LA-UR- 01-1949

Approved for public release;
distribution is unlimited.

Title: Enhanced Sensitivity for Hyperspectral Infrared Chemical Detection

Author(s): Phillip L. Jacobson
Roger R. Petrin
Aaron Koskelo
C. Robert Quick
Jerry Romero

Submitted to: SPIE 15th Annual International Symposium on
Aerospace/Defense Sensing, Simulation, and Controls
16-20 April 2001
Orlando, FL

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Enhanced sensitivity for hyperspectral infrared chemical detection

Phillip L. Jacobson, Roger R. Petrin, Aaron Koskelo, C. Robert Quick, and Jerry Romero

Los Alamos National Laboratory
Los Alamos, NM 87545

ABSTRACT

The sensitivity of imaging, hyperspectral, passive remote sensors in the long-wavelength infrared (LWIR) spectral region is currently limited by the ability to achieve an accurate, time-invariant, pixel-to-pixel calibration of the elements composing the Focal Plane Array (FPA). Pursuing conventional techniques to improve the accuracy of the calibration will always be limited by the trade-off between the time required to collect calibration data of improved precision and the drift in the pixel response that occurs on a timescale comparable to the calibration time. This paper will present the results from a study of a method to circumvent these problems. Improvements in detection capability can be realized by applying a quick, repetitive dither of the field of view (FOV) of the imager (by a small angular amount), so that radiance/spectral differences between individual target areas can be measured by a single FPA pixel. By performing this difference measurement repetitively both residual differences in the pixel-to-pixel calibration and $1/f$ detector drift noise can effectively be eliminated. In addition, variations in the atmosphere and target scene caused by the motion of the sensor platform will cause signal drifts that this technique would be able to remove. This method allows improvements in sensitivity that could potentially scale as the square root of the observation time.

Keywords: Passive remote sensing, calibration, enhanced non-uniformity correction

1. INTRODUCTION

Infrared detector arrays have found significant application in remote sensing, including surface temperature monitoring, multi-/hyperspectral geologic categorization, atmospheric absorption, effluent monitoring, and military target detection/identification. Each of these applications is slightly different in demand for sensor performance; however, independent of whether the sensor is an imaging or spectrally dispersed design, detector arrays are required to have uniform response across the detection surface. Residual error in the non-uniformity correction of these arrays typically limits the performance to some level. For this reason, significant efforts have been applied both to detector material improvement and to correction algorithms/processes. For many imaging applications, standard calibration techniques work adequately, since the image resolution is high enough to ignore single-pixel effects. The applications of interest here are more towards long-range where targets are most likely on the order of a pixel or two at most. In such cases, non-uniformity of the array is a critical limitation since a target/background ratio will be significantly affected by offset in either pixel.

The standard technique applied in calibrating an infrared focal plane array is a two-point non-uniformity correction (NUC). This method uses two different radiance levels to calibrate detector response in such a way that after correction all pixels will ideally respond with the same signal output at each of the calibration levels. The most widely used calibration approach is to have a sensor stare at a uniform target (e.g. extended area blackbody) of known temperature and emissivity while measuring the array response. This measurement is then repeated while viewing a target of different effective temperature. In this fashion, a two-point NUC is performed by using these two measured response values, for each pixel, to calculate an array of offset and gain corrections, thus forcing the calibrated detector output to be ideally uniform across the array at these two temperatures. Once the calibration is determined, the detector output is measured, corrected, and saved. By selecting temperatures that bracket the typical radiance levels to be measured in the scene of interest, the sensor is assumed to operate uniformly (and linearly) within that range of response. This calibration technique is simple and straightforward to implement; however, three main drawbacks include:

- 1) the time required to perform the measurements (losing signal averaging time)
- 2) the need for stable, uniform blackbody calibration targets onboard the sensor, and
- 3) the required linearity in detector response.

The measurement time and non-linearity become even more important with some of the long-wavelength arrays in use today, since the stability of NUC results with those arrays are sometimes as short as a few minutes.¹ Commonly used sensors today may spend a few minutes calibrating and possibly only have 10-15 minutes of stable calibration time to measure and signal average.

