Design Considerations for a High-Temperature, High-Dynamic Range IRSP

Joe LaVeigne[a], Breck Sieglinger[b]

[a]Santa Barbara Infrared, 30 S Calle Cesar Chavez, Suite C, Santa Barbara, CA 93103
[b]McAuley-Brown, Shalimar, FL

ABSTRACT

Achieving very high apparent temperatures is a persistent goal in infrared scene projector (IRSP) design. Several programs are currently under way to develop technologies for producing high apparent temperatures. Producing a useful system capable of reproducing high fidelity scenes across a large range of apparent temperatures requires more than just a high temperature source. The entire scene projection system must support the extended dynamic range of the desired scenarios. Supporting this extended range places requirements on the rest of the system. System resolution and non-uniformity correction (NUC) are two areas of concern in the development of a high dynamic range IRSP. We report the results of some initial investigations into the resolution required for acceptable system performance and the effects of moving to a higher dynamic range may put on existing NUC procedures.

Keywords: Infrared, Scene projection, Non-uniformity Correction (NUC), IRSP, HWIL, Ultra high temperature (UHT), MIRAGE

1. INTRODUCTION

Among other qualities, high radiance or apparent temperature is desired for many IR scene projector applications. Current state of the art resistive arrays can achieve apparent temperatures up to 700K in the 3-5 um band[1]. New technologies are currently under development with the goal achieve temperatures in excess of 2000K using high temperature materials for resistive arrays[2] or other sources using narrow band emission. While there has been considerable effort applied to the development of these technologies with the basic goal of being capable of producing very high apparent temperatures, there has been less of a focus on the system level aspects of a high dynamic range projector. Although producing higher temperatures is a worthy goal, a practical IRSP system must also produce accurate radiance with high fidelity at low apparent temperatures as well. Here we present the results of some initial investigations into the development of a high dynamic range IRSP at the system level.

The initial studies described below focus on two areas: System level resolution and Non-uniformity correction (NUC). Resolution becomes more of a challenge as the maximum apparent temperature of the system increases. This is primarily due to the nonlinear relationship between apparent temperature and radiance. A system with a maximum apparent temperature of 2000K has a MWIR output radiance 40 times higher than one with a 700K maximum apparent temperature. The 40 times larger radiance at 2000K means a system that can achieve that radiance will require a higher fidelity by that same factor in order to simulate low temperature objects with the same absolute resolution. The resolution of the system is set by that component or algorithm which has the lowest resolution and may be dependent on the radiance being commanded. The discussions below will address these issues for a resistive array and provide some commentary on applications based on other technologies.

Non-uniformity correction is another important portion of system design that may be affected by achieving higher radiance levels. In order to achieve higher apparent temperatures with a resistive array, more power must be applied to

the emitters. This is typically achieved by increasing the current through the unit cell of the emitter. In order to accomplish this without requiring a higher voltage input to the unit cell, the gain of the current control of the unit cell is increased. This increased gain in the transfer function of the unit cell amplifies any non-uniformity in the individual pixels. In particular, these effects are amplified at low radiance values near the input voltage at which the unit cell begins driving current through the resistor. Current scene projection systems use a 16 point piece-wise linear correction for the NUC. The discussions below will address whether the current NUC process will be acceptable for a high dynamic range system.

2. RESOLUTION

The apparent temperature of an object is defined as the temperature of a blackbody that produces the equivalent integrated radiance over the band of interest. The Planck Function shown in Equation 1 describes the radiance of a blackbody. Consider a system that is linear in radiance with the requirement or have a minimum resolution of 0.1K at an apparent temperature of 300K. Such a system would be capable of producing a reasonable simulation of a typical ambient environment, though it would be far from the ~15mK NETD of a typical mid-wave infrared (MWIR) imager. A 0.1K step at 300K is equivalent to a radiance of 6.8x10-7 W/(cm^{2}sr). A 16 bit system with a 2000K maximum apparent temperature would require have a step size of nearly 10K at 300K. In order to achieve 0.1K resolution at 300K, the system would require approximately 8 million steps, or nearly 23 bits of resolution. Current scene projectors are designed as 16 bit systems and typically operate linear in radiance. Without changes, these systems would not be acceptable for a high apparent temperature array.

$$L(\lambda, T) = \frac{2hc^2}{\lambda^5(e^{hc/\lambda kT} - 1)}\, d\lambda \qquad [1]$$

