Inverse sparse grid flood non-uniformity correction

Mark H. Bowden[*a], James A. Buford, Jr.[b]
[a]Optical Sciences Corporation, 6767 Old Madison Pike NW, Huntsville, AL, USA 35806-2187
[b]AMSRD-AMR-SS-HW Bldg. 5400 Fowler Road, Redstone Arsenal, AL USA 35898-5000

ABSTRACT

We present a technique for the correction of spatial non-uniformity in an infrared emitter array projection system for flood scenes. The technique is a sparse grid approach, but, instead of turning on a sparse grid of emitters to estimate the radiance of each one, the array is commanded uniformly to a constant level, and a sparse grid of emitters is turned off. The resultant loss of radiance in the neighborhood of each emitter is used to estimate its response. Typically, less than one percent of the emitters are turned off in a grid, so flood scene effects, such as substrate heating, are accounted for without the complexity of coupled outputs of a full flood process.

Keywords: NUC, non-uniformity correction, emitter array, infrared projector, flood

1 INTRODUCTION

In November of 2003, the U.S. Army Aviation and Missile Research, Development, and Engineering Center (AMRDEC) began development of a non-uniformity correction (NUC) station for its resistor array projectors. The goal is not only to provide a basic NUC capability but also to explore alternative methods and algorithms in the NUC process. To this end, a configuration was conceived that enables rapid development and testing of algorithms by integrating a high-level numerical computation language with the array drive electronics. This provides the means for reading and writing buffers and lookup tables within the hardware at will. The numerical computation language is used to script algorithms, perform trial computations by hand, and plot results thereby effecting a short design-code-test cycle. As the designs converge, the algorithms are transferred to a C++ program with a user-friendly interface. This setup has been used to develop an initial capability for correction of resistor arrays, and experimentation continues to improve the process. This paper describes the NUC station progress and reports a technique that is under development that combines well-known sparse grid[1] techniques with flood[2] mapping.

2 CONFIGURATION

The initial, tiled NUC station configuration is shown in Fig. 1a. The sensor is mounted on a two-axis motion platform that positions it in front of the blackbody or the emitter array as directed by the control software. This arrangement has its advantages and disadvantages. Many detector pixels are mapped onto each emitter resulting in good sampling of the emitter image[3], but data collection times are longer because of the small footprint of the sensor. On the other hand, the blackbody reference is readily available for sensor NUC updates, and the entire sensor optical train is included in the sensor NUC. This configuration has been used to develop a sparse grid implementation that curve fits the emitter responses. The NUC correction function is derived directly from the fitting function, and the piecewise linear correction tables are generated from these continuous functions.

A second configuration has been adopted to explore non-tiled, flood mapping. Fig. 1b shows the layout of the non-tiled configuration. This setup adds a second blackbody for covering a greater dynamic range, as well as to make camera parameter determination faster by being able to set the blackbodies to the endpoints of the desired range and switch between them readily. The emitter array is mounted on a two-axis motion platform to allow for panning among tiles in tiled configuration and to allow repositioning for dead pixel replacement in non-tiled mode. The zoom collimator controls the detector-to-emitter mapping ratio. The reference blackbody sources are mounted on a rail and are hand-

[*] mbowden@opticalsciences.com; phone 1 256 842-8582 x215; fax 1 256 955-7339; http://opticalsciences.com/.