The concept of correcting a calibration with a scene-based NUC (SBNUC) has been examined before with encouraging results.^{2,3,4,5} In this technique, the motion of the scene across a detector array is used to detect residual offset/gain corrections at the given scene radiance. The procedural concept is based on the same target element, with constant radiance, being measured on two nearby detector elements. Any measured difference in signal output is considered error in the NUC and becomes a correction factor. (Since this assumes constancy in the target radiance, some applications suffer when rapidly changing targets are of interest.) Being a "realtime" update procedure, temporal NUC fluctuations are alleviated. In addition, since the correction occurs at the scene radiance of interest, there is little affect from non-linearity in the detector response. However, little effort has been put into the study of how a sensor might use applied motion (dithering) of the scene on the FPA in order to measure the difference in radiance of two target areas-of-interest on the *same* detector pixel(s), creating a "difference image". In typical measurements, reliance on measuring a target area and comparing to a background area, measured on two sets of pixels, is the underlying premise. These two sets of pixels would most likely have different time-varying response offset (and gain), which would significantly limit the signal averaging that could be achieved. A new method of sensor design and data acquisition, where the scene image is dithered spatially across the detector array, would effectively eliminate any non-uniformity that would otherwise exist between the measurement of two spatial targets on two different pixels. By measuring these dithered images as a function of wavelength, spectral information can be obtained over longer averaging times, thereby creating immunity to drifts in calibration and increasing the potential sensitivity of the sensor.

The proposed technique is a Controlled-Dither Non-Uniformity Correction (CDNUC). Counter to the idea of actually calibrating a detector and then relying on stability during detection or signal averaging, this technique

1. applies a rapid, dithering motion to the detector field-of-view,
2. uses this motion to correct *residual* detector non-uniformities, and
3. controls the dither to measure target and background data on the same detector pixel(s).

The method differs from previous scene-based techniques^{2,3,4,5,6} primarily by the third point, where the desired comparison measurements are made on a single pixel, thereby removing temporal variations in the detector operation/calibration. Previous SBNUC procedures typically relied on normal scene motion or a fixed-pattern, scanning element. As mentioned, applications with single-pixel targets benefit the most from this procedure; however, other applications will realize an improvement in absolute detection accuracy, in addition to the extended averaging time provided.

2. IR CAMERA CALIBRATION

The IR sensor used in this study consists mainly of a cryogenic (LHe) dewar housing a HgCdTe 256x256 detector array manufactured by Amber Engineering (AE-173), a set of relay/steering optics, a galvo-tuned dither mirror (Cambridge Technology, Inc. Model 6900), a telescope, and a large gimbal mirror for targeting. The camera, shown in Figure 1, was operated at a 40 K array temperature and irradiance was limited to a 2 micron spectral bandpass centered around either a 9 or 10 micron center wavelength using a cooled interference filter inside the dewar. The array readout operates in a snapshot mode where all pixels integrate photocurrent for a selected integration time (0.33 msec) and then the signal voltage levels are read out at a slower rate (9 ms for the 111 Hz frame rate). The control hardware used to run the camera was manufactured by SE-IR Corp., and uses a two-channel 14-bit A/D datastream through a digital signal processor and Coreco framegrabber. A camera control unit is necessary for the standard NUC, since it measures, calculates, and stores the calibration values to be applied to all subsequent data collected.

Calibration of the camera was achieved using a pair of infrared radiation sources. Two different size sets of blackbody calibration sources were available, depending on the target location. The larger set, shown in Figure 2, was used when viewing through the telescope to a remote target. The smaller set, while more accurate and stable in temperature control, could not fill the telescope field-of-view as required for calibration, and was only used when viewing internal targets. The calibration procedure followed was to first measure a set of 25 frames with the camera viewing the cold blackbody source (ϕ_c = typically 10°C), which took approximately 5 sec. A second set of 25 frames was then measured viewing the hot source (ϕ_h = 35°C). The measured frames were averaged to give a two-point set of data ($V_{x,y}$) for each detector pixel at the selected

Figure 1. A five-inch cryogen dewar houses an AE-173 long-wavelength infrared detector array. The lens doublet in front of the dewar re-images the telescope image onto the camera. All optical elements in the camera field-of-view must be included when calibrating.

