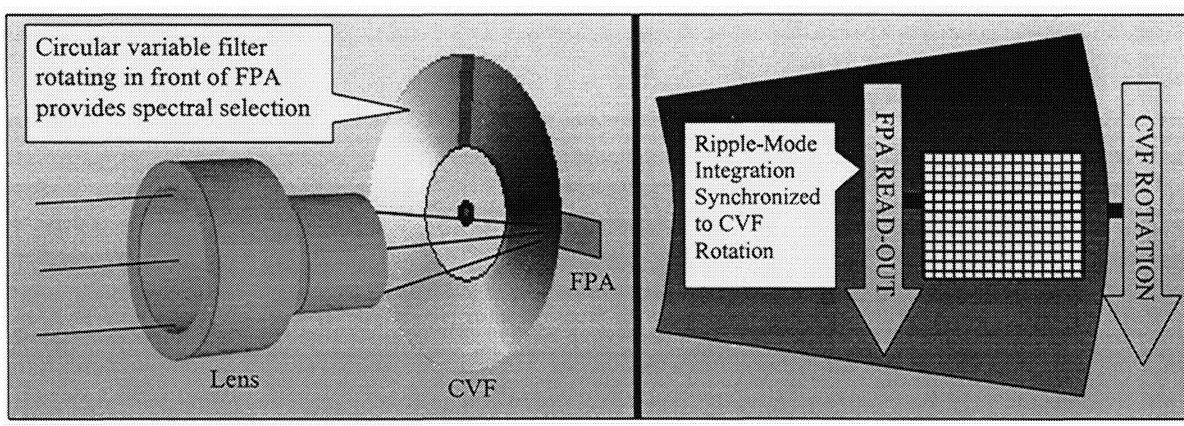

Figure 2-2. CVF-based HS imager- Ripple-mode array read-out synchronized to CVF rotation delivers high-speed, high spectral purity imagery.

from a distance of 300'. Note that outside the atmospheric absorption bands, SNR generally exceeds 150:1. SNR can be increased by reducing the image collection rate, allowing longer integration times at each image. Figure 2-3 also shows the effects in the SWIR/MWIR of changing the output cube rate from 30 cubes per second to 10 cubes per second, and finally to 1 cube per second. In order to achieve the high SNR shown in figure 2-3, cooling of the CVF is required to minimize its radiance, even though its emissivity is very low. Since the CVF lies immediately in front of the detector, unwanted background radiances are filtered to the same narrow bandwidth that scene radiances are, thereby dramatically reducing the shot noise associated with those background radiances.

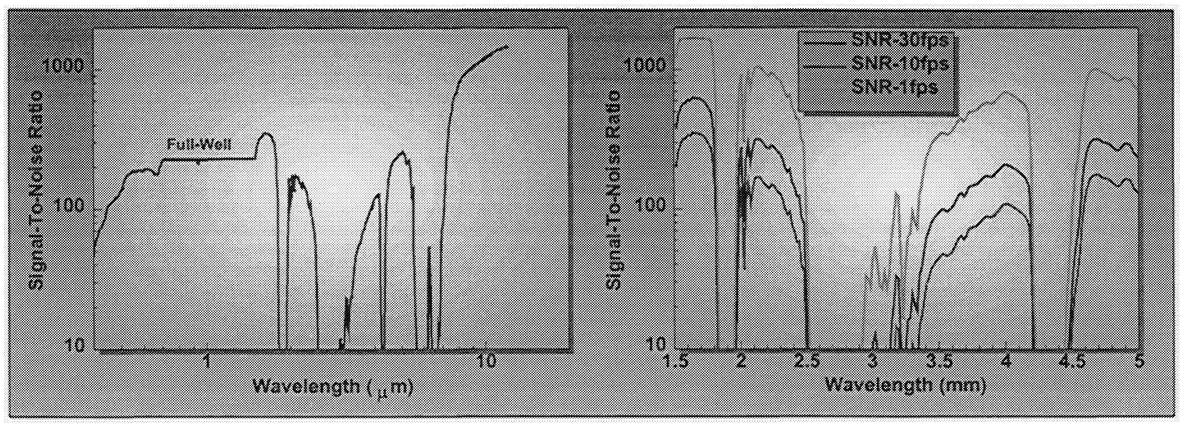

Figure 2-3. Spectral SNR – Left: MIDIS provides exceptional SNR across the entire spectrum, yielding high-quality imagery. Right: Variable cube collection rate allows trading temporal resolution for SNR