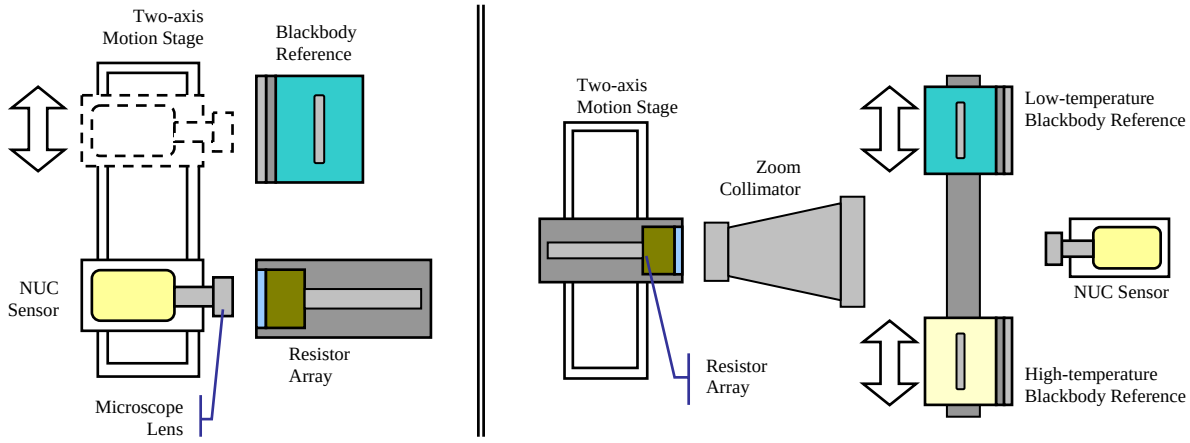

Figure 1: Tiled and non-tiled configurations.

positioned in front of the sensor for non-uniformity correction and calibration of the NUC sensor. In this configuration, the following hardware comprises the NUC station:

- CSA projector electronics
- BRITE I emitter array
- Amber AE187 512x512 MWIR camera
- Two-axis motion stage
- Janos zoom collimator
- EOI blackbody (15°C to 210°C)
- CI blackbody (50°C to 600°C)
- Dell workstation PC

3 NUC SENSOR

3.1 The NUC before the NUC

The first step in correcting an emitter array for non-uniformities is correcting the non-uniformities of the NUC sensor. Initially, the built-in two point NUC of the IR camera was used. For ordinary imaging use, the response of the camera is close enough to linear to be corrected with a single linear correction function. However, for radiometric measurements, the non-linearity of each detector response, especially near the ends of the range, induce considerable measurement error. One could expand the camera range so that the desired measurement range was well clear of the endpoints, but this discards much of the limited precision of the analog-to-digital conversion. A better way is to set the camera to produce raw, uncorrected output and perform the sensor NUC in software.

The authors chose a NUC approach for the sensor that also calibrates the sensor. For each detector, several data points of in-band radiance versus counts are collected across the range of the camera. These data points are interpolated to yield a curve that maps counts to radiance. When an image is captured, the counts-to-radiance curve for each detector is used to convert the image from counts to radiance. Because all of the detectors are being converted into the same absolute scale, non-uniformities are automatically corrected. The output is an image calibrated in radiance.

The type of interpolating curve is of considerable importance. Initially, a natural spline was used to interpolate the data points. This type of spline passes through every data point, which is intuitively desirable since we know those points characterize the camera. However, the natural spline can behave pathologically in some cases, particularly when the data points are close together. Other spline bases exist, though they typically require adjustment to produce acceptable behavior. Given that for a 512x512 FPA, 262,144 interpolating functions are required, “tweaking” is undesirable. In the end, a simple piecewise linear interpolating function was chosen. The knots are concentrated primarily near the endpoints of the curve where the non-linearity is the greatest. The computational and storage burden is much less than for spline functions. Sensor non-uniformities can be held typically to less than 0.5%.

3.2 Handling dead emitters

Virtually any NUC sensor will have some detectors with responsivity outside acceptable limits. One way of handling this is to curve fit the emitter image and use the interpolating curve[4] to replace dead detector values. This also provides for noise rejection as well. A Gaussian curve is a natural choice for the interpolant, but on closer examination, it does not fit that well. The emitter image has greater "shoulders" than the Gaussian curve. The following empirically derived function is used for the interpolant in the AMRDEC NUC station with good results in tiled mode.

$$\phi(x,y) = a_0 e^{-d^2/2} + \frac{a_1}{d^5 e^{a_2/d}}, \text{ where } d = \sqrt{a_x^2(x-x_0)^2 + a_y^2(y-y_0)^2} . \quad (1)$$

There are typically too few detectors per emitter in non-tiled mode for curve fitting to work well.