The minimum incremental step in radiance of a system is limited by the step sizes of its subsystems. In the case of an IRSP, there are several subsystems to be considered when determining system resolution. Scenarios are produced with a scene generator. The scene data is then transferred to the command and control electronics (C&CE) where it is processed before being transmitted to the digital emitter engine (DEE) and the array. It is insightful to consider the resolution of each of the subsystems independently. Modern scene generator packages use floating point calculations, so resolution on the scene generator side is not an issue. Most current systems use a digital transfer standard such as DVI or DVP2 to transfer the scene data to the C&CE, so the transfer standard is not a limiting factor either. The current C&CEs including SBIR's MIRAGE product line and the KHILS PACE systems operate on a 16 bit fixed point basis. Some of the processing including translation and rotation compensation require the system to operate in linear radiance in order to in to do some spatial calculations. The final step in the processing for the C&CE is to convert the data stream into raw counts sent out to the array. The array has a non-linear transfer function between commanded drive and output radiance. The non-linear response is due to both the quadratic dependence of the power dissipated by the emitter with a constant current unit cell and the non-linear radiance output of a pixel as a function of physical temperature. This non-linear response provides improved resolution at low radiance levels. Figure 1 shows a plot of the measured MWIR response of a MIRAGE-XL resistive array as function of raw drive commanded to the array. Figure 2 plots the thermal resolution as a function of apparent temperature for that same system operated with a 16 bit linearization table. For the MIRAGE-XL system, the native resolution is finer than the linear resolution below 400K. Thus, for a linearized system, the resolution below 400K would be set by the linearization table. Above approximately 400K, the native resolution of the system is the limiting factor.

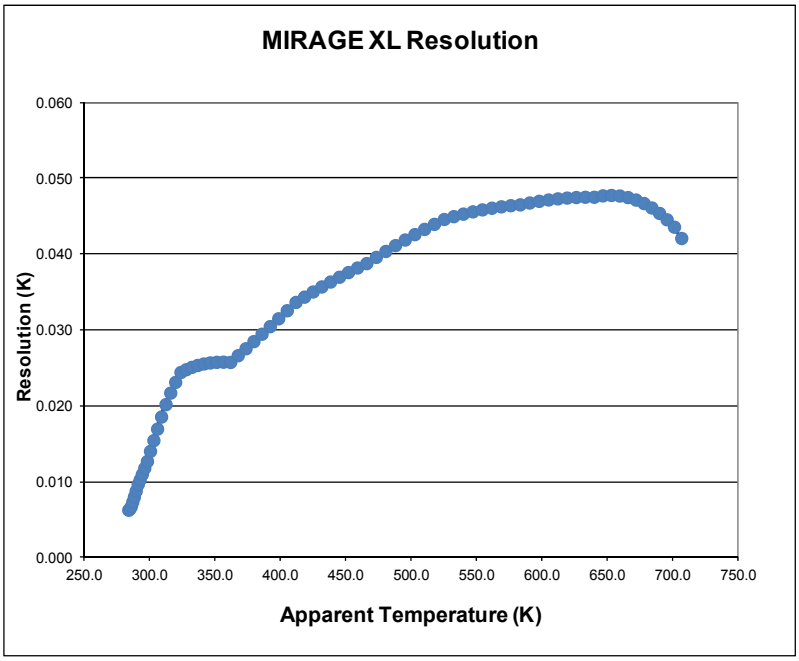

Figure 1 Native Resolution of MIRAGE-XL Array. The plot above shows measured resolution for a typical MIRAGE-XL array. Resolution is given in apparent temperature as a function of apparent temperature. The system has a 16 bit input, but the true resolution is approximately 14 bit due to overlap in the lower and higher bytes of the DACs. The shape is determined by the non-linear transfer function between commanded current and output radiance.