2.2 Processing Electronics

Generation of spectral images - as described above or by any other technique - is only a small part of the real goal of generating useful discriminants which can be used to classify, identify, track, or otherwise monitor a desired object. As alluded to earlier, data in general is only useful for the information it contains. The MIDIS processor is capable of accepting spectral imagery at rates up to 160 million pixels-per-second per head, for an aggregate input rate, if all three heads are attached, of 480 million pixels-per-second. At two bytes per pixel, this translates to an input rate of nearly one gigabyte-per-second. Such flood of data would be nearly useless without some way to effectively extract the information contained therein at the same rate that the data is collected.

To the end of extracting information from the spectral/spatial cube of data in real-time, MIDIS includes a very powerful real-time processor based on software-reconfigurable hardware architectures. Making copious use of high-density field-programmable gate arrays (FPGAs) and complex programmable logic devices (CPLDs), this processor runs at speeds unobtainable by microproccessors, unless they are massively paralleled with resultant size, weight, and power penalties. The processor first converts raw imagery to calibrated spectroradiometric imagery on the memory/calibration board. Output from multiple memory/calibration boards (when multiple heads are connected to the processor) is fed into the input multiplexer and image re-mapping board to select the bands from the multiple heads to send to the rest of the processor, and to re-map images into a common spatial reference frame. Re-mapping corrects for spatial mis-alignments from head-to-head, and allows correction of distortion from optics. Selected data is then fed to the pre-processor which first corrects calibrated imagery for atmospheric effects, and then translates corrected imagery into either reflectance (Visible-SWIR bands) or emittance and surface temperature(MWIR and LWIR bands). Corrected imagery then passes to the main processor, which implements three channels of response curve integration, and one channel of spectral correlation. A block diagram of this architecture was presented in Figure 2-1. Details of the processor are provided below.

Spectral response integration allows the instrument to emulate any desired SWIR/MWIR imager by integrating the measured spectral radiances with any arbitrary response curve, so that the output at a given pixel, SE(p), is given by

$$SE(p) = \int L(\lambda, p)\, R(\lambda)\, d\lambda$$

where R(λ) is the spectral response curve of the sensor being emulated. The MIDIS processor computes three such integrals simultaneously in real-time. An even more powerful processing operation implemented by the instrument is spectral correlation. Here, a desired relative spectral radiance to filter for is loaded into the instrument, either by direct measurement or through the software interface. Many different filtering algorithms with varying degrees of sensitivity then filter each pixel's spectral radiance against the sought after radiance, with matching pixels appearing bright, while mismatched pixels are dim. Multiple runs of the processing, each taking $1/30^{th}$ second for a 256 x 256 x 80 point cube, can extract the spectral content of the scene as shown in figure 2-4.