4 INVERSE SPARSE GRID FLOOD MAPPING

4.1 Methodology

Correction of array non-uniformities requires first an accurate mapping of the response of each emitter to its input. For normal sparse grid mapping at high output levels, sparse grids of emitters are turned on and the image of the array around each emitter is processed to yield a measure of the emitter output. Sparse grids are used to avoid coupling of emitter outputs and to avoid driving large portions of the array at high levels, which can exceed drive current limitations and cause secondary effects such as substrate heating and bus bar robbing. For lower level output scenes where the expected scene content will drive large portions of the array simultaneously, a flood approach is sometimes taken. In this case, it is desirable to account for the secondary effects in the mapping process. Flood mapping is complicated by the fact that emitter outputs overlap each other making it more difficult to relate changes in input to the corresponding changes in output.

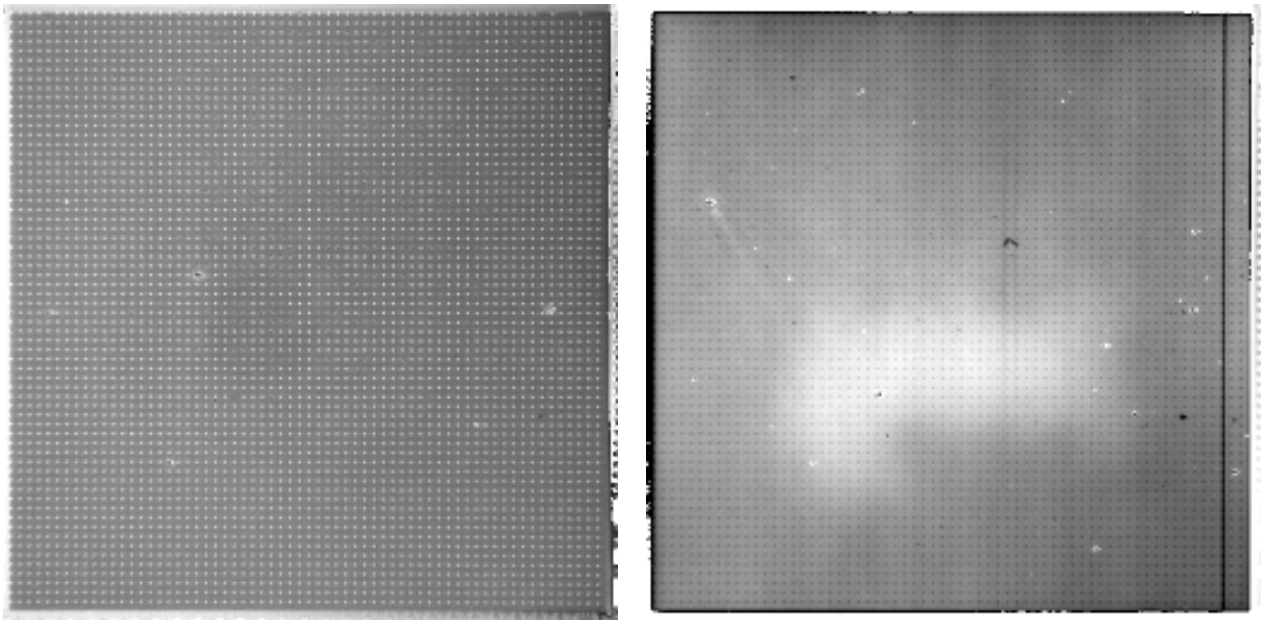

Normal Sparse Grid Inverse Sparse Grid

Figure 2: Sparse grid types.

Inverse sparse grid flood (ISGF) mapping is a technique intended to combine the advantages of uncoupled emitter outputs of sparse grids with the compensation for secondary effects of flood mapping. Fig. 2 shows images of a normal sparse grid and an inverse sparse grid. Note the area of higher output near the lower central region in the ISGF image due to substrate heating. With the normal sparse grid, typically an image of the array in the off state is subtracted from the sparse grid image to give the output above background of the emitters. With the ISGF approach, the inverse sparse grid is subtracted from an image of the array with all emitters on. This approach in effect estimates the output of the emitter by noting the contribution lost when it is turned off.