Figure 2. Two infrared sources are shown here mounted on the inside wall of the sensor trailer, with the cold source below the hot source. The sensor is pointed to each source in turn by tilting the large gimbal mirror up or down. The hot source can be as high as $\sim 50^{\circ}\text{C}$, while the cold source can be $< 10^{\circ}\text{C}$ by running chilled water through a reservoir behind the black plate. Temperatures are controlled and monitored with eight separate heaters and thermocouples.

temperature extremes. The entire procedure can take at least a couple minutes. Under the assumption that all detectors respond uniformly and linearly between the selected temperatures, a linear fit to the measured detector voltages should completely calibrate the array response to be uniform. An array of gain and offset corrections was calculated based on

$$V_{\text{cal}}(\phi) = O_{x,y} + G_{x,y} \cdot V_{x,y}(\phi) \quad (1)$$

The assumption of uniform response sets V_{cal} to be identical for all pixels at the calibration radiance levels. Initially, at least at the two selected radiance levels, this calibration removes any offset and gain variations between detector pixels, which may be due to detector bias fluctuation, temperature variation, differences in the two readout amplifier chains, offsets in the A/D channels, and other sources.

Temporal instability in the calibration gain and offset factors can be a limiting factor in the sensitivity of an infrared camera. Characterization of the AE-173 camera was performed to better understand the temporal behavior of calibrations. The small infrared sources were used to measure/calculate a NUC. Once the corrections were determined, calibrated data was recorded as a function of time, one frame every ten seconds for over two hours. The standard deviation of each frame of data was calculated and plotted as a function of time in Figure 3, showing a clear increase of the spatial noise to almost 3 times that of the temporal noise. The total noise for any single frame of data is given by σ , where

$$\sigma^2 = \sigma_t^2 + \sigma_{sp}^2(t) + \delta_{res}^2 + \dots \quad (2)$$

Assuming independent contribution, each term adds in quadrature as shown, including the inherent temporal noise (σ_t) such as photon shot noise, electronic noise, etc..., the time dependent spatial noise (σ_{sp}), the residual uncertainty in calibration (δ), and any other terms not considered here. In most situations, the temporal noise is desired to be the dominant noise term and, being random, should average down as $1/n^{1/2}$, for n frames measured. The residual calibration uncertainty is reduced through a balance of calibrating during a short time period (not too long between measuring different calibration temperatures) and averaging long enough to reduce the temporal noise contribution to the measured detector response. In other words, the longer the averaging time at each temperature value the more accurate the measurement, but the longer between the two temperatures the more inaccurate the two calibration points may be, due to detector offset drift during each measurement.

Figure 3. The first figure shows a plot of signal level, for three selected pixels, as a function of the time since calibration. Data was taken at a rate of one frame every ten seconds. Over the ~2 hr time period, while staring at a uniform calibration target, the detector pixels drift in offset together, but with some spread in response. This spread in relative offset appears as a spatial “fixed pattern” noise (FPN). The second figure shows the calculated standard deviation, in temperature units, for each frame of data. The increase in noise is due to the increased FPN. Where $c=1$, a time constant is measured to be 70 minutes (including the 20 minutes between calibration and start of the data set). (This assumes the standard deviation is dominated by temporal noise to begin with.)

For any given dataset, the temporal noise can be calculated by¹

$$\sigma_{i,j}^2 = \sum_{j=1}^n \frac{(x_{i,j} - \langle x_i \rangle)^2}{n-1} \quad (3)$$

for the i^{th} pixel and n frames of data. The mean temporal noise is calculated by averaging this variance over the N pixels in the array. Assuming the calibration is accurate enough that any residual uncertainties are negligible, along with any other contributing factors, the temporal and spatial noise are dominant. A “correctability” figure of merit, defined in detail elsewhere¹, can be used to describe measurements of a single blackbody source. This factor, c , is simplified to

$$c = \sqrt{\frac{\sigma^2}{\sigma_i^2} - 1} \quad (4)$$

The correctability (c) is defined here in such a way that when the spatial and temporal noise contributions are equal, $c=1$. A line is drawn in Figure 3 to illustrate the time constant associated with this array calibration (~ 38 min.). This time constant has been defined as the long-term stability. To visualize the effect of such spatial degradation in a calibration, two images are shown in Figure 4. First is an image captured immediately after calibration, illustrating the uniform “white” noise behavior of the temporal contribution. Second, an image taken 2 hours after calibration is presented to illustrate the fixed-pattern noise (FPN) that arises, both on local scale (pixel-to-pixel crosshatch) and across the array (top-left to bottom-right corners). In addition, both images show inoperable pixels that do not pass calibration criteria (saturated, flickering, unresponsive, etc...). These pixels are designated “bad pixels” and turned off so that the output of the array at those locations is zero. (An existing option in the controller to replace these bad pixels with a “nearest-neighbor” average value was not used.)