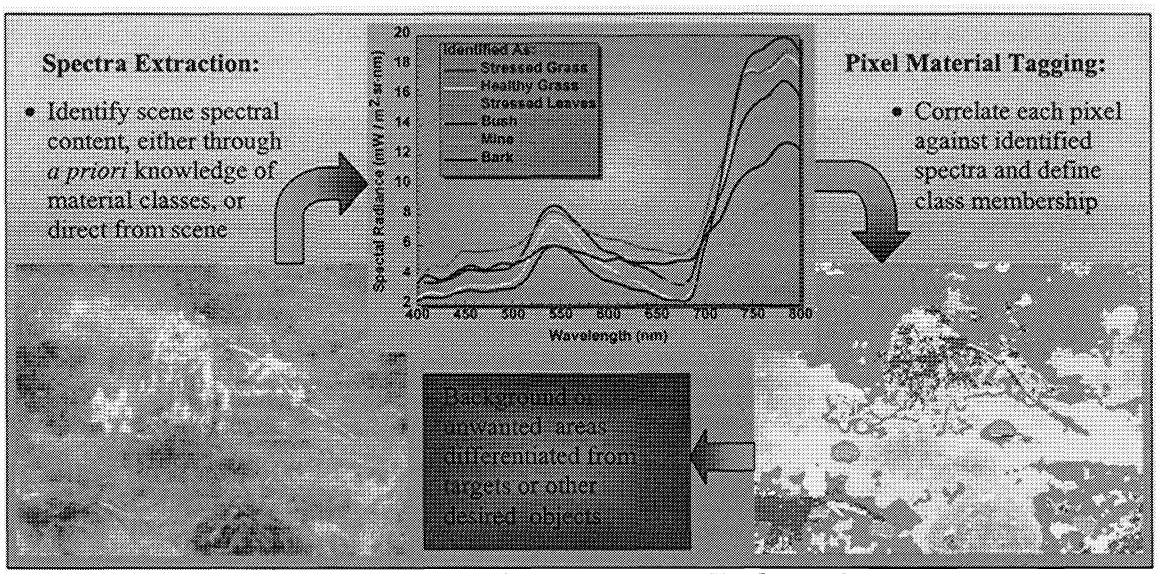

Figure 2-4. Material-Specific Scene Segmentation – Each pixel in an image is assigned to a material class based on spectral content. The MIDIS processor can identify 30 material classes per second from a 256 x 256 x 80 point cube

The development of spectral extraction/identification algorithms is an area of great activity and interest within the hyperspectral imaging community. Many high-performance algorithms have been developed, some operating to identify fully resolved objects without any knowledge of the background, and some operating to identify sub-resolved objects by either knowing or identifying both the background and target spectra. MIDIS Millenium, the hyperspectral image processor portion of the MIDIS system, implements a wide variety of these algorithms. An example of a background-independent algorithm designed to identify fully resolved (or slightly sub-resolved) objects is the zero-mean differential area correlator. This algorithm is formed by normalizing both the measured and filter radiances such that each encloses unity area, and subtracting the area "trapped" between the two curves from 1. For two identical spectral radiances, zero area will be "trapped", and the correlation value will be 1. As more and more area lies between the two curves, the correlation value will become

$$C_{ZMDA} = 1 - \sqrt{\sum\left[\frac{L(\lambda_n) - \langle L(\lambda_n)\rangle}{\sqrt{\Sigma\left(L(\lambda_n) - \langle L(\lambda_n)\rangle\right)^2}} - \frac{F(\lambda_n) - \langle F(\lambda_n)\rangle}{\sqrt{\Sigma\left(F(\lambda_n) - \langle F(\lambda_n)\rangle\right)^2}}\right]^2}$$

smaller and smaller. The zero-mean differential area correlator is given formally by

Where $L(\lambda)$ represents any given pixel's measured spectral radiance (or reflectance/emittance if so desired), $F(\lambda)$ represents the "filter" radiance, and <> represents the first moment or mean operator. Using the derivative of spectral radiance creates yet tighter filtering algorithms. Operation of the zero-mean differential area correlation algorithm is shown pictorially in figure 2-5

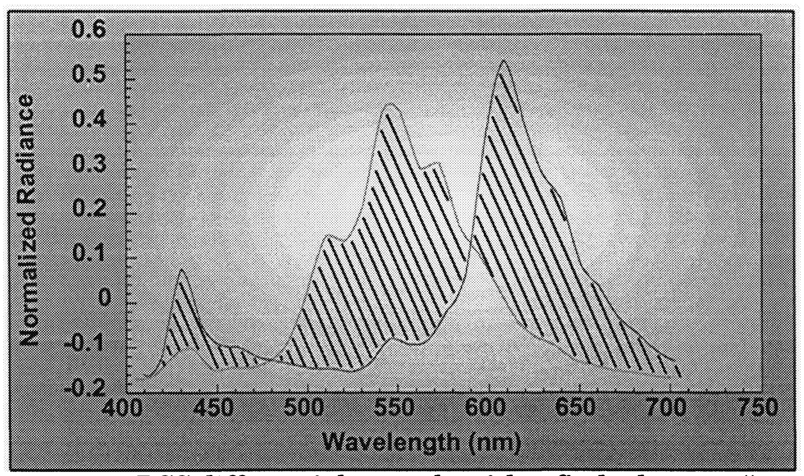

Figure 2-5. The zero-mean RSS differential area algorithm finds the area "trapped" between two normalized radiance curves

An example of a sub-pixel discrimination algorithm is the maximally noise-tolerant spectral matched filter, perhaps the simplest of this class of algorithm, given formally by

$$C_{MF} = \sum L(\lambda_n) \cdot M(\lambda_n)$$

where $M(\lambda)$ represents the matched filter, generated such that it is orthogonal to each of the basis elements found in the scene, but not orthogonal to the sought-after target. Since typical hyperspectral image sets contain upwards of a hundred bands, but typical scenes are adequately described by a handful of basis elements (materials), the system of equations used to define the matched filter is highly underconstrained. To adequately constrain the matched filter vector generation to produce the optimum vector, the final constraints come from requiring that the filter vector produce minimum residual in the presence of the noise and spectral variation present in the collected spectral imagery.