Once the emitter responses are determined, any suitable correction technique can be used. In the authors' case, a curve fitting approach was used. The response at several different levels (typically five) was fitted for each emitter. The function used to model the response for emitter (*i*,*j*) is

$$f_{ij}(x) = A_{ij}(e^{\gamma_{ij}x} - 1), \tag{2}$$

where x is the incoming pixel value from the scene generator and A_{ij} and γ_{ij} are the fit coefficients. The desired response can be specified as

$$f_d(x) = A_d\left(e^{\gamma_d x} - 1\right). \tag{3}$$

The NUC hardware maps its input values to output values in the same domain. In other words, it is merely applying a function $g_{ij}(x)$ to the incoming pixel values. Thus the output of emitter (*i*,*j*) with NUC applied is $f_{ij}(g_{ij}(x))$. However, in order to have every emitter give the same output for the same input, the following condition must hold:

$$f_{ij}\left(g_{ij}(x)\right) = f_d(x)\ . \tag{4}$$

Solving Eq. 4 for $g_{ij}(x)$ yields

$$g_{ij}(x) = \frac{\ln\left[A_d\left(e^{\gamma_d x} - 1\right) + A_{ij}\right] - \ln A_{ij}}{\gamma_{ij}}\ . \tag{5}$$

Note that this is a continuous function independent of the hardware implementation. It can be approximated by the 16-segment, piecewise linear function of the CSA RNUC board. If other NUC hardware in the future implements arbitrary floating-point calculations instead of breakpoints, gains, and offsets, this correction function could be applied more precisely.

Once the correction is applied, ideally all of the emitters have the desired response, $f_d(x)$. The CSA RNUC board follows the NUC calculations with a global gain lookup table for linearization. This table maps input values to output values and can be considered to apply a function $h(x)$ to the output of the NUC stage. Therefore, for linearization the

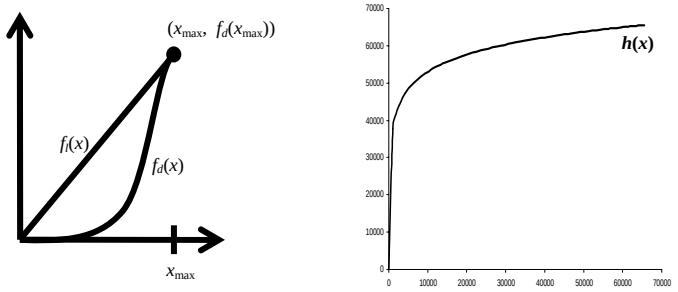

(a) (b)

Figure 3: Linearization.

following condition must hold:

$$f_d\left(h(x)\right) = f_l(x), \tag{6}$$

where $f_l(x)$ is the desired linear response. Fig. 3a shows that the endpoints of the desired linear response at zero and maximum input correspond to the endpoints of the actual response. We want zero input to map to zero output, and we want the maximum input value to map to the maximum output value, therefore the desired linear response function is

$$f_l(x) = m_l x = \frac{f_d\left(x_{\max}\right)}{x_{\max}} x. \tag{7}$$

This implies that $h(x)$ satisfies

$$A_d\left(e^{\gamma_d h(x)} - 1\right) = m_l x. \tag{8}$$

Solving this equation for $h(x)$ yields

$$h(x) = \frac{1}{\gamma_d} \ln\left(\frac{m_l}{A_d} x + 1\right). \tag{9}$$

This equation is used directly to generate the linearization table for the hardware, as shown in Fig. 3b.