Assuming a perfectly stable calibration, even the inherent non-linearity of the detector response degrades the application of only a two-point NUC. When only two calibration source temperatures are used, any radiance level other than those selected for calibration will lead to deviations from the ideal cases. Detector response will vary from the expected linear fit and lead to an increased spatial noise. This can be seen in Figure 5, where the average noise level is plotted as a function of target source temperature. The two endpoints are the calibration temperatures. After calibration, while staring at the cold source, the temperature was raised. The spatial noise is seen to increase significantly above the baseline temporal noise (~ 15 units).

(The line is a simple fit to the data to aid the reader.) The increase in spatial noise as a function of varied irradiance indicates a non-linear response in the detector array that is not calibrated for with the two-point NUC.

Figure 4. The figure on the left shows a typical image of a calibration source immediately after calibration. The black pixels (mainly one large cluster) are inoperable pixels, set to zero signal output, and ignored when calculating noise characteristics. The “white” noise illustrates that the frame noise is dominated by temporal noise in the detectors. The figure on the right is a frame taken of the same calibration source after ~ two hours. Two features of interest are the offset between the top-left and bottom-right corners and the “cross-hatch” fixed-pattern noise evident. This increasing spatial noise is typical in IR cameras.

Figure 5. A single blackbody source is monitored as the temperature is raised from 10-25°C and cooled back down. A set of 900 frames is measured at each temperature. The average standard deviation from each frame is plotted vs. temperature. The value at 35°C is measured with the other calibration source. At the calibration temperatures, spatial noise is negligible and only temporal noise contributes to the measured standard deviation. In order to account for the long time period to increase the source temperature up to 25°C, a modeled temporal drift (determined from the data shown in Figure 3) was subtracted from the calculated σ . Resulting values are plotted here.

3. DITHERED SCENE METHOD

The Controlled-Dither NUC is a technique of both correcting residual non-uniformities in an array, based on scene motion, and measuring comparison data from two different target points-of-interest on the same detector pixel. This latter aspect of the technique is the most important and effectively eliminates any fluctuation that can occur in a signal due to varying relative pixel responses. For applications where targets may only extend over one or two pixels, non-uniform response in an array may be a limiting factor in staring systems. A complex scene⁷, such as the one shown in Figure 6, may have such items

Figure 6. A complex scene is shown here, where several targets that may be of interest are of only single-pixel size. Also, a number of bad pixels can be tentatively identified, but may offer false positive or negative classification.

Figure 7. Two dithered images are shown with an aluminum target in the field of view. A synthetic plume is added to illustrate how selected pixels can measure both the background target and the plume on alternating frames. A line is drawn to aid the reader in identifying a single row of detector pixels common to the two images.

as vehicles, weapon systems, or chemical stacks only a few pixels in image size. Also seen in this image, are several that may be bad detector pixels, but absolute certainty is difficult to achieve without the aid of scene motion. By moving the scene back-and-forth, defective pixels can be identified, corrected, and/or avoided. In addition, if target and background locations can be determined, the dither direction and magnitude can be selected such that a single pixel can measure both target locations for comparison frame-to-frame.

For dithered-scene motion, an accurate, repeatable optical motion must be achieved, using either a transmissive or reflective element in the optical train. Rates of motion should be selected such that the field-of-view (FOV) settles before frame capture. The frequency of motion can be as fast as the frame rate of the camera in use. The ideal case would be to have alternating sequential image frames at each FOV position, which allows difference measurements to be made as close in time as possible. The actual rates will depend on the imager used, the application and target variability, and whether spectral information is to be gathered at each location. Magnitudes must be variable and two-dimensional motion is desired. Sub-pixel accuracy and repeatability is necessary for achieving the desired results.