An interesting question in attempting to find sub-resolved targets is, especially for a space-based HS sensor, what is the optimum ground sample distance (GSD) to maximize detection rates and minimize false alarm rates, if indeed an optimum does exist. A simple analysis identifies the optimum GSD for a given target size, using the following argument. The most important metric in being able to identify a target embedded in a background is not simply the sensor's signal-to-noise ratio, but rather the *target* signal to *total* noise ratio. Only the target's signal is germain to determining the presence of the target or not, but the total sensor noise plays in corrupting the reliability of detection. Target-signal-to-total-noise ratio is easily generated for a given target size using the following proceedure. First the sensor SNR must be determined. In the case of a well designed system, limited neither by detector read-noise nor background (not signal) shot noise on the low signal end, nor by full-well capacity on the high signal end, then the total number of photons, and hence electrons, collected by a detector increases as the cube of the GSD. The AΩ product of the sensor increases as the square of the GSD – the projected pixel size – and the allowable integration time increases linearly with GSD, resulting in the aforementioned cube dependance of detector output with GSD. If the system is well designed such that it is near signal shot-noise limited, then SNR increases as the 3/2 power of the GSD. A given size target will remain fully resolved until the GSD increases to ¼ the smallest dimension of the target, beyond which point no single pixel can be guaranteed to be fully filled by the target. As GSD increases, the average pixel fill will drop, ultimately as the square of the GSD (simple area ratios in the limit of large GSD). The product of sensor SNR and pixel fill factor gives the desired target-signal-to-total-noise ratio. As shown in figure 2-6 for a tank sized target (15 m^2), this metric rises

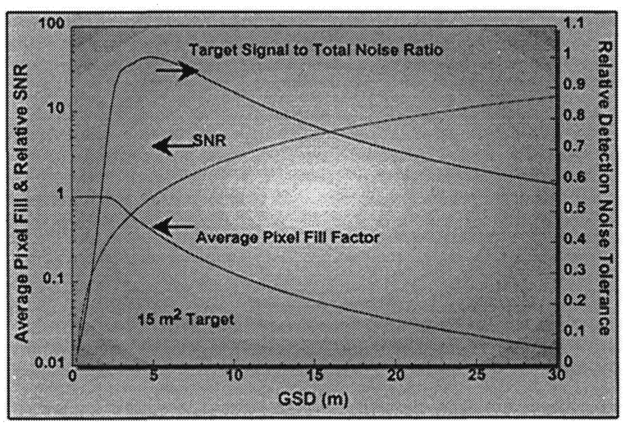

Figure 2-6. SNR increases with GSD, but ultimately not as rapidly as average pixel fill decreases, resulting in an optimum GSD to find given sized targets.

rapidly as GSD increases from zero due to the rapid increas in SNR and large target pixel fill factor, peaks, and then decreases, ultimately as the –1/2 power of the GSD. For the tank-sized object, target detectability is maximum at a 5 meter GSD, translating to an average pixel fill of 40%. Figure 2-7 shows this effect by plotting the relative false alarm rate for detection of a tank as GSD is increased, with probability of detection held constant. Note that false alarm rate falls rapidly, and then increases more slowly beyond the 5 m optimum GSD.