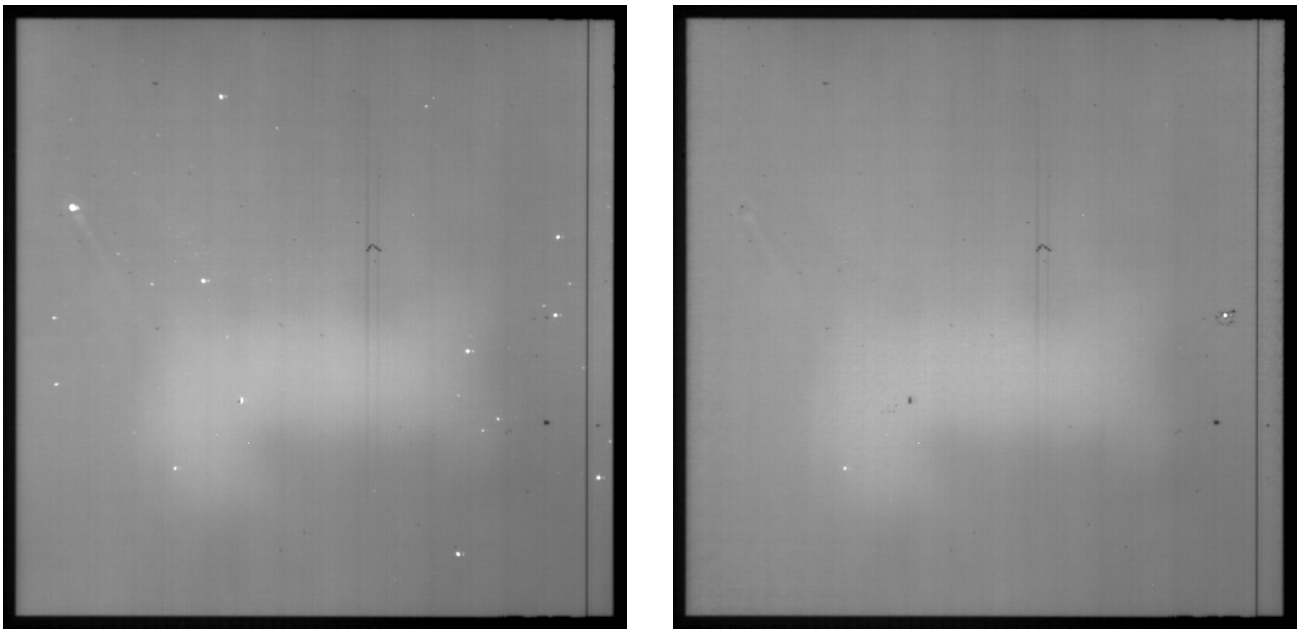

Raw Corrected

Figure 4: Raw and corrected flood array images.

4.2 Results

The ISGF method has been implemented with mixed results to date. Fig. 4 shows images of the raw and corrected array outputs. Note that the high spatial frequency non-uniformities are removed (except, of course, for dead and stuck-on emitters), but the low spatial frequency shading remains. This is a direct result of subtracting the background emitter image. This applies, in effect, a high-pass spatial filter to the measured emitter response, thus filtering out the large area spatial variations from the map. Some sensors are not sensitive to the low-frequency shading. For such, this technique

would be adequate. An advantage of using ISGF over regular sparse grid mapping is that stuck-on emitters do not corrupt the surrounding emitter data collection to the extent that they do with positive sparse grids. If a sensor is insensitive to shading, the shading should not be corrected. Doing so would reduce the dynamic range of the array unnecessarily.

5 NUC HARDWARE ISSUES

5.1 Inter-stage arithmetic overflow

There were several issues encountered related to the NUC electronics hardware and their operation. During evaluation of post-run imagery, it was noticed that some emitters would be driven to maximum output seemingly at random. It was determined that the emitters that did go to full output did so every time above a certain value. But, the value above which the emitter went full-scale varied from emitter to emitter. This effect was traced to the behavior of the CSA RNUC board when arithmetic overflows occur in the intermediate stage of the NUC gain and offset calculation. The product of the incoming pixel value and the gain is computed first. In the original version of AMRDEC's CSA RNUC board, if this product exceeded 65,535, the offset was not applied, even if it would have brought the result below 65,535. The final pixel value was left at the maximum value. This resulted in a "popcorn" effect in which increasing input levels caused pixels scattered about the scene to jump suddenly to the maximum output. This jump depended on the scene content as well as the NUC gain for each emitter. Fig. 5 shows in black the emitters that are stuck at maximum output because of this effect for three different input levels. As the input level increases, more of the emitters overflow in the intermediate stage. Not all overflow because many have gains less than one. Dynetics, Inc., which absorbed CSA's Simulator Group, was contracted to reprogram the RNUC board firmware to apply the offset calculation regardless of the intermediate value. With this change, the phenomenon no longer appears.