For the purpose of the work presented here, a galvo-tuned mirror was placed in a semi-collimated section of the IR sensor optical train and was capable scanning at up to 100 Hz with a repeatability of < 0.2 pixels ($20 \mu\text{rad}$). By simply setting the mirror to scan repeatedly, a simple procedure of collecting images and then registering based on a scene fiducial was possible and quite successful. With a more advanced control system in place, a mirror feedback loop to the camera control would replace the need for a spatial fiducial. With only a single dither mirror, motion along a detector array axis was selected and the fitted position of the fiducial gave the required pixel shift. Once registered, the data set was then processed to achieve two types of "difference" imagery, one with the scene spatially registered to overlap again and the other spatially offset by the dither magnitude, but registered temporally for image subtraction. In the first set of images, each scene area is measured by two different detector pixels. With perfect motion and constant scene radiance between measured images, all scene information is subtracted out and detector non-uniformities stand out significantly. In the second set of images, temporal registration is determined only to segregate images into the different FOV positions. No spatial registration is performed. In this step, each detector pixel measures two separate target locations, defined by the dither motion. The difference image gives the radiance contrast between different points in the scene, but on single detector elements.

4. EXPERIMENTAL RESULTS

The measured datasets were recorded using the infrared camera control system. The output was windowed to a specific 112×128 section of the array, and each of these pixels gave a voltage level stored as a 16-bit integer for each frame. Typical files were 900 frames in size and could be repeated consecutively to concatenate files to obtain longer measurements. When viewing outdoor scenes, a large gimbal mirror (0.75 m diameter) is used for targeting. Files were stored and analyzed using IDL 5.4. Once custom pre-processing has registered, segregated, and subtracted the frames, standard image analysis could be used for statistics and pattern recognition on the dithered data.

Initial experiments have been limited to understanding basic methods of data acquisition and analysis. Preliminary data were taken staring at a uniform aluminum plate target at a 2 km range; however, due to the complexity of atmospheric effects (turbulence, weather, etc.), more recent experiments have been moved indoors. This effectively concentrates our study more on camera instability than target or application specific variability. In addition, actual chemical releases and spectral

information have been ignored until now, since the real solution to any problem is measuring radiance levels accurately, whether at one wavelength or many.

When measuring a target signal, comparison with a background level is typically necessary for reference. This is illustrated in Figure 7, with a uniform target plate viewed by the IR camera, and a dither motion applied vertically between capture of the two frames. A synthetic plume is used to illustrate how this target of interest would be measured by first measuring an expected background signal, then measuring the signal with the plume, and finally subtracting the two measurements, which were taken on the same detector pixel. A line is drawn between the two frames to illustrate how a common row of detector pixels would measure background in the first frame and plume absorption on the second frame. This line also gives a reference to the ~ 10 pixel magnitude of the dither in this dataset.

As described above, the first use of the CDNUC would be to identify residual non-uniformity in a detector array calibration. A typical image is shown in Figure 8, where a large outdoor target appears, surrounded by objects of various radiance levels.

Figure 8. A sample image of a target from a dithered data set is shown here, along with an enlarged section to the right. Two features are of specific interest. First, the black spots are “bad” pixels, defined as inoperable after the NUC. These are left uncorrected and unused in the data analysis. Second, the target has non-uniform features, both single-pixel and area, even after the NUC was performed. These features are highlighted with arrows.

The first aspect to notice in the image is the set of residual “bad” pixels that failed the standard two-point NUC. These pixels are all set to have zero response (turned off) and are seen as black pixels. A few less obvious features are highlighted with arrows, where the array seemed to have both single-pixel offsets (left and right arrows) and one large area with non-uniformity (middle arrow). An enlarged section of the image is shown to the right with the non-uniform area encircled.

While staring at the target with the IR sensor, the gimbal mirror was dithered at a 1 Hz repetition rate and the camera captured data at a 111 Hz frame rate. Under this mode of operation, several sets of data were captured. For the mirror dithering over a range of approximately 10 pixels (and ~ 5 in the other dimension with coupling), a set of 3600 frames were measured. To register this data set, the calculated target centroid was applied to a subset of the full 112x128 images captured. Each frame of this subset has the centroid of the target at the same location. The registered set of data was then

Figure 9. An image depicting the standard deviation (σ) for each element in the dithered data set is shown here. The target has streaks of non-uniformity caused by both single-pixel and area defects in the detector NUC. The dithering motion essentially magnifies the effect of the non-uniformity, which allows for easy identification and possible correction.