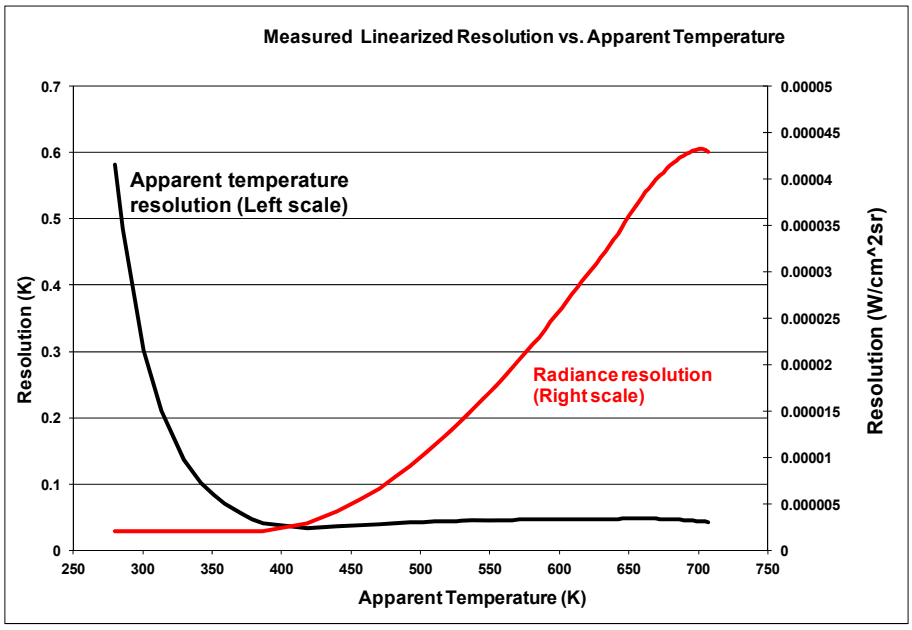

Figure 2. Plots of the resolution of a typical MIRAGE-XL system as a function of apparent temperature after linearization for operation in the MWIR band. Below approximately 400K, the resolution is limited by the system linearization to 16 bit resolution. Above 400K, the system resolution is limited by that of the array.

For a system with a 2000K MWIR maximum apparent temperature, the native resolution will become coarser assuming a similar digital to analog converter (DAC) with an effective resolution of 14 to 15 bits is used. The predicted resolution for such a system is shown in Figure 3. Though the step size is larger than existing arrays, it does not exceed 0.1K until nearly 500K. Based on these predictions, a 14 bit native resolution for a resistive array would be adequate depicting low radiance scenes. Figure 3 also contains a plot of the resolution of the same system after a 16 bit linearization has been applied. In this case, the step size near ambient temperatures increases to 10K or more. Such a coarse resolution is not acceptable for low radiance scenes. In order to operate with a response that is linear in radiance, a different representation must be used in the system. This will lead to new firmware and potentially new hardware being developed to support the new representation. A 24 bit fixed point number would be just acceptable for a MWIR projector with a 2000K maximum temperature. A floating point representation would also be acceptable. Given the flexibility of the floating point representation for future growth, it is the recommended format for the next generation for control electronics.

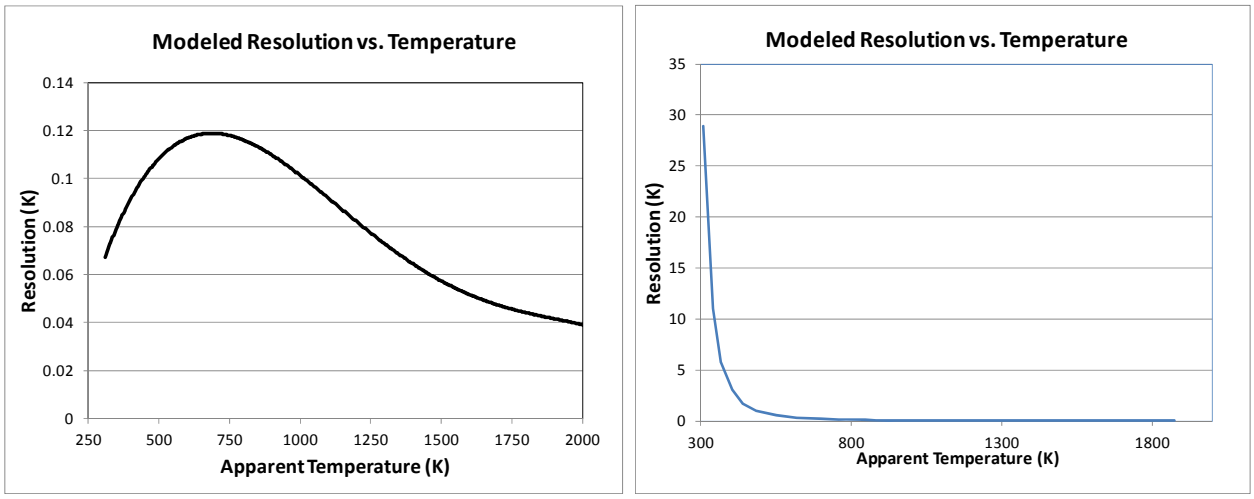

Figure 3. Modeled resolution of a high radiance scene projector. The plot on the right is the native resolution with effective 15 bit input to the RIIC. The plot on the left is the resolution assuming the system operates linearly in radiance with 16 bit resolution. Despite having an effective 15 bit input to the RIIC, the native resolution in apparent temperature is acceptable, due to the non-linear current to radiance transfer function of the unit cell.

The resolution issues described above apply to any emitter array. Consider a light emitting diode (LED) based array. In that case the native radiance versus drive function is much closer to being linear than the resistive arrays. If the native bit depth of a system based on a LED emitter array is 16 bits, then that will set the limit on the system resolution. For high temperature LED arrays, a higher resolution circuit at low radiance levels will be required in order to simulate low radiance scenes. The issues with digital micro-mirror devices (DMDs) are related. For a DMD to produce adequate resolution it must be capable of flipping back and forth very rapidly. For a 23 bit system operating at 400 Hz, in order to display a single bit of radiance, the mirrors would have to switch at over 3 GHz. This is far beyond the maximum mirror frequency of nearly 100MHz for DMD devices currently in use.