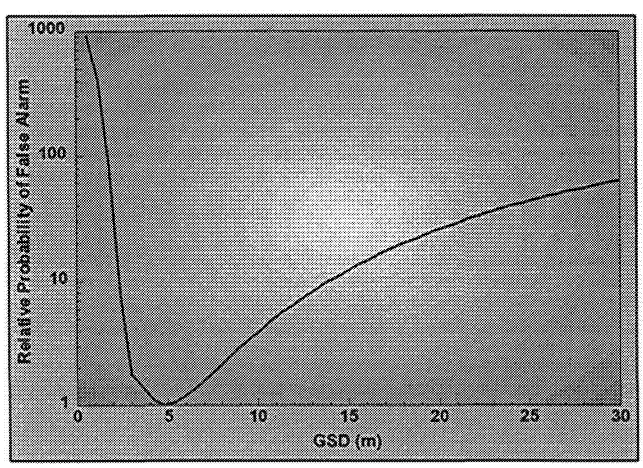

Figure 2-7. While holding probability of detection constant, increasing GSD brings false alarm rate through a minimum at 5m, when searching for 15 m^2 tanks

3.0 EXAMPLE REAL-TIME APPLICATIONS

Presented at the outset of this paper (section 1.0) were two examples of military applications that require real-time collection and processing of hyperspectral imagery, and also require forward-looking imaging capability. These two, extraction of a threat (tank) from background, and identification of land mines (surface in this case), require that the threats be found now, not two weeks from now. A tank commander who must wait two weeks, or even just hours or minutes to find out if an enemy tank is targeting him will find out much quicker in a much less pleasant way. He must have the information immediately, especially if he is to scan a scene to identify the threats. Similarly, an advancing convoy or patrol must know now where the mines are that the must destroy or avoid.

As tremendous as the military potential for this technology is, the industrial, medical, and commercial possibilities are even larger. For instance, the ability to identify surface contaminants throughout an image would be a boon to a wide variety of applications, including coatings and semiconductor production, aircraft/vehicle painting, anti-terrorism (in the form of identifying explosive compounds on vehicles or packages), etc. Since production or processing capability translates directly to profitability, any system designed for such applications must produce immediate indication of surface cleanliness. Figure 3-1 shows an example of composite coupons viewed by a broad-band IR imager, and the results of filtering spectral imagery to identify areas of mold-release contamination where subsequent coatings would not adhere.

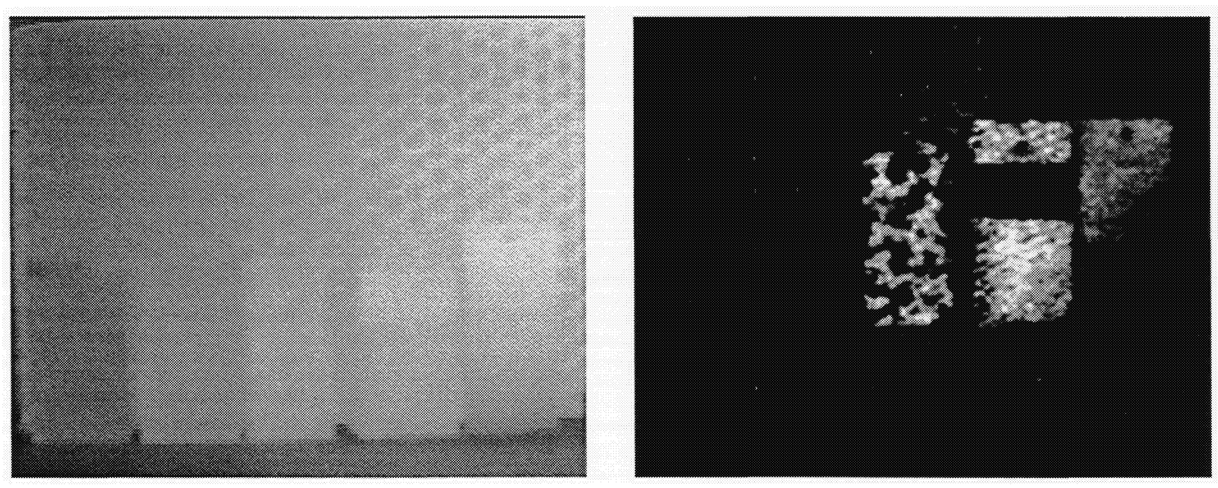

Figure 3-1. Broadband IR imagery (left) fails to show the mold-release contamination on composite coupons. Real-time spectral correlation on hyperspectral imagery (right) reveals the contamination.

Medical applications of real-time hyperspectral imagery are also tremendous, from aiding the pathologist in identifying cancerous cells on slides, to *invivo* identification of such cells and tumors, to identifying necrotic tissues. Any medical need to differentiate between tissues, whether within the body or without, would benefit from this technology.