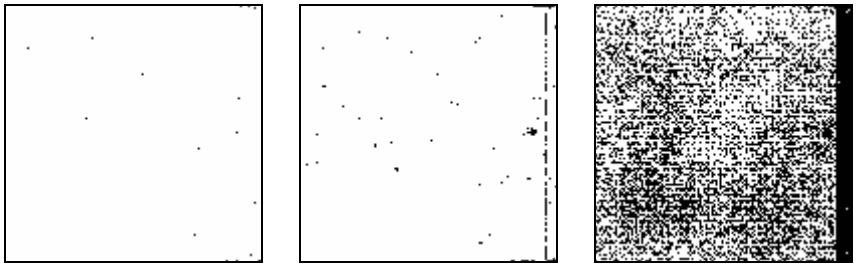

40,000 60,000 65,535

Figure 5: Inter-stage arithmetic overflow

5.2 Need for greater than unity gain

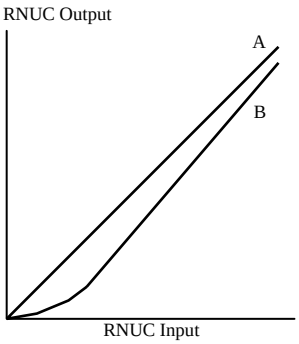

Figure 6: NUC functions

AMRDEC purchased a new model of projector electronics similar to the CSA electronics in architecture. However, the new electronics did not support NUC gains greater than one. The idea being that if a maximum output drive level is chosen such that virtually all emitters can achieve this level, all the NUC gains will be less than or equal to one. However, this is not the case. The RNUC section implements via its gains and offsets a 16-segment, piecewise linear curve that transforms each emitter's input to the DACs such that all emitters produce the same output for the same input. Two such emitter correction functions are shown in Fig. 2. Typically, the desired emitter response is chosen so that almost all the emitters can achieve that response over the entire input range. Curve A is a straight line with a gain of one and offset of zero. This would apply to an emitter that already has the desired response. Its output to the DACs simply equals the input. Curve B is a typical (though exaggerated) curve for an overachieving emitter. The input is reduced so that the output counts are not so high. However, note that for this emitter, not quite as much reduction is required at the higher input settings. This means the slope of B towards the end of the range is greater than

one, which implies that the gains in the NUC table are greater than one—even though the emitter output is always being reduced. Therefore, greater than unity gains are required even if the vast majority of emitters are attenuated.

5.3 Optimal distribution of piecewise linear knots

Given that the RNUC board models the correction functions with 16-bit, piecewise linear curves, a question that arises is "What is the optimal distribution of knots (breakpoints)?" In Fig. 7a we see eight correction functions plotted along with the line $y = x$, which represents the "correction" function for an emitter that has the desired response. We observe that all of these emitters are overachievers, because their inputs are being reduced relative to the desired response line. These are the curves that the RNUC board will model with a 16-segment, piecewise linear curve. Intuitively, we wish to concentrate the knots in the region of highest curvature. We can see from Fig. 7a that the maximum curvature appears to occur in the lower portion of the curves. If we subtract the ideal $y = x$ line from each of the curves, we get the deviation of each curve from the ideal linear curve. Fig. 7b shows this result. The curvature is much more apparent in this figure. We can see that the region of strongest curvature occurs from about 5,000 to 35,000 counts on the input scale. Distributing the majority of our knots in this range will greatly improve the fidelity of the 16-segment, piecewise linear model. Better results could be achieved by using a non-linear least squares curve fitting algorithm to fit a typical correction curve, where the locations of the knots are used as parameters to the curve fit in addition to the slope and offset of each segment. This distribution of knots would then be used to set the breakpoint table in the hardware.