3. NON-UNIFORMITY CORRECTION

Non-uniformity correction is an important part of the calibration of an IRSP. For a high temperature array, NUC must cover a larger radiance range than that of arrays currently in use. At higher radiance levels, the NUC must accommodate a transition from a regime where the primary loss mechanism in the pixels is conductance down the legs to a regime where the primary loss mechanism is radiation. In addition, the higher gain in the unit cell transfer

function makes the emitters more sensitive to variation in the drive circuit, especially the voltage where the unit cell begins to drive current through the emitter.

3.1 Statistical Model

In order to study the possible implications of a high temperature system prior to its fabrication a statistical model of pixel response was employed. Output radiance of a given pixel depends on a number of factors including the resistance and emissivity of the emitter, the conductance of the legs, and the voltage to current transfer function of the individual unit cell. A simplified statistical pixel model was produced to help guide the study. This model allowed for variation in pixel emissivity, resistance, leg conductance and unit cell transfer function. The emissivity, resistance and leg conductance were given a normal distribution around nominal levels based on measured data and best estimates for process variability. The unit cell transfer function was modeled as a linear function with variable slope and offset. The distributions of the slope and offset were derived from models of RIIC performance. A Monte Carlo simulation was performed where numerous pixels were simulated while varying the parameters described above within their expected distributions. Predicted radiance as a function of commanded drive was then calculated for each pixel. In order to validate the model, the analysis was performed on an existing emitter array currently in production and those results compared to measured response curves for a typical array of that variety.

The state of the art resistive emitter arrays produced by SBIR are all based around constant current unit cells. The physical temperature attained by a pixel at a particular drive is determined by the power into the pixel, and the losses due to conduction to the substrate and radiation into the ambient environment. The power into the pixel is a function of the resistance of the pixel and the current supplied by the unit cell:

$$P_{in} = I^2 R \qquad [2]$$

The loss due to conductance down the leg is a function of the substrate temperature, the physical temperature of the pixel:

$$P_{leg} = 2 * \mathrm{K} * (T_{emitt} - T_{sub}) \qquad [3]$$

Where K is the conductance of one leg and T_{emitt} and T_{sub} are the temperatures of the emitter and substrate respectively. Losses due to radiation can be calculated by integrating the product of the Planck Function for the physical temperature of the pixel and the emissivity of the pixel, taking into account the fill factor of the active area of the pixel body, taking into account the radiation from the ambient environment. Because the conductive losses are much larger than the radiative losses near normal operational temperatures, losses due to radiation do not become significant until physical temperatures are typically several hundred K above the normal operational temperature of 273K. Because of this, any contribution to radiative losses from the temperature of the ambient environment can be considered insignificant. Thus the radiative losses can be given as:

$$L_{\mathrm{total}} = \pi * \eta \int_0^\infty \frac{2hc^2}{\lambda^5 (e^{hc/\lambda k T_{emitt}} - 1)} * \epsilon(\lambda)\, d\lambda \qquad [4]$$

Where η is the fill factor of the pixel. The steady-state physical temperature of a pixel occurs when the sum of the loss terms is balanced by the input power. Using Equations 2-4 the physical temperature of a pixel can be numerically derived. Once the physical temperature of a pixel is known, it can be used to calculate the apparent temperature of the pixel in a given band. The radiance of the pixel is calculated by integrating Equation 1 over the band of interest for the radiance of the pixel, the substrate and the reflected radiance of the ambient environment:

$$L = \int_{\lambda 1}^{\lambda 2} [\eta L(\lambda, T_{emitt})\, \epsilon_{emitt} + (1-\eta) L(\lambda, T_{sub}) \epsilon_{sub} + (1-\eta) L(\lambda, T_{amb})(1-\epsilon_{sub}) + \eta L(\lambda, T_{amb})(1-\epsilon_{emitt})] d\lambda \quad [5]$$

Where T_{amb} is the temperature of the ambient environment, η is the fill factor of the pixel and ϵ_{emitt} and ϵ_{sub} are the wavelength dependent emissivities of the emitter and substrate respectively. To simplify the model, a constant emissivity over the band of interest was assumed. Using the above equations, a series of pixels with randomly assigned total emissivity, in-band emissivity, resistance and leg conductance were modeled and a response curve for each pixel was generated. An average response curve was then calculated from the sample of pixel responses produced in the simulation. This curve was used to generate a linearization table. Once this table was constructed, the ratio of pixel radiance to average radiance was calculated for each pixel in the simulation. Figure 4 shows the relative response of 30 emitters from the simulation of the existing system. Figure 5 shows the measured relative response for a sample of pixels in a fielded system. The range of simulated responses is quite similar in shape and magnitude to the measures response of the real system. This similarity gives some confidence that the modeled results are a reasonable approximation.