processed to determine the segmented averages and standard deviation of the mean (σ_{sdom}) as a function of N. As discussed

above, 12 different values of N were used in this calculation. For $N=2$, the calculated standard deviations across the subarray are presented as an image in Figure 9. This image is of particular interest to address the issue of residual non-uniformity. The highlighted single-pixel locations in Figure 8 are obvious in this image as the solid “streaks” near the center of the image. The large non-uniform area in the center of the target is also clear. These residual NUC errors can be identified, corrected, and/or avoided for later acquisitions. The images shown here also illustrate how a dither without sub-pixel control will effect the precision of the difference imagery.

Recent experiments have included a precision galvo-tuned mirror, which allows for simpler interpretation of the residual non-uniformity identification measurements. A set of images shown in Figure 10 were measured using a two-pixel horizontal dither. The measured images were spatially registered to overlap and subtracted in alternating fashion for each dither location. In this method, all spatial scene information is removed and the only remaining features are due to “bad” pixels and residual non-uniformity in the detector array. Two sub-arrays are shown, one to better illustrate the presence of the desired “binary” signals and the other to show how processing can pull out the associated detector pixel locations and give the direction of their offset from the local array mean. In this application, the binary signal is recognizable as a pair of pixels separated by the dither magnitude (in this case, a 2-pixel dither). The two-valued signal is indicative of single-detector-pixel non-uniformity being *dithered* in the original dataset. In essence, this technique relies on the spatial registration, where the dithering is effectively reflected back onto the array in such a manner that abnormal features in the array are displaced spatially and become easier to recognize. A row of data from the difference image of Figure 10 is shown in Figure 11. Here the binary signal is obvious, being significantly larger than the observed noise level. Since this signal results in a measurement twice as large as the single offset measurement (400 counts vs. 200 counts in this case), this method allows for

Figure 10. The first image shows a target scene with spatial structure due to a radiance contrast of approximately 200 counts. When two dithered images are spatially registered and subtracted, all scene structure is removed and only residual detector fixed pattern noise remains. One such difference image is shown here, with a box drawn around the region expanded in the sub-images. The expanded sub-image shows the binary (positive/negative) pixel pairs indicative of a single detector pixel with an abnormal offset level. The duplicate black pairs are due to inoperable pixels determined during the standard two-point NUC. The second sub-image is an identification of the specific detector pixel and direction of the offset, white for positive and black for negative.

Figure 11. A single row of the difference image shown in Figure 10 is plotted here to illustrate a typical “binary” signal indicative of a single detector pixel with a residual offset.

detection of offset errors smaller than otherwise possible and with greater reliability. In addition to the black/white binary-signal pixel pairs, pairs of black pixels are also observed in Figure 10. These are indicative of a single “bad” or inoperable detector pixel.

For increased stability and precision in characterization measurements, an indoor target, exhibiting both spectral absorption and a spatial fiducial for registration, was installed. A thin strip of polystyrene and a 0.01" diameter wire were held in front of a uniform aluminum plate at room temperature. Calibrations were performed using the smaller blackbody sources held at 10 and 35°C. The camera NUC procedure set the calibration temperatures to be at approximately 11500 and 7200 counts on the 14-bit A/D output. This implies a temperature resolution of 5.8 mK/count. For the measured temporal noise (9 cts.) associated with this system, $NETD=52$ mK. For illustration of the target used and the dither magnitudes involved, Figure 12 shows a pair of images taken at the two positions of the mirror motion. Using the wire position, the raw data file is processed to determine fiducial position as a function of frame number. The images are then separated into three sets, position A, position B, and unsettled motion (discarded). Temporal registration is maintained such that, for each period of dither motion, the same number of frames from each position are saved, thereby ensuring that subtracted frames must be from within the same period of mirror motion. This procedure maintains a short time delay between relative measurements of the different target positions. The images shown here are one from each position, A (lower) and B (upper). As opposed to the spatial registration applied in the non-uniformity identification, these images are subtracted without any spatial registration. In this process, each detector pixel will measure the radiance from two different target areas. A “difference” image is shown in Figure 12. The point-of-interest here is the polystyrene edge. Since the same uniform background extends behind the polystyrene, the gray band in the center of the image is indicative of a measured absorption. At mirror position A, those specific detector pixels measured the background and at position B they only measured the radiance not absorbed by the polystyrene. This difference measurement is directly related to the broadband absorption (752 cts.) of the polystyrene sheet.