It can also be noted in Fig. 7b that segments of the curves with slopes greater than zero correspond to segments in the correction functions that have slopes greater than one in Fig. 7a. Even though all of these emitters are being attenuated, several of them require greater than unity gains for correction.

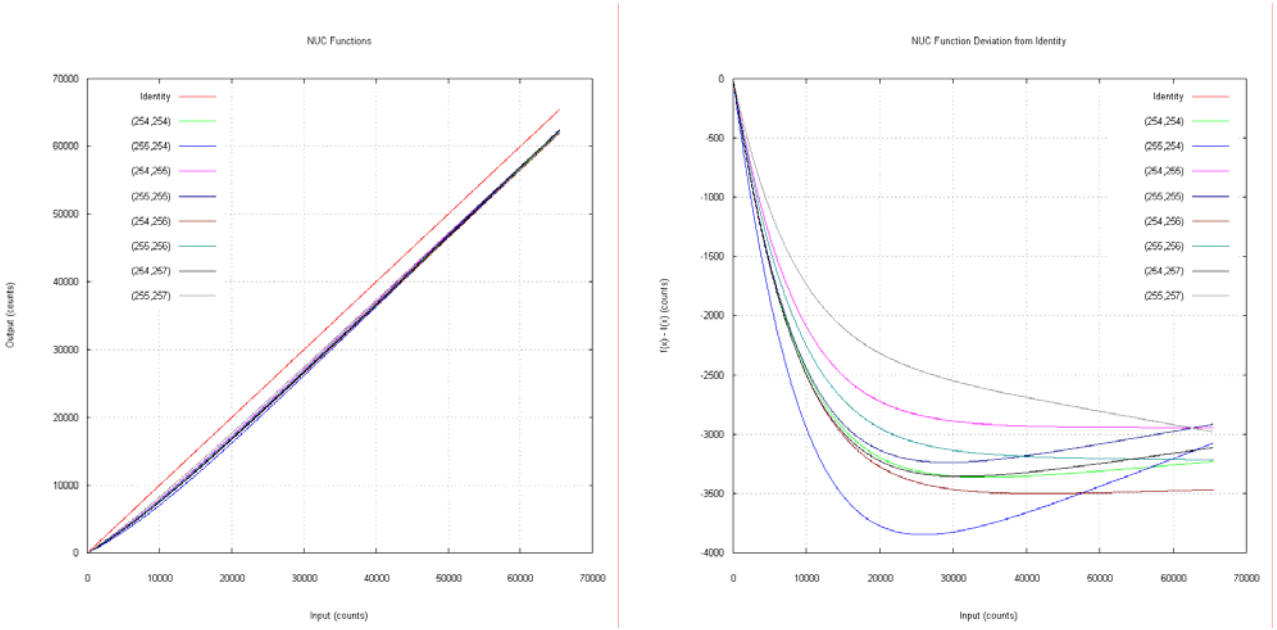

(a) (b)

Figure 7: NUC function curvature.

5.4 Substrate heating stabilization

When a flood scene is applied to the array, it takes several minutes for the output of the array to stabilize. Fig. 8 shows the array output as a flood input is applied to it. During ISGF mapping, the AMRDEC software waits an operator-specified amount of time after turning on the array before data collection begins (typically six minutes). This implies that in the testing environment, the array should have the scene applied to it long enough before testing begins for stabilization to occur.