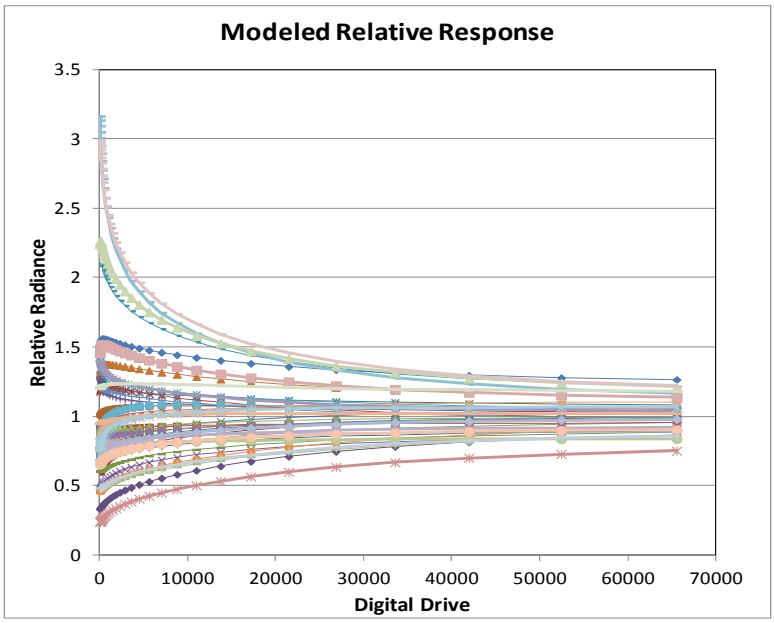

Figure 4. Simulation of existing system. The plots above graph the relative response of the modeled of the existing IRSP system. The overall character of the measured curves in Figure4 is reproduced in the model.

Based on these results the model was used to simulate the response of a high temperature array. The results are shown in Figure 6. The high temperature array exhibits a larger variance at low radiance levels. This is primarily due to the higher slope in the voltage to current transfer function of the unit cell combined with the variation in the voltage at which the individual unit cells begin to drive current through the emitter.

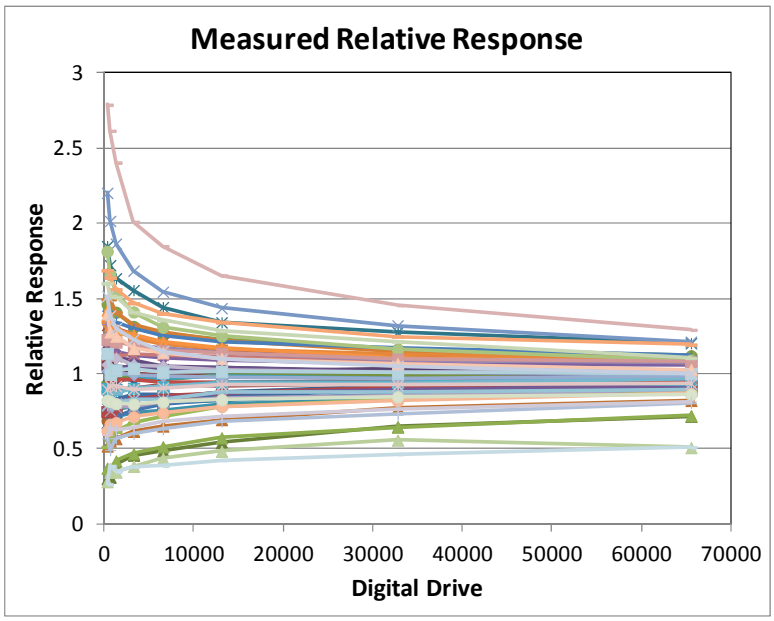

Figure 5 Relative measured response of an existing IRSP system. The curves above show the relative response of an uncorrected array. The distribution is typical for resistive arrays currently in use.