Figure 12. Two frames frame an unprocessed dataset are shown here. Each frame has the dither mirror at one of the two, selected target positions. A thin sheet of polystyrene absorber is on the left and a 0.10" diameter wire on the right. The image on the right is a “difference” image, the result of subtracting the two raw images. The absorption measurement is found in the gray band in the middle of this image, where each pixel has measured both the polystyrene and the background sequentially.

The typical approach to an imaging sensor absorption measurement would be to measure radiance values of a scene containing both the target-of-interest and a background similar to that expected to be under the target. A detector pixel(s) is identified for each area and the pixels are compared either by subtracting the values or calculating the ratio. In the top image of Figure 12, the polystyrene is in front of a uniform target. The standard method would compare pixels left of the edge (in the center of the image) to pixels right of the edge. An analysis was performed here on this data and an absorption value was measured to be 748 counts (± 70 , for an area selected along the edge of the polystyrene). A similar analysis on the other dither position (the lower image) yielded a value of 746(63) counts. In contrast to this standard method, analysis of the difference images, where comparison measurements are made using the same detector pixel for each spatial target area, yielded a comparable absorption level (752 counts), but a reduced noise level of $\sigma=24$ cts. This reduction in measurement noise is due to residual offsets being subtracted out of the difference images.

A set of images, shown in Figure 13, illustrate a raw image and a “difference” image, where the first shows significant FPN and the second shows the subtraction of this pattern noise from the resulting imagery. Expanded views are given to offer clarity and to show the area of interest surrounding the edge of the polystyrene absorption target. In the raw image on the left, a crosshatch pattern is seen and several pixels are noted to have a significant offset. When the FOV is dithered, both effects are removed from the difference image, as evidenced by the sub-image on the right. In addition, for this two-pixel dither magnitude, the absorption region is seen as a narrow band in the difference image.

Figure 13. A raw image is shown on the left, with an expanded sub-image. Besides the “cross-hatch” pattern noise, several pixels are observed to have large offsets. This noise significantly limits the sensitivity of typical absorption measurements. A differenced image is shown on the right, with the same sub-image area shown. The pattern noise has been subtracted off so that all pixels show approximately the same signal level, except the area with a true absorption signal at the edge of the polystyrene sheet. The black pixels remaining are inoperable pixels identified by the two-point NUC.

5. CONCLUSION

A new technique, the Controlled-Dither NUC, has been presented here for application to infrared array-based remote sensing platforms, both existing and new designs. This method has been shown to identify residual calibration errors in an array, measure correction factors to improve array uniformity, and, most importantly, negate the effects of any non-uniformity on actual radiance comparison measurements, such as chemical absorption. With a dither of the camera field-of-view, controlled scene motion is created, allowing single-detector-pixel measurements of multiple points on the ground. Subtracting these measurements gives a difference image set, in which all relative offsets between pixels are removed. The elimination of both residual NUC errors and spatial noise due to detector drift should reduce the detector array noise to the ideal limit of pure temporal noise. This is the best that any sensor can perform. Measurements presented here, on data taken 30 minutes after calibration, showed dramatic reduction of spatial noise by a factor of 3 over the alternate method of using pairs of detector pixels to measure absorption differences. This improvement in single-frame noise is a factor of 9 in reduced averaging time to reach a given signal-to-noise measurement. Such improvements are of great interest in many applications.

The ultimate sensitivity of target detection is determined by the total signal-to-noise. Where signal levels are limited, noise becomes a significant driver in system design and application. The CDNUC technique is most important in having both the ability to *correct* detector array calibrations, thereby reducing systematic noise levels, and to *bypass* calibration deficiencies, which lead to measurement inaccuracy, for single pixel targets. The dither method offers the means of significantly improving the capability of both existing sensors and new designs, by eliminating the need to constantly re-calibrate a sensor and shortening the required averaging time of an acquisition period.