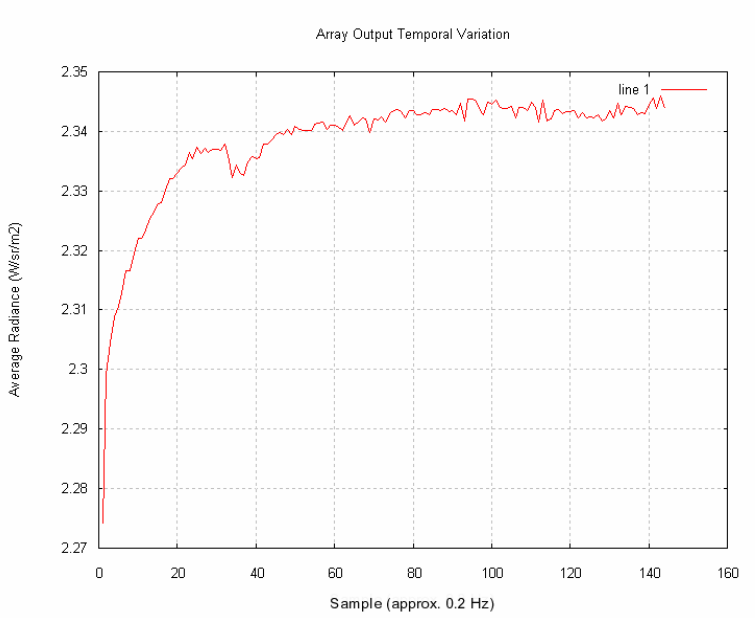

Figure 8: Array output stabilization

5.5 Scene generation server for efficient mapping

In order to load a new image into the CSA hardware via the VME bus, the boards have to be placed in a non-running state. This turns off the array output allowing the array to cool from its stabilized value. After a new grid image was loaded, it was found that a delay of several seconds was necessary before collecting the data for that grid in order to allow the array to re-stabilize. This delay added a considerable amount of time over the course of the entire mapping process.

To remove this stabilization delay, a simple image generator server was developed for a Silicon Graphics Onyx2. The NUC software communicates with the server via UDP/IP packets and requests the necessary grid and flood imagery during the mapping process. The output of the Onyx2 drives the emitter array via the DD02 digital video port. Changes in the command image take place from one frame to the next without the need to turn off the array for a VME transfer.

6 CONCLUSIONS

We have seen that a simple piecewise linear interpolation of the response of each emitter provides a computationally efficient and accurate means for NUC sensor correction and calibration. A fitting function for emitter images has been presented that conforms better to the shape of the typical emitter image than does a Gaussian curve and is effective for replacement of dead detectors in tiled configuration. Inverse sparse grid flood mapping shows promise for providing for flood mapping without the complexity of dealing with coupled outputs. The method is suitable for use with curve fitting or iterative NUC processes. Currently, only high spatial frequency non-uniformities are accounted for, though it is expected that low frequency shading will be dealt with in future development. It was shown that NUC hardware must not stop the gain and offset computations if the intermediate product overflows and that greater than unity gains are required, even if emitters are only attenuated. We have seen that the CSA RNUC board models NUC functions by a 16-segment, piecewise linear curve and that the knots of the curve can be distributed so as to minimize modeling error. It was observed that, for flood mapping, substrate heating can require considerable pre-heating time for the array to stabilize before mapping or scene projection and that a simple grid image server can eliminate the cooling between array command changes.

7 ACKNOWLEGEMENTS

The authors wish to acknowledge the assistance and suggestions of Danny Saylor of Optical Sciences Corporation in the development of the NUC station.

REFERENCES

1. E. Olson, R. Murrer, “Non-uniformity correction of a resistor array infrared scene projector”, *Technologies for Synthetic Environments: Hardware-in-the-Loop Testing IV*, R. Murrer, **3697**, 403-413, SPIE, Bellingham, 1999.
2. E. Olson, T. Bergin, “Non-uniformity correction using a flood technique and 1:1 mapping”, *Technologies for Synthetic Environments: Hardware-in-the-Loop Testing VI*, R. Murrer, **4366**, 501-509, SPIE, Bellingham, 2001.
3. D. Moore, C. Stanek, C. Peterson, “Performance limitations and practical considerations associated with resistive emitter array non-uniformity correction (NUC)”, *Technologies for Synthetic Environments: Hardware-in-the-Loop Testing II*, R. Murrer, **3084**, 120-130, SPIE, Bellingham, 1997.
4. D. Flynn, S. Marlow, E. Olson, “Projector nonuniformity and spatial effects modeling”, *Technologies for Synthetic Environments: Hardware-in-the-Loop Testing III*, R. Murrer, **3368**, 189-201, SPIE, Bellingham, 1998.