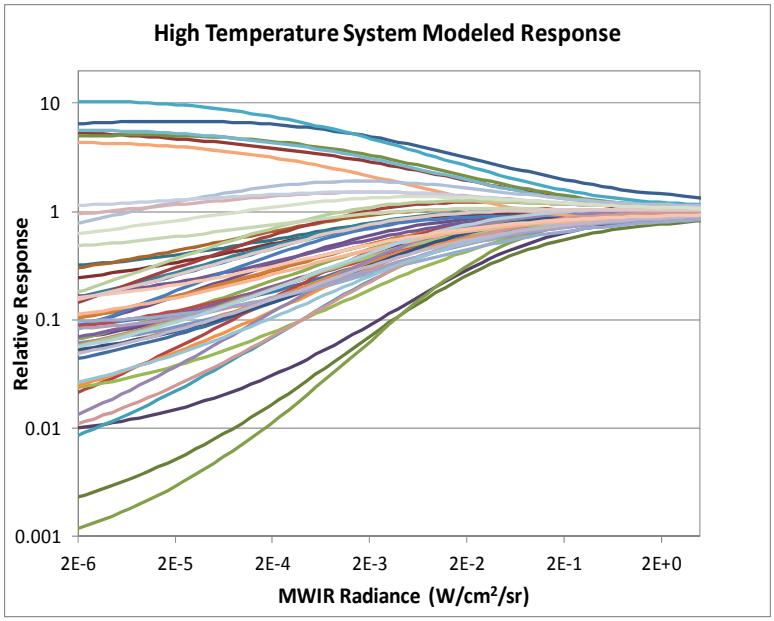

Figure 6 Modeled response of 50 High temperature pixels as a function of radiance. The modeled high temperature pixels show more variation at low radiance than the existing pixels. This increased variation will make NUC of the high temperature arrays more difficult at low radiance values compared to existing pixels.

3.2 Overview of Current Generation NUC

The NUC process for resistive arrays has a simple underlying idea, but presents several challenges to implement, even with current-generation pixels[3-5]. The NUC consists of finding the voltage, or the corresponding digital value, for each pixel, that causes it to produce a desired output. The "NUC tables" that perform the correction are accessed in real-time, which is referred to as the RNUC (real-time NUC) implementation.

Because of the nonlinear radiance-versus-voltage response of the pixels, the projector NUC has some aspects which may be counter-intuitive. First, if some pixel output is, say, "too low" by a factor x, the NUC tables do not simply scale the command to this emitter by the factor 1/x. This notion would apply if the system components downstream from the NUC tables were linear. Instead the NUC tables must produce an output that causes the nonlinear system (emitter radiance versus voltage) to generate the desired radiance. That is, the NUC table finds the input voltage or corresponding digital value that will produce the desired output radiance. In fact, over some range of the a pixel's response curves, very small changes in the drive voltage can produce very large (fractional) changes in the output radiance, depending on the particular voltage or output being considered.

Figure 7 shows the simulated response curves for two pixels – that is, their output radiance versus drive voltage. If these two pixels are commanded to produce the radiance corresponding to the horizontal black line below, the NUC tables would have to "choose" the corresponding drive voltages as illustrated. Notice that the ratio of the response curve values is not related in an obvious way to the drive voltages required to produce this output from the two different pixels. The response curve for each pixel, as shown here, is the essential and minimum information needed to perform the NUC.

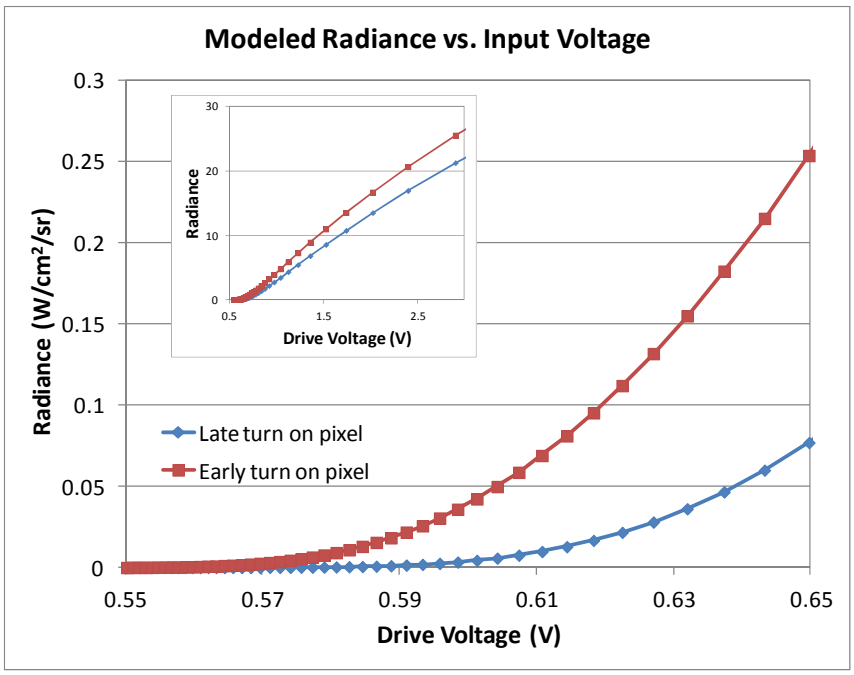

Figure 7. Two simulated response curves without linearization. The main plot shows the output radiance as the unit cell is just beginning to respond to the drive voltage. The inset shows the entire range.