ACKNOWLEDGEMENTS

I would like to acknowledge the help of all those involved with this project, from hardware maintenance and operation to data analysis and technical discussion. Their support helps ensure the success of the program. Especially, I would like to thank Bob Nemzek for his efforts on the Fabry-Perot sensor hardware and the improvements he was able to make in the imaging quality of the infrared camera.

This work was performed under the auspices of the U.S. Department of Energy by Los Alamos National Laboratory under contract W-7405-Eng-36. Funding support is from the National Reconnaissance Office through the Director’s Innovative Initiative 2000 Program.

REFERENCES

-
- ¹ W. Gross, T. Hierl, and M. Schulz, "Correctability and long-term stability of infrared focal plane arrays", *Opt. Eng.* **38**(5), 862, (1999).
- ² W.B. Kendall and A.D. Stocker, "Self-adaptive calibration for staring infrared sensors", *SPIE* **1954**, 181, (1993).
- ³ D.A. Scribner, K.A. Sarkady, J.T. Caulfield, M.R. Kruer, G. Katz, and C.J. Gridley, "Nonuniformity correction for staring IR focal plane arrays using scene-based techniques", *SPIE* **1308**, 224, (1990).
- ⁴ C.J. Young, "Scene Based Non-uniformity Correction", *Proceedings of the MSS Passive Sensors*, 2001.
- ⁵ S. Tzimopoulou and A.H. Lettington, "Scene based techniques for nonuniformity correction of infrared focal plane arrays", *SPIE* **3436**, 172, (1998).
- ⁶ S.R. Horman, K.C. Hepfer, and M. Zurasky, "Uniformity Compensation for High Quantum Efficiency Focal Plane Arrays", *SPIE* **2744**, 154, (1996).
- ⁷ Private communication with H.A. Fry, Los Alamos National Laboratory.

Figure 13. A raw image is shown on the left, with an expanded sub-image. Besides the “cross-hatch” pattern noise, several pixels are observed to have large offsets. This noise significantly limits the sensitivity of typical absorption measurements. A differenced image is shown on the right, with the same sub-image area shown. The pattern noise has been subtracted off so that all pixels show approximately the same signal level, except the area with a true absorption signal at the edge of the polystyrene sheet. The black pixels remaining are inoperable pixels identified by the two-point NUC.

5. CONCLUSION

A new technique, the Controlled-Dither NUC, has been presented here for application to infrared array-based remote sensing platforms, both existing and new designs. This method has been shown to identify residual calibration errors in an array, measure correction factors to improve array uniformity, and, most importantly, negate the effects of any non-uniformity on actual radiance comparison measurements, such as chemical absorption. With a dither of the camera field-of-view, controlled scene motion is created, allowing single-detector-pixel measurements of multiple points on the ground. Subtracting these measurements gives a difference image set, in which all relative offsets between pixels are removed. The elimination of both residual NUC errors and spatial noise due to detector drift should reduce the detector array noise to the ideal limit of pure temporal noise. This is the best that any sensor can perform. Measurements presented here, on data taken 30 minutes after calibration, showed dramatic reduction of spatial noise by a factor of 3 over the alternate method of using pairs of detector pixels to measure absorption differences. This improvement in single-frame noise is a factor of 9 in reduced averaging time to reach a given signal-to-noise measurement. Such improvements are of great interest in many applications.

The ultimate sensitivity of target detection is determined by the total signal-to-noise. Where signal levels are limited, noise becomes a significant driver in system design and application. The CDNUC technique is most important in having both the ability to *correct* detector array calibrations, thereby reducing systematic noise levels, and to *bypass* calibration deficiencies, which lead to measurement inaccuracy, for single pixel targets. The dither method offers the means of significantly improving the capability of both existing sensors and new designs, by eliminating the need to constantly re-calibrate a sensor and shortening the required averaging time of an acquisition period.

ACKNOWLEDGEMENTS

I would like to acknowledge the help of all those involved with this project, from hardware maintenance and operation to data analysis and technical discussion. Their support helps ensure the success of the program. Especially, I would like to thank Bob Nemzek for his efforts on the Fabry-Perot sensor hardware and the improvements he was able to make in the imaging quality of the infrared camera.

This work was performed under the auspices of the U.S. Department of Energy by Los Alamos National Laboratory under contract W-7405-Eng-36. Funding support is from the National Reconnaissance Office through the Director’s Innovative Initiative 2000 Program.