Counterfeiting is another area where real-time collection and exploitation of hyperspectral imagery will have great benefits. As counterfeiting techniques improve, the need for better identification of such counterfeit bills also grows. Use of new bills with better anti-counterfeiting measures makes the counterfeiter's job more difficult, but not impossible. Surface Optics demonstrated the ability to differentiate real from counterfeit bills based on differences in the inks used on the bills. We extended local banks the opportunity to use this technology to screen for counterfeits, offering to find and return all counterfeit bills from any mixture of real and counterfeit money delivered to us. None accepted the offer.

Industrial process control is yet another application where spectral imaging has broad applications. With the ability to differentiate between various chemical components throughout an area, this technology is ideal for monitoring production of webs, mixtures in tanks, application of coatings, etc. Figure 1-1 presented a color image generated from a hyperspectral image cube of a group of spectrally diverse objects. Not apparent in the scene is the fact that the dark green board in the background has an anomalous paint on it, which matches the color of the dominant paint. Such is analogous to a flawed application of materials in a production process. Although the anomalous paint is not visible to the eye, subtle spectral differences allow its identification. Figure 3-2 shows the identification of the bad paint using real-time spectral correlation.

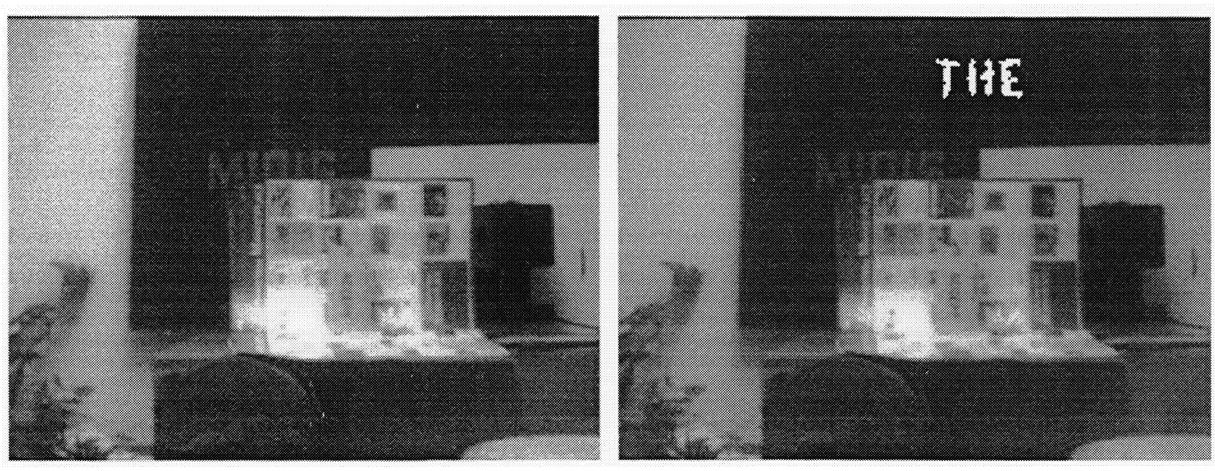

Figure 3-2. Areas of anomalous paint, not visible in the color image at left, are revealed by real-time spectral correlation.

4.0 CONCLUSION

Hyperspectral imaging is the latest advent in the imaging field. For this technology to take its place as a widely used, advantageous successor to panchromatic or multichromatic (e.g., color) imaging, it must move from a scientific realm to a commercial realm. Timely exploitation of the imagery is key to making this transition, for real-world applications in general require real-time exploitation. By using custom, software-reconfigurable hardware processing boards, coupled to high-speed spectral imagers, Surface Optics Corporation has achieved the ability to gather and exploit hyperspectral imagery at video rates, thereby producing the first patented, compact, hyperspectral video camera.

5.0 ACKNOWLEDGEMENTS

Surface Optics Corporation thanks Dr. Paul Willson at US Army ARDEC for his continued support in furthering development of the technology originally devised and patented at SOC. The Jet Propulsion Laboratory has also been vital in development of the spectral imagers for MIDIS, lending their many years of expertise in this field to ensure optimum performance. Work done by JPL in atmospheric correction is also key to operation of parts of the processor.