Because of previously mentioned differences from pixel to pixel, the response curve for every pixel must be determined by direct measurement. The measurements must usually be performed using the spectral response and operating conditions very similar to the desired sensor test conditions.

One challenge of this step is due to the large number of pixels that must be measured at several different drive voltages. A careful strategy is required to limit the time and data collected. A second challenge is related to the precision of these measurements. Unavoidable errors in making these measurements, especially at very low drive voltages impose one limit on the final NUC performance.

Given the pixel outputs measured for several drive voltages, a continuous curve must be constructed so that a drive value can be associated with every commanded radiance, recognizing that it is not practical to measure the output for each of the different voltages produced by the DACs. This interpolation or smooth fitting of the measurements can be done using one of a variety of methods: linear interpolation, piecewise polynomials, global curve fit of some sort, to name a few. The choice of method used for this interpolation depends on factors such as the number of response values measured for each pixel, the underlying shape of the response curve, and the noise in each measurement.

Finally, given the interpolated response curve, the RNUC tables must be implemented in the C&CE. With current-generation hardware and firmware, a look up table (LUT) for each pixel is constructed using up to 16 linear segments, each of which is represented by a "gain" and "offset" value. The gains and offsets are stored as fixed-point numbers that limit the range and precision of the values. Again we mention that the gains and offsets for any segment of the RNUC curves do not represent directly how much too high or low the pixel output radiance is, but rather how to modify the input to the pixel's nonlinear response curve to achieve the desired output.

How will the current process and implementation work with high-output pixels? First consider measuring the pixel outputs for several drive voltages. Key to the number of measurements required is the underlying shape of the nominal response curve, and the uncorrected variations within the population of pixels. Simulations suggest that the variation in output within the high-output emitter population may be considerably larger than the current resistive arrays, particularly in the low-output regime. In this case, additional low-output measurements might be required of all pixels to ensure that all curves are well sampled in this regime.

Also, for very-high-temperature pixels, the NUC measurements must be made over a very large radiance range, which may require combining data from various sensor settings (e.g. long integration time for low radiance, shorter integration time for high radiance), or the use of neutral density filters in the optical path. These are all factors that complicate the construction of a multipoint response curve for every pixel in the array.

Next, consider interpolating the set of response measurements to produce the full response curves. Here again, the variety of response curves will guide the choice of methods to interpolate the set of measured values. In this step, simulations have not revealed whether the response curves are likely to be inherently more complex, and require more complex procedures for interpolating the response values.

Finally, the implementation of the real-time RNUC tables is expected to be particularly challenging with high-output pixels. The present tables are constructed using gain and offset values for each segment of each curve. The segment slope or gain values are stored as 16-bit fixed-point values between 0 and 4, in 2^16 equal steps. The offset values are also 16-bit quantities and can represent numbers between minus 65536 and 65534 in steps of 2. Simulations suggest that, for high-temperature resistive arrays, the "early turn on" pixels may be tens of times brighter at low drive voltages than the "late turn on" pixels. In contrast, the resistive arrays currently in service have a much smaller spread of outputs at low voltages, usually within a factor of 2 or 3. Although the connection is not direct there is a relationship between the magnitude of the gain values, and the inherent spread in the pixel outputs which can be corrected. Very roughly, it will not be possible to "correct" pixel outputs over the large range expected for high-output pixels using fixed-point gains in the range from zero to four. The following figure shows the same pixel response curves plotted previously, but this time using a linearizing transformation. In this plot, which shows the low-output regime, the relationship between these response curves is clear. The late-turn-on pixel will require a significant "gain" from the NUC table to match the higher output of the early-turn-on pixel.

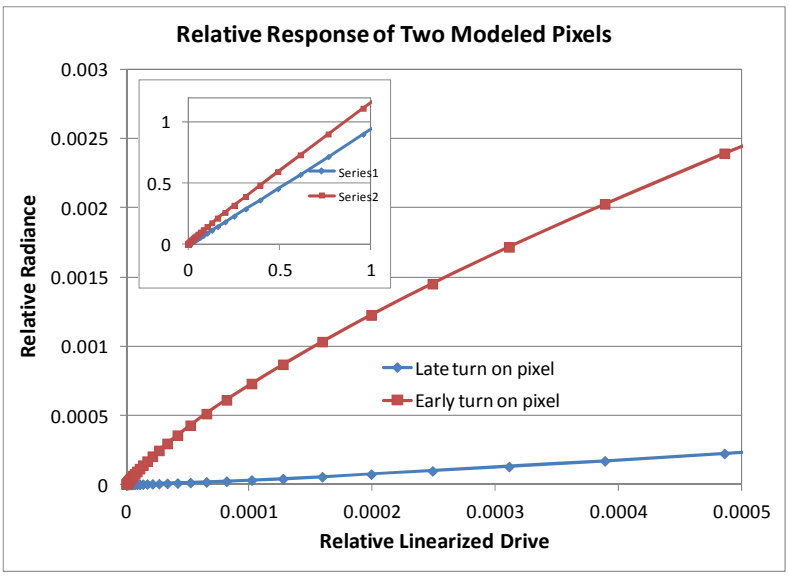

Figure 8. The same two response curves as in Figure 7 after linearization.

The actual situation is not so clear, as the C&CE contains several other "global" lookup tables which can be used to reshape the family of response curves. While this does not alter the spread of values, it does play a role in how the RNUC tables relate to drive voltages. The variety of strategies for implementing the C&CE tables, including the RNUC, is quite large, with options that greatly affect the sensitivity and range of output values (and the NUC performance).

Another consideration with the RNUC tables is more subtle, and is related to whether the C&CE global tables are configured to linearize the nominal response curves. This allows the RNUC tables to use the piecewise-linear segments to approximate curves that are nearly linear. This greatly reduces errors associated with the RNUC process over much of the range of the response curve. It also reduces the penalty incurred and planning difficulties associated with choosing the "breakpoint" values over which the RNUC segments are defined.

Within the scope of this paper, we note that the C&CE tables, including RNUC, will not provide suitable NUC performance without changing at least one of the following aspects of the current implementation:

- The range and numerical precision of the gains and offsets
- The common configuration whereby the after-NUC LUT implements a global "linearization" table
- The number of breakpoint segments
- The method of linear interpolation between the breakpoints

Further, it is likely that high-output pixels will require additional effort in the form of more response measurements (sometimes called array maps), and additional manipulation of the camera data to account for different setting or optical filters, as mentioned above.

4. SUMMARY/CONCLUSIONS

An initial study was performed to investigate the system effects of moving from the current state-of-the-art scene projectors to a projector with significantly higher apparent temperatures and a wider dynamic range. The primary conclusions from the study were:

- 16-bit linear drive or inputs to a “high-temperature” system will not allow for acceptable resolution at low output
- Similarly, an inherently linear system driven with 16-bit evenly-spaced inputs will not provide acceptable resolution at low output
- The present C&CE NUC implementation will not allow for acceptable NUC performance if the pixel-to-pixel variations are as large as early simulations suggest
- New numerical formats (probably floating point) and LUT representations or transformations will be required to achieve an acceptable NUC for high-temperature projectors that must also provide good performance at near-ambient temperatures.

These considerations and others should be carefully reviewed as the control systems are designed for high temperature arrays currently under development. It is noted that some of the required changes will extend beyond the design of the control electronics. Changes to improve resolution must be implemented in the entire system, including moving from the existing 16 bit radiance interface, which is most common, to some other interface such as a floating point representation in order to achieve the necessary resolution at low apparent temperatures. Changes in the RNUC algorithm may require alteration of the global look-up tables typically used and will likely require additional low radiance measurements to address the higher variability expected in the new systems.

5. ACKNOWLEDGEMENTS

The authors would like to thank J.P. Curzan, from Nova Sensors, and Eileen Cutaia, from Raytheon, for providing helpful discussions and information regarding the transfer functions of the unit cell designs considered above.

6. REFERENCES

[1] Sparkman, Kevin, LaVeigne, Joe, McHugh, Steve, Lannon, John, Goodwin, Scott, "Ultra High Temperature (UHT) Infrared Scene Projector System Development Status", Proc. SPIE 8356, (2012)

[2] Oleson, Jim, James, Jay, LaVeigne, Joe, Sparkman, Kevin, Matis, Greg, McHugh, Steve, Lannon, John, Goodwin, Scott, Huffman, Alan, Solomon, Steve, Bryant, Paul, "Large format resistive array (LFRA) infrared scene projector (IRSP) performance and production status", Proc. SPIE 6208, (2006)

[3] Sieglinger, Breck A., Norman, James D., Meshell, William M., Flynn, David S., Thompson, Rhoe A. and Goldsmith, George C. II, "Array nonuniformity correction: new procedures designed for difficult measurement conditions", Proc. SPIE 5092, (2003)

[4] Sisko, R. B., Thompson, Rhoe A., Marlow, Steven A., Sieglinger, Breck A. , “Resistor array waveband nonuniformity measurements and RNUC band converter” , Proc. SPIE 6208, (2006)

[5] LaVeigne, Joe, Franks, Greg, Sparkman, Kevin, Prewarski, Marcus, Nehring, Brian, and McHugh, Steve, “LWIR NUC Using an Uncooled Microbolometer Camera”, Proc SPIE 7663, (2010)